

Isilon OneFS

CLI Command Reference

8.2.1

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

1 Introduction to this guide.....	18
About this guide.....	18
Where to go for support.....	18
Remote Support.....	18
2 OneFS isi commands A through C.....	19
isi antivirus policies create.....	22
isi antivirus policies delete.....	24
isi antivirus policies list.....	24
isi antivirus policies modify.....	25
isi antivirus policies start.....	27
isi antivirus policies view.....	27
isi antivirus quarantine.....	27
isi antivirus release.....	28
isi antivirus reports delete.....	28
isi antivirus reports scans list.....	29
isi antivirus reports scans view.....	30
isi antivirus reports threats list.....	30
isi antivirus reports threats view.....	31
isi antivirus scan.....	31
isi antivirus servers create.....	32
isi antivirus servers delete.....	32
isi antivirus servers list.....	33
isi antivirus servers modify.....	33
isi antivirus servers view.....	34
isi antivirus settings modify.....	34
isi antivirus settings view.....	37
isi antivirus status.....	37
isi audit progress global view.....	37
isi audit progress view	37
isi audit settings global modify.....	38
isi audit settings global view.....	40
isi audit settings modify.....	40
isi audit settings view.....	42
isi audit topics list.....	43
isi audit topics modify.....	44
isi audit topics view.....	44
isi auth access.....	44
isi auth ads create.....	45
isi auth ads delete.....	48
isi auth ads list.....	49
isi auth ads modify.....	49
isi auth ads spn check.....	53
isi auth ads spn create.....	54

isi auth ads spn delete.....	54
isi auth ads spn fix.....	55
isi auth ads spn list.....	55
isi auth ads trusts controllers list.....	56
isi auth ads trusts list.....	56
isi auth ads view.....	57
isi auth duo modify.....	57
isi auth error.....	58
isi auth file create.....	58
isi auth file delete.....	61
isi auth file list.....	61
isi auth file modify.....	62
isi auth file view.....	67
isi auth groups create.....	67
isi auth groups delete.....	68
isi auth groups flush.....	69
isi auth groups list.....	69
isi auth groups members list.....	70
isi auth groups modify.....	71
isi auth groups view.....	72
isi auth id.....	73
isi auth krb5 create.....	73
isi auth krb5 delete.....	74
isi auth krb5 domain create.....	74
isi auth krb5 domain delete.....	74
isi auth krb5 domain list.....	75
isi auth krb5 domain modify.....	75
isi auth krb5 domain view.....	76
isi auth krb5 list.....	76
isi auth krb5 realm create.....	76
isi auth krb5 realm delete.....	77
isi auth krb5 realm list.....	77
isi auth krb5 realm modify.....	78
isi auth krb5 realm view.....	78
isi auth krb5 spn check.....	78
isi auth krb5 spn create.....	79
isi auth krb5 spn delete.....	79
isi auth krb5 spn fix.....	80
isi auth krb5 spn import.....	80
isi auth krb5 spn list.....	80
isi auth krb5 view.....	81
isi auth ldap create.....	81
isi auth ldap delete.....	87
isi auth ldap list.....	87
isi auth ldap modify.....	88
isi auth ldap view.....	96
isi auth local list.....	96
isi auth local modify.....	97
isi auth local view.....	98
isi auth log-level modify.....	99

isi auth log-level view.....	99
isi auth mapping create.....	99
isi auth mapping delete.....	100
isi auth mapping dump.....	101
isi auth mapping flush.....	102
isi auth mapping import.....	103
isi auth mapping list.....	103
isi auth mapping modify.....	103
isi auth mapping token.....	104
isi auth mapping view.....	105
isi auth netgroups view.....	105
isi auth nis create.....	106
isi auth nis delete.....	108
isi auth nis list.....	109
isi auth nis modify.....	109
isi auth nis view.....	114
isi auth privileges.....	114
isi auth refresh.....	114
isi auth roles create.....	115
isi auth roles delete.....	115
isi auth roles list.....	116
isi auth roles members list.....	116
isi auth roles modify.....	117
isi auth roles privileges list.....	119
isi auth roles view.....	120
isi auth settings acls modify.....	120
isi auth settings acls view.....	124
isi auth settings global modify.....	124
isi auth settings global view.....	125
isi auth settings krb5 modify.....	126
isi auth settings krb5 view.....	127
isi auth settings mapping modify.....	127
isi auth settings mapping view.....	128
isi auth status.....	128
isi auth users create.....	129
isi auth users delete.....	130
isi auth users flush.....	131
isi auth users list.....	131
isi auth users modify.....	132
isi auth users view.....	134
isi batterystatus list.....	135
isi batterystatus view.....	135
isi certificate authority delete.....	136
isi certificate authority import.....	136
isi certificate authority list.....	136
isi certificate authority modify.....	137
isi certificate authority view.....	137
isi certificate server delete.....	137
isi certificate server import.....	138
isi certificate server list.....	138

isi certificate server modify.....	139
isi certificate server view.....	139
isi certificate settings modify.....	140
isi certificate settings view.....	140
isi cloud access add.....	141
isi cloud access list.....	141
isi cloud access remove.....	142
isi cloud access view.....	143
isi cloud accounts create.....	143
isi cloud accounts delete.....	145
isi cloud accounts list.....	146
isi cloud accounts modify.....	146
isi cloud accounts view.....	148
isi cloud archive.....	148
isi cloud jobs cancel.....	149
isi cloud jobs create.....	149
isi cloud jobs files list.....	150
isi cloud jobs list.....	151
isi cloud jobs pause.....	152
isi cloud jobs resume.....	153
isi cloud jobs view.....	153
isi cloud pools create.....	154
isi cloud pools delete.....	154
isi cloud pools list.....	155
isi cloud pools modify.....	156
isi cloud pools view.....	157
isi cloud proxies create.....	157
isi cloud proxies delete.....	158
isi cloud proxies list.....	158
isi cloud proxies modify.....	159
isi cloud proxies view.....	160
isi cloud recall.....	161
isi cloud restore_coi.....	161
isi cloud settings modify.....	162
isi cloud settings regenerate-encryption-key.....	163
isi cloud settings view.....	164
isi cluster atime modify.....	164
isi cluster atime view.....	165
isi cluster contact modify.....	165
isi cluster contact view.....	166
isi cluster encoding list.....	166
isi cluster encoding modify.....	166
isi cluster encoding view.....	167
isi cluster identity modify.....	167
isi cluster identity view.....	167
isi cluster internal-networks modify.....	168
isi cluster internal-networks view.....	168
isi cluster join-mode modify.....	169
isi cluster join-mode view.....	169
isi cluster Innset modify.....	169

isi cluster Innset view.....	170
isi cluster reboot.....	170
isi cluster shutdown.....	170
isi cluster time modify.....	171
isi cluster time view.....	171
isi cluster time timezone modify.....	171
isi cluster time timezone view.....	172
isi compression stats list.....	172
isi compression stats view.....	173
isi compression settings modify.....	173
isi compression settings view.....	174
isi config.....	174

3 OneFS isi commands D through L..... 178

isi dedupe inline settings modify.....	181
isi dedupe inline settings view.....	181
isi dedupe reports list.....	182
isi dedupe reports view	182
isi dedupe settings modify.....	183
isi dedupe settings view.....	184
isi dedupe stats.....	185
isi devices add.....	185
isi devices config modify.....	186
isi devices config view.....	187
isi devices drive add.....	187
isi devices drive firmware list.....	187
isi devices drive firmware update list.....	188
isi devices drive firmware update start.....	189
isi devices drive firmware update view.....	189
isi devices drive firmware view.....	189
isi devices format.....	190
isi devices list.....	190
isi devices node add.....	191
isi devices node list.....	191
isi devices node smartfail.....	192
isi devices node stopfail.....	192
isi devices purpose.....	193
isi devices purposelist.....	193
isi devices smartfail.....	194
isi devices stopfail.....	194
isi devices suspend.....	195
isi devices view.....	195
isi diagnostics gather settings modify.....	196
isi diagnostics gather settings view.....	197
isi diagnostics gather start.....	197
isi diagnostics gather status.....	198
isi diagnostics gather stop.....	198
isi diagnostics netlogger settings modify.....	199
isi diagnostics netlogger settings view.....	200
isi diagnostics netlogger start.....	200

isi diagnostics netlogger status.....	201
isi diagnostics netlogger stop.....	201
isi email settings modify.....	201
isi email settings view.....	202
isi esrs modify.....	203
isi esrs view.....	204
isi esrs dataitems list.....	204
isi esrs dataitems modify.....	205
isi esrs download list.....	205
isi esrs download start.....	206
isi esrs download view.....	206
isi esrs telemetry modify.....	207
isi esrs telemetry view.....	207
isi event alerts create.....	207
isi event alerts delete.....	209
isi event alerts list.....	209
isi event alerts modify.....	210
isi event alerts view.....	212
isi event channels create.....	212
isi event channels delete.....	214
isi event channels list.....	215
isi event channels modify.....	215
isi event channels view.....	218
isi event events list.....	218
isi event events view.....	218
isi event groups bulk.....	219
isi event groups list.....	219
isi event groups modify.....	220
isi event groups view.....	221
isi event settings modify.....	221
isi event settings view.....	222
isi event test create.....	222
isi event types list.....	223
isi fc settings list.....	223
isi fc settings modify.....	224
isi fc settings view.....	225
isi file-filter settings modify.....	225
isi file-filter settings view.....	226
isi filepool apply.....	226
isi filepool default-policy modify.....	229
isi filepool default-policy view.....	231
isi filepool policies create.....	231
isi filepool policies delete.....	235
isi filepool policies list.....	236
isi filepool policies modify.....	236
isi filepool policies view.....	240
isi filepool templates list.....	241
isi filepool templates view.....	242
isi_for_array.....	242
isi ftp settings modify.....	244

isi ftp settings view.....	248
isi_gather_info.....	249
isi get.....	254
isi hardening apply.....	255
isi hardening revert.....	255
isi hardening status.....	256
isi hdfs crypto encryption-zones create.....	256
isi hdfs crypto settings modify.....	256
isi hdfs crypto encryption-zones list.....	257
isi hdfs crypto settings view.....	257
isi hdfs fsimage job settings modify.....	257
isi hdfs fsimage job settings view.....	258
isi hdfs fsimage job view.....	258
isi hdfs fsimage latest delete.....	258
isi hdfs fsimage latest view.....	259
isi hdfs fsimage settings modify.....	259
isi hdfs fsimage settings view.....	260
isi hdfs inotify settings modify.....	260
isi hdfs inotify settings view.....	261
isi hdfs inotify stream reset.....	261
isi hdfs inotify stream view.....	261
isi hdfs log-level modify.....	262
isi hdfs log-level view.....	262
isi hdfs proxyusers create.....	262
isi hdfs proxyusers delete.....	263
isi hdfs proxyusers list.....	264
isi hdfs proxyusers members list.....	265
isi hdfs proxyusers modify.....	266
isi hdfs proxyusers view.....	267
isi hdfs racks create.....	268
isi hdfs racks delete.....	268
isi hdfs racks list.....	269
isi hdfs racks modify.....	269
isi hdfs racks view.....	270
isi hdfs ranger-plugin settings modify.....	270
isi hdfs ranger-plugin settings view.....	271
isi hdfs settings modify.....	271
isi hdfs settings view.....	272
isi http settings modify.....	273
isi http settings view.....	273
isi job events list.....	274
isi job jobs cancel.....	276
isi job jobs list.....	276
isi job jobs modify.....	277
isi job jobs pause.....	278
isi job jobs resume.....	279
isi job jobs start.....	279
isi job jobs view.....	281
isi job policies create.....	282
isi job policies delete.....	282

isi job policies list.....	283
isi job policies modify.....	284
isi job policies view.....	285
isi job reports list.....	286
isi job reports view.....	287
isi job statistics view.....	288
isi job status.....	289
isi job types list.....	290
isi job types modify.....	291
isi job types view.....	292
isi license add.....	292
isi license generate.....	293
isi license list.....	294
isi license view.....	294

4 OneFS isi commands N through R..... 295

isi ndmp contexts delete.....	297
isi ndmp contexts list.....	298
isi ndmp contexts view.....	298
isi ndmp dumpdates delete.....	299
isi ndmp dumpdates list.....	299
isi ndmp sessions delete.....	300
isi ndmp sessions list.....	301
isi ndmp sessions view.....	301
isi ndmp settings diagnostics modify.....	302
isi ndmp settings diagnostics view.....	302
isi ndmp settings global modify.....	303
isi ndmp settings global view.....	304
isi ndmp settings preferred-ips create.....	304
isi ndmp settings preferred-ips delete.....	305
isi ndmp settings preferred-ips list.....	305
isi ndmp settings preferred-ips modify.....	306
isi ndmp settings preferred-ips view.....	306
isi ndmp settings variables create.....	307
isi ndmp settings variables delete.....	307
isi ndmp settings variables list.....	308
isi ndmp settings variables modify.....	308
isi ndmp users create.....	309
isi ndmp users delete.....	309
isi ndmp users list.....	310
isi ndmp users modify.....	310
isi ndmp users view.....	311
isi network dnscache flush.....	311
isi network dnscache modify.....	311
isi network dnscache view.....	313
isi network external modify.....	314
isi network external view.....	315
isi network groupnets create.....	315
isi network groupnets delete.....	316
isi network groupnets list.....	316

isi network groupnets modify.....	317
isi network groupnets view.....	318
isi network interfaces list.....	319
isi network pools create	320
isi network pools delete.....	323
isi network pools list.....	323
isi network pools modify.....	325
isi network pools rebalance-ips.....	329
isi network pools sc-resume-nodes.....	329
isi network pools sc-suspend-nodes.....	330
isi network pools view.....	330
isi network rules create.....	331
isi network rules delete.....	331
isi network rules list.....	332
isi network rules modify.....	333
isi network rules view.....	334
isi network sc-rebalance-all.....	334
isi network subnets create.....	334
isi network subnets delete.....	335
isi network subnets list.....	336
isi network subnets modify.....	337
isi network subnets view.....	339
isi nfs aliases create.....	339
isi nfs aliases delete.....	340
isi nfs aliases list.....	340
isi nfs aliases modify.....	341
isi nfs aliases view.....	342
isi nfs exports check.....	343
isi nfs exports create.....	344
isi nfs exports delete.....	349
isi nfs exports list.....	350
isi nfs exports modify.....	352
isi nfs exports reload.....	361
isi nfs exports view.....	361
isi nfs log-level modify.....	361
isi nfs log-level view.....	362
isi nfs netgroup check.....	362
isi nfs netgroup flush.....	363
isi nfs netgroup modify.....	363
isi nfs nlm locks list.....	364
isi nfs nlm locks waiters.....	365
isi nfs nlm sessions check.....	366
isi nfs nlm sessions delete.....	366
isi nfs nlm sessions list.....	366
isi nfs nlm sessions refresh.....	367
isi nfs nlm sessions view.....	368
isi nfs settings export modify.....	368
isi nfs settings export view.....	375
isi nfs settings global modify.....	377
isi nfs settings global view.....	377

isi nfs settings zone modify.....	377
isi nfs settings zone view.....	379
isi ntp servers create.....	379
isi ntp servers delete.....	379
isi ntp servers list.....	380
isi ntp servers modify.....	380
isi ntp servers view.....	381
isi ntp settings modify.....	381
isi ntp settings view.....	382
isi performance datasets create.....	382
isi performance datasets delete.....	382
isi performance datasets list.....	383
isi performance datasets modify.....	384
isi performance datasets view.....	384
isi performance filters apply.....	384
isi performance filters list.....	385
isi performance filters modify.....	385
isi performance filters remove.....	386
isi performance filters view.....	386
isi performance metrics list.....	386
isi performance metrics view.....	387
isi performance settings modify.....	387
isi performance settings view.....	387
isi performance workloads list.....	388
isi performance workloads modify.....	388
isi performance workloads pin.....	389
isi performance workloads unpin.....	389
isi performance workloads view.....	389
isi_phone_home.....	390
isi quota quotas create.....	391
isi quota quotas delete.....	393
isi quota quotas list.....	394
isi quota quotas modify.....	396
isi quota quotas notifications clear.....	398
isi quota quotas notifications create.....	400
isi quota quotas notifications delete.....	402
isi quota quotas notifications disable.....	404
isi quota quotas notifications list.....	405
isi quota quotas notifications modify.....	406
isi quota quotas notifications view.....	409
isi quota quotas view.....	410
isi quota reports create.....	411
isi quota reports delete.....	412
isi quota reports list.....	413
isi quota settings mappings create.....	413
isi quota settings mappings delete.....	414
isi quota settings mappings list.....	414
isi quota settings mappings modify.....	415
isi quota settings mappings view.....	415
isi quota settings notifications clear.....	415

isi quota settings notifications create.....	416
isi quota settings notifications delete.....	418
isi quota settings notifications list.....	418
isi quota settings notifications modify.....	419
isi quota settings notifications view.....	421
isi quota settings reports modify.....	422
isi quota settings reports view.....	423
isi readonly list.....	423
isi readonly modify.....	424
isi readonly view.....	424
isi remotesupport connectemc modify.....	425
isi remotesupport connectemc view.....	426

5 OneFS isi commands S through Z.....427

isi servicelight list.....	431
isi servicelight modify.....	431
isi servicelight view.....	432
isi services.....	432
isi set.....	432
isi smb log-level filters create.....	435
isi smb log-level filters delete.....	436
isi smb log-level filters list.....	436
isi smb log-level filters view.....	437
isi smb log-level modify.....	437
isi smb log-level view.....	438
isi smb openfiles close.....	438
isi smb openfiles list.....	439
isi smb sessions delete.....	439
isi smb sessions delete-user.....	440
isi smb sessions list.....	441
isi smb settings global modify.....	441
isi smb settings global view.....	444
isi smb settings shares modify.....	444
isi smb settings shares view.....	448
isi smb settings zone modify.....	448
isi smb settings zone view.....	450
isi smb shares create.....	450
isi smb shares delete.....	453
isi smb shares list.....	453
isi smb shares modify.....	454
isi smb shares permission create.....	459
isi smb shares permission delete.....	460
isi smb shares permission list.....	460
isi smb shares permission modify.....	461
isi smb shares permission view.....	462
isi smb shares view.....	462
isi snapshot aliases create.....	463
isi snapshot aliases delete.....	463
isi snapshot aliases list.....	464
isi snapshot aliases modify.....	464

isi snapshot aliases view.....	465
isi snapshot locks create.....	465
isi snapshot locks delete.....	466
isi snapshot locks list.....	467
isi snapshot locks modify.....	467
isi snapshot locks view.....	468
isi snapshot schedules create.....	469
isi snapshot schedules delete.....	470
isi snapshot schedules list.....	471
isi snapshot schedules modify.....	472
isi snapshot schedules pending list.....	473
isi snapshot schedules view.....	474
isi snapshot settings modify.....	475
isi snapshot settings view.....	476
isi snapshot snapshots create.....	476
isi snapshot snapshots delete.....	477
isi snapshot snapshots list.....	478
isi snapshot snapshots modify.....	479
isi snapshot snapshots view.....	480
isi snmp settings modify.....	480
isi snmp settings view.....	482
isi auth settings modify.....	482
isi statistics client.....	484
isi statistics data-reduction.....	487
isi statistics data-reduction view.....	488
isi statistics drive.....	488
isi statistics heat.....	490
isi statistics list keys.....	492
isi statistics list operations.....	493
isi statistics protocol.....	494
isi statistics pstat.....	497
isi statistics query current.....	498
isi statistics query history.....	499
isi status.....	501
isi storagepool compatibilities active create.....	501
isi storagepool compatibilities active delete.....	502
isi storagepool compatibilities active list.....	503
isi storagepool compatibilities active view.....	504
isi storagepool compatibilities available list.....	504
isi storagepool compatibilities class active create.....	505
isi storagepool compatibilities class active delete.....	506
isi storagepool compatibilities class active list.....	506
isi storagepool compatibilities class active view.....	507
isi storagepool compatibilities class available list.....	508
isi storagepool compatibilities ssd active create.....	509
isi storagepool compatibilities ssd active delete.....	510
isi storagepool compatibilities ssd active list.....	511
isi storagepool compatibilities ssd active view.....	512
isi storagepool compatibilities ssd available list.....	512
isi storagepool health.....	513

isi storagepool list.....	513
isi storagepool nodepools create.....	514
isi storagepool nodepools delete.....	514
isi storagepool nodepools list.....	515
isi storagepool nodepools modify.....	515
isi storagepool nodepools view.....	517
isi storagepool settings modify.....	517
isi storagepool settings modify.....	518
isi storagepool settings view.....	520
isi storagepool tiers create.....	520
isi storagepool tiers delete.....	521
isi storagepool tiers list.....	521
isi storagepool tiers modify.....	522
isi storagepool tiers view.....	522
isi storagepool unprovisioned view.....	523
isi swift accounts create.....	523
isi swift accounts delete.....	524
isi swift accounts list.....	524
isi swift accounts modify.....	525
isi swift accounts view.....	526
isi sync certificates peer delete.....	526
isi sync certificates peer import.....	526
isi sync certificates peer list.....	527
isi sync certificates peer modify.....	527
isi sync certificates peer view.....	528
isi sync certificates server delete.....	528
isi sync certificates server import.....	529
isi sync certificates server list.....	529
isi sync certificates server modify.....	530
isi sync certificates server view.....	530
isi sync jobs cancel.....	531
isi sync jobs list.....	531
isi sync jobs pause.....	532
isi sync jobs reports list.....	532
isi sync jobs reports view.....	533
isi sync jobs resume.....	533
isi sync jobs start.....	533
isi sync jobs view.....	534
isi sync policies create.....	534
isi sync policies delete.....	540
isi sync policies disable.....	541
isi sync policies enable.....	541
isi sync policies list.....	542
isi sync policies modify.....	543
isi sync policies reset.....	551
isi sync policies resolve.....	552
isi sync policies view.....	552
isi sync recovery allow-write.....	552
isi sync recovery resync-prep.....	553
isi sync reports list.....	553

isi sync reports rotate.....	555
isi sync reports subreports list.....	555
isi sync reports subreports view.....	556
isi sync reports view.....	557
isi sync rules create.....	557
isi sync rules delete.....	558
isi sync rules list.....	559
isi sync rules modify.....	559
isi sync rules reports list.....	560
isi sync rules reports view.....	561
isi sync rules view.....	561
isi sync service policies create.....	561
isi sync service policies delete.....	566
isi sync service policies disable.....	567
isi sync service policies enable.....	567
isi sync service policies list.....	567
isi sync service policies modify.....	568
isi sync service policies reset.....	573
isi sync service policies resolve.....	574
isi sync service policies view.....	574
isi sync service recovery allow-write.....	574
isi sync service recovery resync-prep.....	575
isi sync service target break.....	576
isi sync service target cancel.....	576
isi sync service target list.....	576
isi sync service target view.....	577
isi sync settings modify.....	577
isi sync settings view.....	579
isi sync target break.....	579
isi sync target cancel.....	580
isi sync target list.....	580
isi sync target reports list.....	581
isi sync target reports subreports list.....	583
isi sync target reports subreports view.....	584
isi sync target reports view.....	584
isi sync target view.....	585
isi tape delete.....	585
isi tape list.....	586
isi tape modify.....	586
isi tape rename.....	587
isi tape rescan.....	587
isi tape view.....	588
isi upgrade cluster add-nodes.....	588
isi upgrade cluster add-remaining-nodes.....	589
isi upgrade cluster archive.....	589
isi upgrade cluster assess.....	589
isi upgrade cluster commit.....	589
isi upgrade cluster firmware.....	590
isi upgrade cluster from-version.....	590
isi upgrade cluster nodes firmware.....	591

isi upgrade cluster nodes list.....	591
isi upgrade cluster nodes view.....	592
isi upgrade cluster retry-last-action.....	593
isi upgrade cluster rollback.....	593
isi upgrade cluster rolling-reboot.....	593
isi upgrade cluster settings.....	594
isi upgrade cluster start.....	594
isi upgrade cluster to-version.....	594
isi upgrade cluster view.....	595
isi upgrade patches abort.....	595
isi upgrade patches install.....	595
isi upgrade patches list.....	596
isi upgrade patches uninstall.....	596
isi upgrade patches view.....	597
isi version.....	597
isi worm cdate set.....	597
isi worm cdate view.....	598
isi worm create.....	598
isi worm domains create.....	598
isi worm domains list.....	601
isi worm domains modify.....	602
isi worm domains view	605
isi worm files delete.....	605
isi worm files view.....	605
isi zone restrictions create.....	606
isi zone restrictions delete.....	606
isi zone restrictions list.....	607
isi zone zones create.....	608
isi zone zones delete.....	609
isi zone zones list.....	609
isi zone zones modify.....	610
isi zone zones view.....	611

Introduction to this guide

This chapter describes this reference guide, and provides information about how to get assistance from Isilon technical support.

Topics:

- [About this guide](#)
- [Where to go for support](#)

About this guide

This guide lists and describes all OneFS-specific commands that extend the standard UNIX command set.

NOTE: The documentation for the `isi healthcheck` CLI command has its own OneFS document and Dell EMC Info Hub, which can be found at <https://community.emc.com/docs/DOC-71575>.

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to <https://www.research.net/s/isi-docfeedback>. If you cannot provide feedback through the URL, send an email message to docfeedback@isilon.com.

Where to go for support

This topic contains resources for getting answers to questions about Isilon products.

Online support	<ul style="list-style-type: none"> • Live Chat • Create a Service Request <p>For questions about accessing online support, send an email to support@emc.com.</p>
Telephone support	<ul style="list-style-type: none"> • United States: 1-800-SVC-4EMC (1-800-782-4362) • Canada: 1-800-543-4782 • Worldwide: 1-508-497-7901 • Local phone numbers for a specific country are available at Dell EMC Customer Support Centers.
Isilon Community Network	The Isilon Community Network connects you to a central hub of information and experts to help you maximize your current storage solution. From this site, you can demonstrate Isilon products, ask questions, view technical videos, and get the latest Isilon product documentation.
Isilon Info Hubs	For the list of Isilon info hubs, see the Isilon Info Hubs page on the Isilon Community Network . Use these info hubs to find product documentation, troubleshooting guides, videos, blogs, and other information resources about the Isilon products and features you're interested in.

Remote Support

For information about remote support, see the [OneFS 8.2.0 Web Administration Guide](#).

OneFS isi commands A through C

This chapter contains documentation of the OneFS CLI commands `isi antivirus policies create` through `isi config`.

Topics:

- `isi antivirus policies create`
- `isi antivirus policies delete`
- `isi antivirus policies list`
- `isi antivirus policies modify`
- `isi antivirus policies start`
- `isi antivirus policies view`
- `isi antivirus quarantine`
- `isi antivirus release`
- `isi antivirus reports delete`
- `isi antivirus reports scans list`
- `isi antivirus reports scans view`
- `isi antivirus reports threats list`
- `isi antivirus reports threats view`
- `isi antivirus scan`
- `isi antivirus servers create`
- `isi antivirus servers delete`
- `isi antivirus servers list`
- `isi antivirus servers modify`
- `isi antivirus servers view`
- `isi antivirus settings modify`
- `isi antivirus settings view`
- `isi antivirus status`
- `isi audit progress global view`
- `isi audit progress view`
- `isi audit settings global modify`
- `isi audit settings global view`
- `isi audit settings modify`
- `isi audit settings view`
- `isi audit topics list`
- `isi audit topics modify`
- `isi audit topics view`
- `isi auth access`
- `isi auth ads create`
- `isi auth ads delete`
- `isi auth ads list`
- `isi auth ads modify`
- `isi auth ads spn check`
- `isi auth ads spn create`
- `isi auth ads spn delete`
- `isi auth ads spn fix`
- `isi auth ads spn list`
- `isi auth ads trusts controllers list`
- `isi auth ads trusts list`
- `isi auth ads view`
- `isi auth duo modify`
- `isi auth error`

- isi auth file create
- isi auth file delete
- isi auth file list
- isi auth file modify
- isi auth file view
- isi auth groups create
- isi auth groups delete
- isi auth groups flush
- isi auth groups list
- isi auth groups members list
- isi auth groups modify
- isi auth groups view
- isi auth id
- isi auth krb5 create
- isi auth krb5 delete
- isi auth krb5 domain create
- isi auth krb5 domain delete
- isi auth krb5 domain list
- isi auth krb5 domain modify
- isi auth krb5 domain view
- isi auth krb5 list
- isi auth krb5 realm create
- isi auth krb5 realm delete
- isi auth krb5 realm list
- isi auth krb5 realm modify
- isi auth krb5 realm view
- isi auth krb5 spn check
- isi auth krb5 spn create
- isi auth krb5 spn delete
- isi auth krb5 spn fix
- isi auth krb5 spn import
- isi auth krb5 spn list
- isi auth krb5 view
- isi auth ldap create
- isi auth ldap delete
- isi auth ldap list
- isi auth ldap modify
- isi auth ldap view
- isi auth local list
- isi auth local modify
- isi auth local view
- isi auth log-level modify
- isi auth log-level view
- isi auth mapping create
- isi auth mapping delete
- isi auth mapping dump
- isi auth mapping flush
- isi auth mapping import
- isi auth mapping list
- isi auth mapping modify
- isi auth mapping token
- isi auth mapping view
- isi auth netgroups view
- isi auth nis create
- isi auth nis delete
- isi auth nis list
- isi auth nis modify

- isi auth nis view
- isi auth privileges
- isi auth refresh
- isi auth roles create
- isi auth roles delete
- isi auth roles list
- isi auth roles members list
- isi auth roles modify
- isi auth roles privileges list
- isi auth roles view
- isi auth settings acls modify
- isi auth settings acls view
- isi auth settings global modify
- isi auth settings global view
- isi auth settings krb5 modify
- isi auth settings krb5 view
- isi auth settings mapping modify
- isi auth settings mapping view
- isi auth status
- isi auth users create
- isi auth users delete
- isi auth users flush
- isi auth users list
- isi auth users modify
- isi auth users view
- isi batterystatus list
- isi batterystatus view
- isi certificate authority delete
- isi certificate authority import
- isi certificate authority list
- isi certificate authority modify
- isi certificate authority view
- isi certificate server delete
- isi certificate server import
- isi certificate server list
- isi certificate server modify
- isi certificate server view
- isi certificate settings modify
- isi certificate settings view
- isi cloud access add
- isi cloud access list
- isi cloud access remove
- isi cloud access view
- isi cloud accounts create
- isi cloud accounts delete
- isi cloud accounts list
- isi cloud accounts modify
- isi cloud accounts view
- isi cloud archive
- isi cloud jobs cancel
- isi cloud jobs create
- isi cloud jobs files list
- isi cloud jobs list
- isi cloud jobs pause
- isi cloud jobs resume
- isi cloud jobs view
- isi cloud pools create

- isi cloud pools delete
- isi cloud pools list
- isi cloud pools modify
- isi cloud pools view
- isi cloud proxies create
- isi cloud proxies delete
- isi cloud proxies list
- isi cloud proxies modify
- isi cloud proxies view
- isi cloud recall
- isi cloud restore_coi
- isi cloud settings modify
- isi cloud settings regenerate-encryption-key
- isi cloud settings view
- isi cluster atime modify
- isi cluster atime view
- isi cluster contact modify
- isi cluster contact view
- isi cluster encoding list
- isi cluster encoding modify
- isi cluster encoding view
- isi cluster identity modify
- isi cluster identity view
- isi cluster internal-networks modify
- isi cluster internal-networks view
- isi cluster join-mode modify
- isi cluster join-mode view
- isi cluster Innset modify
- isi cluster Innset view
- isi cluster reboot
- isi cluster shutdown
- isi cluster time modify
- isi cluster time view
- isi cluster time timezone modify
- isi cluster time timezone view
- isi compression stats list
- isi compression stats view
- isi compression settings modify
- isi compression settings view
- isi config

isi antivirus policies create

Creates an antivirus scan policy.

Syntax

```
isi antivirus policies create <name>
  [--description <string>]
  [--enabled {true | false}]
  [--schedule <schedule>]
  [--impact <impact-policy>]
  [--force-run {yes | no}]
  [--paths <path>...]
  [--recursion-depth <integer>]
  [--verbose]
```

Options

<name>

Specifies a name for the policy.

--description <string>

Specifies a description for the policy.

{--enabled | -e} {true | false}

Determines whether the policy is enabled or disabled. If set to `true`, the policy is enabled. The default value is `false`.

{--schedule | -s} <schedule>

Specifies when the policy is run.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify *<frequency>* in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "Saturday" and "sat" are valid.

{--impact | -i} <impact_policy>

Specifies an impact policy for the antivirus scan jobs. You can specify `LOW`, `MEDIUM`, `HIGH`, `OFF_HOURS`, or a custom impact policy that you created.

{--force-run | -r} {yes | no}

Determines whether to force policy scans. If a scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

--paths <path>

Specifies directories to scan when the policy is run. To specify multiple paths, repeat the `--path` option. For example:

```
--paths /ifs/data/directory1 --paths /ifs/data/directory2
```

`--recursion-depth <integer>`

 **NOTE: This option has been deprecated and will not impact antivirus scans if specified.**

Specifies the depth of subdirectories to include in the scan.

`{--verbose | -v}`

Displays a message confirming that the antivirus policy was created.

isi antivirus policies delete

Deletes an antivirus scan policy.

Syntax

```
isi antivirus policies delete {<name> | --all}
  [--force]
  [--verbose]
```

Options

`{<name> | --all}`

Deletes the specified policy or all policies.

`{--force | -f}`

Does not prompt you to confirm that you want to delete the policy.

`{--verbose | -v}`

Displays a message confirming that the antivirus policy was deleted.

isi antivirus policies list

List antivirus scan policies.

Syntax

```
isi antivirus policies list
  [--limit <integer>]
  [--sort {name | description | enabled}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--sort <attribute>`

Sorts output displayed by the specified attribute.

The following values are valid:

name	Sorts output by the URL of the server.
description	Sorts output by the description of the server.
enabled	Sorts output by the state of the server.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi antivirus policies modify

Modifies an antivirus scan policy.

Syntax

```
isi antivirus policies modify <id>
  [--name <new-name>]
  [--description <string>]
  [--enabled {true | false}]
  [--schedule <schedule>]
  [--clear-schedule]
  [--impact <impact-policy>]
  [--clear-impact]
  [--force-run {true | false}]
  [--paths <path>]
  [--clear-paths]
  [--add-paths <path>]
  [--remove-paths <path>]
  [--recursion-depth <integer>]
```

Options

<id>

Modifies the policy with the specified policy identification number.

{--name | -n} <new-name>

Specifies a new name for this policy.

--description <string>

Specifies a description for the policy.

{--enabled | -e} {true | false}

Determines whether this policy is enabled or disabled. If set to `true`, the policy is enabled. The default value is `false`.

{--schedule | -s} <schedule>

Specifies when the policy is run.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify `<interval>` in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`

- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "Saturday" and "sat" are valid.

--clear-schedule

Deletes the current schedule for the policy.

{--impact | -i} <impact_policy>

Specifies an impact policy for the antivirus scan jobs. You can specify LOW, MEDIUM, HIGH, OFF_HOURS, or a custom impact policy that you created.

--clear-impact

Clears the current impact policy for antivirus scan jobs.

{--force-run | -r} {yes | no}

Determines whether to force policy scans. If a scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

--paths <path>

Specifies directories to scan when the policy is run. To specify multiple paths, repeat the --path option. For example:

```
--paths /ifs/data/directory1 --paths /ifs/data/directory2
```

i **NOTE:** If you specify this option, the specified paths will replace all previously specified paths in the list.

--clear-paths

Clears the list of paths to scan.

--add-paths <path>

Adds the specified path to the list of paths to scan.

--remove-paths <path>

Removes the specified path from the list of paths to scan.

{--verbose | -v}

Displays a message confirming that the antivirus policy was modified.

isi antivirus policies start

Runs an antivirus policy.

Syntax

```
isi antivirus policies start <policy>
  [--report-id <id>]
  [--force-run {true | false}]
  [--update {yes | no}]
```

Options

<policy>

Runs the specified policy.

--report-id <id>

Assigns the specified ID to the report generated for this run of the avscan policy. If you do not specify an ID, OneFS will automatically assign one.

{ --force-run | -r } {true | false}

Determines whether to force the scan. If the scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

--update {yes | no}

Specifies whether to update the last run time in the policy file. The default value is `yes`.

isi antivirus policies view

Displays information about antivirus scan policies.

Syntax

```
isi antivirus policies view <name>
```

Options

<name>

Displays information about the named policy.

isi antivirus quarantine

Quarantines a file manually. Quarantined files cannot be read or written to.

Syntax

```
isi antivirus quarantine <file>
  [--verbose]
```

Options

<file>

Quarantines the specified file. Specify as an absolute file path within the `/ifs` file system.

{--verbose | -v}

Displays a message confirming that the file has been quarantined.

isi antivirus release

Removes a file from quarantine. Quarantined files cannot be read or written to.

Syntax

```
isi antivirus release <file>
  [--verbose]
```

Options

<file>

Removes the specified file from quarantine. Specify as a file path within the /ifs file system.

{--verbose | -v}

Displays a message confirming that the file was removed from quarantine.

isi antivirus reports delete

Deletes antivirus reports.

Syntax

```
isi antivirus reports delete {<scan-id> | --all}
  [--age <integer><time>]
  [--verbose]
  [--force]
```

Options

<scan-id>

Deletes the antivirus report with the specified ID.

--all

Deletes all antivirus reports.

--age <integer><time>

Deletes all reports older than the specified age.

The following *<time>* values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

{--verbose | -v}

Displays a message confirming that the reports have been deleted.

{--force | -f}

Does not display a confirmation prompt.

isi antivirus reports scans list

Displays information about recent antivirus scans.

Syntax

```
isi antivirus reports scans list
  [--policy-id <string>]
  [--status <status>]
  [--limit <integer>]
  [--offset <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--policy-id <string>

Filters output based on the ID of the policy.

--status <status>

Filters output based on the current status of the scan job.

The following values are valid:

Finish	Displays only completed jobs.
Succeeded	Displays only successfully completed jobs.
Failed	Displays only failed jobs.
Cancelled	Displays only cancelled jobs.
Started	Displays only running jobs.
Paused	Displays only paused jobs.
Resumed	Displays only jobs that were paused, then resumed.
Pending	Displays only pending jobs.

{--limit | -l} <integer>

Displays no more than the specified number of items.

{--offset | -o} <integer>

Specifies the number of entries to bypass from the beginning of the scan report.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

id	Sorts output by the ID of the antivirus report.
policy_id	Sorts output by the ID of the policy that created the report.
status	Sorts output by the status of the antivirus scan.
start	Sorts output by the time that the antivirus scan started.
files	Sorts output by the number of files that were scanned by the antivirus scan.

infections Sorts output by the number of threats detected by the antivirus scan.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi antivirus reports scans view

Displays an antivirus report.

Syntax

```
isi antivirus reports scans view <id>
```

Options

<id>

Displays the antivirus report of the specified ID.

isi antivirus reports threats list

List antivirus threat reports.

Syntax

```
isi antivirus reports threats list
  [--scan-id <string>]
  [--file <string>]
  [--remediation <string>]
  [--limit <integer>]
  [--offset <integer>]
  [--sort {scan_id | file | remediation | threat | time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--scan-id <string>

Unique identifier for the scan report.

--file <string>

Name of the file containing the threat report.

--remediation <string>

Description of action taken to remediate threat.

{--limit | -l} <integer>

Maximum number of antivirus threats to display.

{--offset | -o} <integer>

Number of threat report entries to bypass from the beginning.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

scan-id	Unique identifier of threat report scans.
file	Files containing threats.
remediation	Actions taken to alleviate threats.
threat	Sorts by specific threats.
time	Sorts by times of threats.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi antivirus reports threats view

Displays information about a detected threats.

Syntax

```
isi antivirus reports threats view <id>
```

Options

<id>

Displays information about the threat with the specified ID.

isi antivirus scan

Manually scans a file for viruses.

Syntax

```
isi antivirus scan <file>  
  [--policy <id>]  
  [--report-id <id>]  
  [--force-run {yes | no}]
```

Options

<file>

Scans the specified file.

{--policy | -p} <id>

Assigns a policy ID for this scan. The default ID is MANUAL.

--report-id <id>

Assigns the specified ID to the report generated for this antivirus scan. If you do not specify an ID, OneFS will automatically assign one.

{ --force-run | -r} {true | false}

Determines whether to force the scan. If the scan is forced, all files are scanned regardless of whether OneFS has marked files as having been scanned, or if global settings specify that certain files should not be scanned.

isi antivirus servers create

Adds and connects to an ICAP server.

Syntax

```
isi antivirus servers create <url>  
  [--description <string>]  
  [--enabled {yes | no}]  
  [--verbose]
```

Options

<url>

Specifies the URL of the ICAP server.

--description <string>

Specifies an optional description for the policy.

{--enabled | -n} {yes | no}

Determines whether the ICAP server is enabled.

{--verbose | -v}

Displays a message confirming that the server has been added.

isi antivirus servers delete

Deletes antivirus servers.

Syntax

```
isi antivirus servers delete {<url> | --all}  
  [--verbose]  
  [--force]
```

Options

<url>

Deletes the specified antivirus server.

--all

Deletes all antivirus servers.

{--verbose | -v}

Displays a message confirming that OneFS has disconnected from the ICAP server.

{--force | -f}

Does not display a confirmation prompt.

isi antivirus servers list

Displays a list of antivirus servers that OneFS is currently connected to.

Syntax

```
isi antivirus servers list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

url	Sorts output by the URL of the server.
description	Sorts output by the description of the server.
enabled	Sorts output by the state of the server.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi antivirus servers modify

Modifies the settings of an ICAP server.

Syntax

```
isi antivirus servers modify <url>
  [--enabled {yes | no}]
```

```
[--description <string>]
[--verbose]
```

Options

<url>

Specifies the URL of the ICAP server.

--description <string>

Specifies an optional description for the policy.

{--enabled | -n} {yes | no}

Determines whether the ICAP server is enabled.

{--verbose | -v}

Displays a message confirming that the server has been added.

isi antivirus servers view

Displays information about an ICAP server.

Syntax

```
isi antivirus servers view <url>
```

Options

<url>

Displays information about the specified ICAP server.

isi antivirus settings modify

Sets and displays global configuration settings for anti-virus scanning.

Syntax

```
isi antivirus settings modify
  [--fail-open {true | false}]
  [{--glob-filters <string>... | --clear-glob-filters
  | --add-glob-filters <string> | --remove-glob-filters <string>}]
  [--glob-filters-enabled {true | false}]
  [--glob-filters-include {true | false}]
  [--path-prefixes <path>... | --clear-path-prefixes
  | --add-path-prefixes <path> | --remove-path-prefixes <path>}]
  [--repair {true | false}]
  [--report-expiry <integer><time>]
  [--scan-cloudpool-files {true | false}]
  [--scan-on-close {true | false}]
  [--scan-on-open {true | false}]
  [--scan-size-maximum <integer>{k | M | G | T | P}]
  [--service {true | false}]
  [--quarantine {true | false}]
  [--truncate {true | false}]
  [--verbose]
```

Options

`--fail-open {true | false}`

If `--scan-on-open` is set to `true`, determines whether users can access files that cannot be scanned. If this option is set to `false`, users cannot access a file until the file is scanned by an ICAP server.

If `--scan-on-open` is set to `true`, this option has no effect.

`--glob-filter <string>`

Specifies a file name or extension. To specify multiple filters, you must include multiple `--glob-filter` options within the same command. Specifying this option will remove any existing glob filters.

You can include the following wildcards:

Wildcard character	Description
*	Matches any string in place of the asterisk. For example, specifying "m*" would match "movies" and "m123"
[]	Matches any characters contained in the brackets, or a range of characters separated by a dash. For example, specifying "b[aei]t" would match "bat", "bet", and "bit" For example, specifying "1[4-7]2" would match "142", "152", "162", and "172" You can exclude characters within brackets by following the first bracket with an exclamation mark. For example, specifying "b[!ie]" would match "bat" but not "bit" or "bet" You can match a bracket within a bracket if it is either the first or last character. For example, specifying "[[c]at" would match "cat", and "[at" You can match a dash within a bracket if it is either the first or last character. For example, specifying "car[-s]" would match "cars", and "car-"
?	Matches any character in place of the question mark. For example, specifying "t?p" would match "tap", "tip", and "top"

 **NOTE:** If you specify this option, the specified filters will replace all previously specified filters in the list.

`--clear-glob-filters`

Clears the list of filters.

`--add-glob-filters <string>`

Adds the specified filters to the list of filters.

`--remove-glob-filters <string>`

Removes the specified filters to the list of filters.

`--glob-filters-enabled {true | false}`

Determines whether glob filters are enabled. If no glob filters are specified, glob filters will remain disabled even if this option is set to `true`.

`--glob-filters-include {true | false}`

Determines how glob filters are interpreted by OneFS. If set to `true`, OneFS will scan only files that match a glob filter. If set to `false`, OneFS will scan only files that do not match any glob filters.

`--path-prefix <path>`

If specified, only files contained in the specified directory path will be scanned. This option affects only on-access scans. To specify multiple directories, you must include multiple `--path-prefix` options within the same command. Specifying this option will remove any existing path prefixes.

 **NOTE: If you specify this option, the specified filters will replace all previously specified filters in the list.**

--clear-path-prefixes

Clears the list of paths.

--add-path-prefixes <path>

Adds the specified paths to the list of paths.

--remove-path-prefixes <path>

Removes the specified paths to the list of paths.

--repair {true | false}

Determines whether OneFS attempts to repair files that threats are detected in.

--report-expiry <integer> <time>

Determines how long OneFS will retain antivirus scan reports before deleting them.

The following <time> values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

--scan-cloudpool-files {true | false}

Determines whether cloudpool files are scanned for antiviruses..

--scan-on-close {true | false}

Determines whether files are scanned after the files are closed.

--scan-on-open {true | false}

Determines whether files are scanned before the files are sent to users.

--scan-size-maximum <integer>{k | M | G | T | P}

If specified, OneFS will not send files larger than the specified size to an ICAP server to be scanned.

 **NOTE: Although the parameter accepts values larger than 2GB, OneFS does not scan files larger than 2GB.**

--service {true | false}

Determines whether the antivirus service is running.

--quarantine {true | false}

Determines whether OneFS quarantines files that threats are detected in. If `--repair` is set to `true`, OneFS will attempt to repair the files before quarantining them. If both `--truncate` and `--quarantine` are set to `true`, the `--truncate` option is ignored.

--truncate {true | false}

Determines whether OneFS truncates files that threats are detected in. If `--repair` is set to `true`, OneFS will attempt to repair the files before truncating them. If both `--truncate` and `--quarantine` are set to `true`, the `--truncate` option is ignored.

{--verbose | -v}

Displays a message confirming that the settings have been modified.

isi antivirus settings view

Displays antivirus settings.

Syntax

```
isi antivirus settings view
```

Options

There are no options for this command.

isi antivirus status

Displays information about the scan status of files.

Syntax

```
isi antivirus status <file>
```

Options

<file>

Displays information about the specified file. Specify the full pathname of the file. The file must be within the `/ifs` file system.

{--verbose | -v}

Displays more detailed information.

isi audit progress global view

Displays the latest protocol audit event log time for the cluster. It also displays the time of the oldest unsend protocol audit event to the CEE server and the time of the oldest non-forwarded protocol audit event to syslog in the cluster.

Syntax

```
isi audit progress global view
```

OneFS displays log times for cluster audit events.

isi audit progress view

Displays the progress of delivery of the protocol audit events to the CEE server and syslog for the current node. This is the default view. You can also specify the logical node number to view the progress of delivery of the protocol audit events for the current node.

Syntax

```
isi audit progress view  
[--lnn <integer>]
```

Options

--lnn<integer>

Displays a logical node number view of the progress of delivery of the protocol audit events to the CEE server and syslog. The view includes the timestamp of the last captured protocol audit event and the timestamp of the last event sent to the CEE server and syslog corresponding to the node.

The following command displays the progress of delivery of the protocol audit events to the CEE server and syslog for the current node:

```
isi audit progress view
```

OneFS displays the progress report.

The following command displays a logical node number view of the progress of delivery of the protocol audit events to the CEE server and syslog:

```
isi audit progress view --lnn=2
```

OneFS displays the progress report by logical node number view.

isi audit settings global modify

Enables or disables global auditing configuration changes and protocol access, and configures additional protocol-auditing settings on your cluster.

Syntax

```
isi audit settings global modify
  [--protocol-auditing-enabled {yes | no}]
  [--audited-zones <zones>]
  [--clear-audited-zones]
  [--add-audited-zones <zones>]
  [--remove-audited-zones <zones>]
  [--cee-server-uris <uris>]
  [--clear-cee-server-uris]
  [--add-cee-server-uris <uris>]
  [--remove-cee-server-uris <uris>]
  [--hostname <string>]
  [--config-auditing-enabled {yes | no}]
  [--config-syslog-enabled {yes | no}]
  [--cee-log-time <string>]
  [--syslog-log-time <string>]
  [--config-syslog-servers <string>]
  [--clear-config-syslog-servers]
  [--add-config-syslog-servers <string>]
  [--remove-config-syslog-servers <string>]
  [--protocol-syslog-servers <string>]
  [--clear-protocol-syslog-servers]
  [--add-protocol-syslog-servers <string>]
  [--remove-protocol-syslog-servers <string>]
  [--verbose]
```

Options

--protocol-auditing-enabled {yes | no}

Enables or disables the auditing of data-access requests through the SMB, NFS, and HDFS protocols.

--audited-zones <access zones>

Specifies one or more access zones, separated by commas, which will be audited if protocol auditing is enabled. This option overwrites all entries in the list of access zones; to add or remove access zones without affecting current entries, use `--add-audited-zones` or `--remove-audited-zones`.

--clear-audited-zones

Clears the entire list of access zones to be audited if protocol auditing is enabled.

--add-audited-zones <access zones>

Adds one or more access zones, separated by commas, to the list of zones that will be audited if protocol auditing is enabled.

--remove-audited-zones <access zones>

Removes one or more access zones, separated by commas, which will be audited if protocol auditing is enabled.

--cee-server-uris <uris>

Specifies one or more CEE server URIs, separated by commas, where audit logs will be forwarded if protocol auditing is enabled. The OneFS CEE export service uses round robin load-balancing when exporting events to multiple CEE servers. This option overwrites all entries in the list of CEE server URIs. To add or remove URIs without affecting current entries, use `--add-cee-server-uris` or `--remove-cee-server-uris`.

--clear-cee-server-uris

Clears the entire list of CEE server URIs to which audit logs are forwarded if protocol auditing is enabled.

--add-cee-server-uris <uris>

Adds one or more CEE server URIs, separated by commas, to the list of URIs where audit logs are forwarded if protocol auditing is enabled.

--remove-cee-server-uris <uris>

Removes one or more CEE server URIs, separated by commas, from the list of URIs where audit logs are forwarded if protocol auditing is enabled.

--hostname <string>

Specifies the name of the storage cluster to use when forwarding protocol events—typically, the SmartConnect zone name. When SmartConnect is not implemented, the value must match the hostname of the cluster as your third-party audit application recognizes it. If the field is left blank, events from each node are filled with the node name (clustername + Inn). This setting is required only if needed by your third-party audit application.

--config-auditing-enabled {yes | no}

Enables or disables the auditing of requests made through the API for system configuration changes.

--config-syslog-enabled {yes | no}

Enables or disables the forwarding of system configuration changes to syslog.

--cee-log-time <date>

Specifies a date after which the audit CEE forwarder will forward protocol access logs. Specify *<date>* in the following format:

```
[protocol]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

--syslog-log-time <date>

Specifies a date after which the audit syslog forwarder will forward logs. To forward SMB, NFS, and HDFS traffic logs, specify `protocol`. To forward configuration change logs, specify `config`. Specify *<date>* in the following format:

```
[protocol|config]@<YYYY>-<MM>-<DD> <HH>:<MM>:<SS>
```

--config-syslog-servers <string>

Specifies a list of remote servers to forward audit configuration change logs to. You must specify the `--config-syslog-servers` for each server to forward audit configuration change logs to.

--clear-config-syslog-servers

Clears the list of remote servers to which audit configuration change logs are forwarded for logging in syslog.

--add-config-syslog-servers <string>

Adds servers to the list of remote servers to which audit configuration change logs are forwarded for logging in syslog. You must specify the `--add-config-syslog-servers` option for each additional server to add.

--remove-config-syslog-servers <string>

Removes servers from the list of remote servers to which audit configuration change logs are forwarded for logging in syslog. You must specify the `--remove-config-syslog-servers` option for each server to remove.

--protocol-syslog-servers <string>

Specifies a list of remote servers to which audit protocol logs are forwarded. You must specify the `--protocol-syslog-servers` option for each server to remove.

`--clear-protocol-syslog-servers`

Clears the list of remote servers to which audit protocol logs are forwarded for logging in syslog.

`--add-protocol-syslog-servers <string>`

Adds servers to the list of remote servers to which audit protocol logs are forwarded for logging in syslog. You must specify the `--add-protocol-syslog-servers` option for each server to add.

`--remove-protocol-syslog-servers <string>`

Removes servers from the list of remote servers to which audit protocol logs are forwarded for logging in syslog. You must specify the `--remove-protocol-syslog-servers` option for each server to remove.

`{--verbose | -v}`

Displays the results of running the command.

isi audit settings global view

Displays global audit settings configured on your cluster.

Syntax

```
isi audit settings global view
```

Options

There are no options for this command.

Examples

The following command displays the audit settings configured on the cluster:

```
isi audit settings global view
```

The system displays output similar to the following text:

```
Protocol Auditing Enabled: Yes
  Audited Zones: System, zoneA
  CEE Server URIs: http://example.com:12228/cee
  Hostname: mycluster
Config Auditing Enabled: Yes
Config Syslog Enabled: Yes
```

isi audit settings modify

Enables you to set filters within an access zone for protocol event types that fail or succeed in an access zone, and to specify which event types to forward to syslog.

Syntax

```
isi audit settings modify
  [--audit-failure <event types>]
  [--clear-audit-failure]
  [--add-audit-failure <event types>]
  [--remove-audit-failure <event types>]
  [--audit-success <event types>]
  [--clear-audit-success]
  [--add-audit-success <event types>]
  [--remove-audit-success <event types>]
```



```
[--syslog-audit-events <event types>]
[--clear-syslog-audit-events]
[--add-syslog-audit-events <event types>]
[--remove-syslog-audit-events <event types>]
[--syslog-forwarding-enabled {yes | no}]
[--zone<access zone>]
[--verbose]
```

Options

--audit-failure <event types>

Specifies one or more filters, separated by commas, for auditing protocol event types that failed. The following event types are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

This option overwrites the current list of filtered event types. To add or remove filters without affecting the current list, configure settings with `--add-audit-failure` or `--remove-audit-failure`.

--clear-audit-failure

Clears all filters for auditing protocol event types that failed.

--add-audit-failure <event types>

Adds one or more filters, separated by commas, for auditing protocol event types that failed. Valid event type values are the same as for `--audit-failure`.

--remove-audit-failure <event types>

Removes one or more filters, separated by commas, for auditing protocol event types that failed. Valid event type values are the same as for `--audit-failure`.

--audit-success <event types>

Specifies one or more filters, separated by commas, for auditing protocol event types that succeeded. The following event types are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

This option overwrites the current list of filtered event types. To add or remove filters without affecting the current list, configure settings with `--add-audit-success` or `--remove-audit-success`.

--clear-audit-success

Clears all filters for auditing protocol event types that succeeded.

--add-audit-success <event types>

Adds one or more filters, separated by commas, for auditing protocol event types that succeeded. Valid event type values are the same as for `--audit-success`.

--remove-audit-success <event types>

Removes one or more filters, separated by commas, for auditing protocol event types that succeeded. Valid event type values are the same as for `--audit-success`.

--syslog-audit-events <event types>

Specifies the auditing protocol event types to forward to syslog. Only those events that match both the `syslog-audit-events` and `--audit-success` or `--audit-failure` will be forwarded to syslog. The following event types are valid:

- close
- create
- delete
- get_security
- logoff
- logon
- read
- rename
- set_security
- tree_connect
- write
- all

This option overwrites the current list of forwarded event types. To add or remove event types without affecting the current list, configure settings with `--add-syslog-audit-events` or `--remove-syslog-audit-events`.

--clear-syslog-audit-events

Clears all auditing protocol event types that are forwarded to syslog.

--add-syslog-audit-events <event types>

Adds one or more auditing protocol event types, separated by commas, that are forwarded to syslog. Valid event type values are the same as for `--syslog-audit-events`.

--remove-syslog-audit-events <event types>

Removes one or more auditing protocol event types, separated by commas, that are forwarded to syslog. Valid event type values are the same as for `--syslog-audit-events`.

--syslog-forwarding-enabled {yes | no}


Enables or disables syslog forwarding audit events in the specified access zone.

--zone <access zones>

Specifies the access zone to which event type filters are applied or forwarded to syslog.

{--verbose | -v}

Displays the results of running the command.

 **NOTE: Each audited event consumes system resources; you should only log events that are supported by your auditing application.**

isi audit settings view

Displays audit filter settings in an access zone and whether syslog forwarding is enabled.

Syntax

```
isi audit settings view  
[--zone<access zone>]
```

Options

--zone<access zone>

Specifies the name of the access zone to view.

Examples

The following command displays the audit settings configured in the zoneA access zone:

```
isi audit settings view --zone=zoneA
```

The system displays output similar to the following text:

```
Audit Failure: create, delete, rename, set_security, close
Audit Success: create, delete, rename, set_security, close
Syslog Audit Events: close
Syslog Forwarding Enabled: No
```

isi audit topics list

Displays a list of configured audit topics, which are internal collections of audit data.

Syntax

```
isi audit topics list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi audit topics modify

Modifies the properties of an audit topic.

Syntax

```
isi audit topics modify <name>
  [--max-cached-messages <integer>]
  [--verbose]
```

Options

<name>

Specifies the name of the audit topic to modify. Valid values are `protocol` and `config`.

--max-cached-messages <integer>

Specifies the maximum number of audit messages to cache before writing them to a persistent store. The larger the number, the more efficiently audit events can be processed. If you specify **0**, each audit event is sent synchronously.

{--verbose | -v}

Displays the results of running the command.

isi audit topics view

Displays the properties of an audit topic.

Syntax

```
isi audit topics view <name>
```

Options

<name>

Specifies the name of the audit topic whose properties you want to view. Valid values are `protocol` and `config`.

isi auth access

Lists the permissions that a user has to access a given file or directory.

Syntax

```
isi auth access {<user> | --uid <integer> | --sid <string>} <path>
  [--zone <string>]
  [--share <string>]
  [--numeric]
  [--verbose]
```

Options

<user>

Specifies the user name.

--sid <string>
Specifies the user by SID.

--uid <integer>
Specifies the user by UID

<path>
Specifies the path of the file or directory under /ifs.

--zone <string>
Specifies the access zone.

--share <string>
Specifies an SMB share name for which to report share configurations and file/directory access information.

{--numeric | -n}
Displays the numeric identifier of the user.

{--verbose | -v}
Displays more detailed information.

isi auth ads create

Configures an Active Directory provider and joins an Active Directory domain.

Syntax

```
isi auth ads create <name> <user>
  [--machine-account <string>]
  [--instance <string>]
  [--password <string>]
  [--organizational-unit <string>]
  [--kerberos-nfs-spn {yes | no} ]
  [--kerberos-hdfs-spn {yes | no} ]
  [--dns-domain <dns-domain>]
  [--groupnet <groupnet>]
  [--allocate-gids {yes | no}]
  [--allocate-uids {yes | no}]
  [--assume-default-domain {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--domain-offline-alerts {yes | no}]
  [--findable-groups <string>...]
  [--findable-users <string>...]
  [--home-directory-template <path>]
  [--ignore-all-trusts {yes | no}]
  [--ignored-trusted-domains <dns-domain>...]
  [--include-trusted-domains <dns-domain>...]
  [--ldap-sign-and-seal {yes | no}]
  [--login-shell <path>]
  [--lookup-domains <dns-domain>...]
  [--lookup-groups {yes | no}]
  [--lookup-normalize-groups {yes | no}]
  [--lookup-normalize-users {yes | no}]
  [--lookup-users {yes | no}]
  [--machine-password-changes {yes | no}]
  [--machine-password-lifespan <duration>]
  [--node-dc-affinity <string>]
  [--node-dc-affinity-timeout <timestamp>]
  [--nss-enumeration {yes | no}]
  [--restrict-findable {yes | no}]
  [--sfu-support {none | rfc2307}]
  [--store-sfu-mappings {yes | no}]
  [--unfindable-groups <string>...]
  [--unfindable-users <string>...]
  [--verbose]
```

Options

<name>

Specifies the fully-qualified Active Directory domain name, which can be resolved to an IPv4 or an IPv6 address. The domain name will also be used as the provider name.

<user>

Specifies the user name of an account that has permission to join machine accounts to the Active Directory domain.

--machine-account <string>

The machine account name to be used by Active Directory. The default value is the cluster name.

--instance <string>

Sets the Active Directory name for this instance.

--password <string>

Specifies the password of the provided user account. If you omit this option, you will be prompted to supply a password.

--organizational-unit <string>

Specifies the name of the organizational unit (OU) to connect to on the Active Directory server. Specify the OU in the form **OuName** or **OuName1/SubName2**.

--kerberos-nfs-spn {yes | no}

Specifies whether to add SPNs for using Kerberized NFS.

--kerberos-hdfs-spn {yes | no}

Specifies whether to add SPNs for using Kerberized HDFS.

--dns-domain <dns-domain>

Specifies a DNS search domain to use instead of the domain that is specified in the **--name** setting.

--groupnet <groupnet>

Specifies the groupnet referenced by the Active Directory provider. The groupnet is a top-level networking container that manages hostname resolution against DNS nameservers and contains subnets and IP address pools. The groupnet specifies which networking properties the Active Directory provider will use when communicating with external servers.

--allocate-gids {yes | no}

Enables or disables GID allocation for unmapped Active Directory groups. Active Directory groups without GIDs can be proactively assigned a GID by the ID mapper. If this option is disabled, GIDs are not proactively assigned, but when a user's primary group does not include a GID, the system may allocate one.

--allocate-uids {yes | no}

Enables or disables UID allocation for unmapped Active Directory users. Active Directory users without UIDs can be proactively assigned a UID by the ID mapper. If this option is disabled, UIDs are not proactively assigned, but when a user's identity does not include a UID, the system may allocate one.

--assume-default-domain {yes | no}

Enables lookup of unqualified user names in the primary domain.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format **<integer>{Y|M|W|D|H|m|s}**.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time that a user logs in, if a home directory does not already exist for the user.

--domain-offline-alerts {yes | no}

Specifies whether to send an alert if the domain goes offline. If this option is set to **yes**, notifications are sent as specified in the global notification rules. The default value is **no**.

--findable-groups <string>...

Specifies a list of groups that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--findable-users <string>...

Specifies a list of users that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--home-directory-template <path>

Specifies the template path to use when creating home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

(i) NOTE: If you are using Active Directory with Services for UNIX (SFU), spaces in Windows-created directory names are converted to underscores for UNIX compatibility.

--ignore-all-trusts {yes | no}

Specifies whether to ignore all trusted domains.

--ignored-trusted-domains <dns-domain>...

Specifies a list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

--include-trusted-domains <dns-domain>...

Specifies a list of trusted domain to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

--ldap-sign-and-seal {yes | no}

Specifies whether to use encryption and signing for LDAP requests to a domain controller.

--login-shell <path>

Specifies the full path to the login shell to use if the Active Directory server does not provide login-shell information. This setting applies only to users who access the file system through SSH.

--lookup-domains <string>...

Specifies a list of domains to which user and group lookups are to be limited. Repeat this option to specify multiple list items.

--lookup-groups {yes | no}

Specifies whether to look up Active Directory groups in other providers before allocating a GID.

--lookup-normalize-groups {yes | no}

Specifies whether to normalize Active Directory group names to lowercase before looking them up.

--lookup-normalize-users {yes | no}

Specifies whether to normalize Active Directory user names to lowercase before looking them up.

--lookup-users {yes | no}

Specifies whether to look up Active Directory users in other providers before allocating a UID.

--machine-password-changes {yes | no}

Specifies whether to enable periodic changes of the machine account password for security purposes.

--machine-password-lifespan <duration>

Sets the maximum age of the machine account password, in the format `<integer>{Y|M|W|D|H|m|s}`.

{--node-dc-affinity | -x} <string>

Specifies the domain controller that the node should exclusively communicate with (affinitize to). This option should be used with a timeout value, which is configured using the `--node-dc-affinity-timeout` option. Otherwise, the default timeout value of 30 minutes is assigned.

(i) NOTE: This setting is for debugging purposes and should be left unconfigured during normal operation. To disable this feature, use a timeout value of 0.

{--node-dc-affinity-timeout} <timestamp>

Specifies the timeout setting for the local node affinity to a domain controller, using the date format `<YYYY>-<MM>-<DD>` or the date/time format `<YYYY>-<MM>-<DD>T<hh>:<mm>[:<ss>]`.

(i) NOTE: A value of 0 disables the affinity. When affinitization is disabled, communication with the specified domain controller may not end immediately. It may persist until another domain controller can be chosen.

--nss-enumeration {yes | no}

Specifies whether to allow the Active Directory provider to respond to getpwent and getgrent requests.

--restrict-findable {yes | no}

Specifies whether to check the authentication provider for filtered lists of findable and unfindable users and groups.

--sfu-support {none | rfc2307}

Specifies whether to support RFC 2307 attributes for Windows domain controllers. RFC 2307 is required for Windows UNIX Integration and for Services For UNIX (SFU) technologies.

--store-sfu-mappings {yes | no}

Specifies whether to store SFU mappings permanently in the ID mapper.

--unfindable-groups <string>...

Specifies a list of groups that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.

--unfindable-users <string>...

Specifies a list of users that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.

{--verbose | -v}

Displays the results of running the command.

isi auth ads delete

Deletes an Active Directory provider, which includes leaving the Active Directory domain that the provider is joined to. Leaving an Active Directory domain disrupts service for users who are accessing the domain. After you leave an Active Directory domain, users can no longer access the domain from the cluster.

Syntax

```
isi auth ads delete <provider-name>
  [--force]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the provider to delete.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays the results of running the command.

Examples

To leave an Active Directory domain named some.domain.org and delete the authentication provider that is associated with it, run the following command:

```
isi auth ads delete some.domain.org
```

At the confirmation prompt, type **y**.

isi auth ads list

Displays a list of Active Directory providers.

Syntax

```
isi auth ads list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To view a list of all the Active Directory providers that the cluster is joined to, run the following command:

```
isi auth ads list
```

The system displays a list of authentication providers.

isi auth ads modify

Modifies an Active Directory authentication provider.

Syntax

```
isi auth ads modify <provider-name>
  [--reset-schannel {yes | no}]
  [--domain-controller <string>]
  [--authentication {yes | no}]
  [--allocate-gids {yes | no}]
  [--allocate-uids {yes | no}]
  [--assume-default-domain {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--domain-offline-alerts {yes | no}]
  [--findable-groups <string>...]
  [--clear-findable-groups]
  [--add-findable-groups <string>...]
  [--remove-findable-groups <string>...]
  [--findable-users <string>...]
```

```

[--clear-findable-users]
[--add-findable-users <string>...]
[--remove-findable-users <string>...]
[--home-directory-template <path>]
[--ignore-all-trusts {yes | no}]
[--ignored-trusted-domains <dns-domain>]
[--clear-ignored-trusted-domains]
[--add-ignored-trusted-domains <dns-domain>]
[--remove-ignored-trusted-domains <dns-domain>]
[--include-trusted-domains <dns-domain>]
[--clear-include-trusted-domains]
[--add-include-trusted-domains <dns-domain>]
[--remove-include-trusted-domains <dns-domain>]
[--ldap-sign-and-seal {yes | no}]
[--login-shell <path>]
[--lookup-domains <dns-domain>]
[--clear-lookup-domains]
[--add-lookup-domains <dns-domain>]
[--remove-lookup-domains <dns-domain>]
[--lookup-groups {yes | no}]
[--lookup-normalize-groups {yes | no}]
[--lookup-normalize-users {yes | no}]
[--lookup-users {yes | no}]
[--machine-password-changes {yes | no}]
[--machine-password-lifespan <duration>]
[--node-dc-affinity <string>]
[--node-dc-affinity-timeout <timestamp>]
[--nss-enumeration {yes | no}]
[--restrict-findable {yes | no}]
[--sfu-support {none | rfc2307}]
[--store-sfu-mappings {yes | no}]
[--unfindable-groups <string>...]
[--clear-unfindable-groups]
[--add-unfindable-groups <string>...]
[--remove-unfindable-groups <string>...]
[--unfindable-users <string>...]
[--clear-unfindable-users]
[--add-unfindable-users <string>...]
[--remove-unfindable-users <string>...]
[--verbose]

```

Options

<provider-name>

Specifies the domain name that the Active Directory provider is joined to, which is also the Active Directory provider name.

--reset-schannel {yes | no}

Resets the secure channel to the primary domain.

--domain-controller <dns-domain>

Specifies a domain controller.

--authentication {yes | no}

Enables the use of the provider for authentication and identity.

--allocate-gids {yes | no}

Enables or disables GID allocation for unmapped Active Directory groups. Active Directory groups without GIDs can be proactively assigned a GID by the ID mapper. If this option is disabled, GIDs are not assigned proactively, but when a user's primary group does not include a GID, the system may allocate one.

--allocate-uids {yes | no}

Enables or disables UID allocation for unmapped Active Directory users. Active Directory users without UIDs can be proactively assigned a UID by the ID mapper. If this option is disabled, UIDs are not assigned proactively, but when a user's identity does not include a UID, the system may allocate one.

--assume-default-domain {yes | no}

Enables lookup of unqualified user names in the primary domain.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format `<integer>{Y|M|W|D|H|m|s}`.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--domain-offline-alerts {yes | no}

Specifies whether to send an alert if the domain goes offline. If this option is set to `yes`, notifications are sent as specified in the global notification rules. The default value is `no`.

--findable-groups <string>...

Specifies a list of groups that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--clear-findable-groups

Removes all entries from the list of findable groups.

--add-findable-groups <string>...

Adds an entry to the list of groups that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--remove-findable-groups <string>...

Removes an entry from the list of groups that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--findable-users <string>...

Specifies a list of users that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--clear-findable-users

Removes all entries from the list of findable users.

--add-findable-users <string>...

Adds an entry to the list of users that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--remove-findable-users <string>...

Removes an entry from the list of users that can be resolved by this authentication provider. Repeat this option to specify multiple list items.

--home-directory-template <path>

Specifies the template path to use when creating home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

 **NOTE: If you are using Active Directory with Services for UNIX (SFU), spaces in Windows-created directory names are converted to underscores for UNIX compatibility.**

--ignore-all-trusts {yes | no}

Specifies whether to ignore all trusted domains.

--ignored-trusted-domains <dns-domain>

Specifies a list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

--clear-ignored-trusted-domains

Clears the list of ignored trusted domains if `--ignore-all-trusts` is disabled.

--add-ignored-trusted-domains <dns-domain>

Adds a domain to the list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

--remove-ignored-trusted-domains <dns-domain>

Removes a specified domain from the list of trusted domains to ignore if `--ignore-all-trusts` is disabled. Repeat this option to specify multiple list items.

--include-trusted-domains <dns-domain>

Specifies a list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--clear-include-trusted-domains`

Clears the list of trusted domains to include if `--ignore-all-trusts` is enabled.

`--add-include-trusted-domains <dns-domain>`

Adds a domain to the list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--remove-include-trusted-domains <dns-domain>`

Removes a specified domain from the list of trusted domains to include if `--ignore-all-trusts` is enabled. Repeat this option to specify multiple list items.

`--ldap-sign-and-seal {yes | no}`

Specifies whether to use encryption and signing on LDAP requests to a domain controller.

`--login-shell <path>`

Specifies the path to the login shell to use if the Active Directory server does not provide login-shell information. This setting applies only to users who access the file system through SSH.

`--lookup-domains <string>`

Specifies a list of domains to which user and group lookups are to be limited. Repeat this option to specify multiple list items.

`--clear-lookup-domains`

Clears the list of restricted domains for user and group lookups.

`--add-lookup-domains <string>`

Adds an entry to the restricted list of domains to use for user and group lookups. Repeat this option to specify multiple list items.

`--remove-lookup-domains <string>`

Removes an entry from the list of domains to use for user and group lookups. Repeat this option to specify multiple list items.

`--lookup-groups {yes | no}`

Specifies whether to look up Active Directory groups in other providers before allocating a GID.

`--lookup-normalize-groups {yes | no}`

Specifies whether to normalize Active Directory group names to lowercase before looking them up.

`--lookup-normalize-users {yes | no}`

Specifies whether to normalize Active Directory user names to lowercase before looking them up.

`--lookup-users {yes | no}`

Specifies whether to look up Active Directory users in other providers before allocating a UID.

`--machine-password-changes {yes | no}`


Specifies whether to enable periodic changes of the machine account password for security purposes.

`--machine-password-lifespan <duration>`

Sets the maximum age of the machine account password, in the format `<integer>{Y|M|W|D|H|m|s}`.


`{--node-dc-affinity | -x} <string>`

Specifies the domain controller that the node should exclusively communicate with (affinitize). This option should be used with a timeout value, which is configured using the `--node-dc-affinity-timeout` option. Otherwise, the default timeout value of 30 minutes is assigned.

 NOTE: This setting is for debugging purposes and should be left unconfigured during normal operation. To disable this feature, use a timeout value of 0.

`{--node-dc-affinity-timeout} <timestamp>`

Specifies the timeout setting for the local node affinity to a domain controller, using the date format `<YYYY>-<MM>-<DD>` or the date/time format `<YYYY>-<MM>-<DD>T<hh>:<mm>[:<ss>]`.

 NOTE: A value of 0 disables the affinity. When affinitization is disabled, communication with the specified domain controller may not end immediately. It may persist until another domain controller can be chosen.

- nss-enumeration {yes | no}**
Specifies whether to allow the Active Directory provider to respond to getpwent and getgrent requests.
- restrict-findable {yes | no}**
Specifies whether to check the authentication provider for filtered lists of findable and unfindable users and groups.
- sfu-support {none | rfc2307}**
Specifies whether to support RFC 2307 attributes for domain controllers. RFC 2307 is required for Windows UNIX Integration and for Services For UNIX (SFU) technologies.
- store-sfu-mappings {yes | no}**
Specifies whether to store SFU mappings permanently in the ID mapper.
- unfindable-groups <string>...**
Specifies a list of groups that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- clear-unfindable-groups**
Removes all entries from the list of unfindable groups.
- add-unfindable-groups <string>...**
Adds an entry to the list of groups that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- remove-unfindable-groups <string>...**
Removes an entry from the list of groups that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- unfindable-users <string>...**
Specifies a list of users that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- clear-unfindable-users**
Removes all entries from the list of unfindable users.
- add-unfindable-users <string>...**
Adds an entry to the list of users that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- remove-unfindable-users <string>...**
Removes an entry from the list of users that cannot be resolved by this authentication provider. Repeat this option to specify multiple list items.
- {--verbose | -v}**
Displays the results of running the command.

isi auth ads spn check

Checks valid service principal names (SPNs).

Syntax

```
isi auth ads spn check <provider-name>
```

Options

- <provider-name>**
Specifies the Active Directory provider name.

isi auth ads spn create

Adds one or more service principal names (SPNs) for a machine account. SPNs must be propagated to all domain controllers to make them available to clients.

Syntax

```
isi auth ads spn create <provider-name> <spn>
  [--user <string>]
  [--password <string>]
```

Options

<provider-name>

Specifies the Active Directory provider name.

<spn>

Specifies the service principal name.

{--user | -U} <string>

Specifies an administrative user account name with permission to create SPNs in the Active Directory domain.

{--password | -P} <string>

Specifies the administrative user account password.

isi auth ads spn delete

Deletes one or more SPNs that are registered against a machine account.

Syntax

```
isi auth ads spn delete <provider-name> <spn>
  [--user <string>]
  [--password <string>]
```

Options

<provider-name>

Specifies the Active Directory provider name.

<spn>

Specifies the service principal name.

{--user | -U} <string>

Specifies an administrative user account name with permission modify SPNs in the Active Directory domain.

{--password | -P} <string>

Specifies the administrative user account password.

isi auth ads spn fix

Adds missing service principal names (SPNs) for an Active Directory provider.

Syntax

```
isi auth ads spn fix <provider-name>
  [--spn <string>]
  [--user <string>]
  [--password <string>]
  [--noremove
```

Options

<provider-name>

Specifies the Active Directory provider name.

--spn <string>

Specifies the service principal name.

--user <string>

Specifies an administrative user account name with permission to add SPNs for the Active Directory domain.

--password <string>

Specifies the administrative user account password.

--noremove

Specifies not to remove any unexpected SPNs, but to add missing SPNs.

isi auth ads spn list

Displays a list of service principal names (SPNs) that are registered against a machine account.

Syntax

```
isi auth ads spn list <provider-name>
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<provider-name>

Specifies the Active Directory provider name.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth ads trusts controllers list

Displays a list of domain controllers for a trusted domain.

Syntax

```
isi auth ads trusts controllers list <provider-name>
  [--dc-site <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<provider-name> <string>

Specifies an Active Directory provider.

--dc-site

Specifies a domain controller site.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

The following command displays a list of trusted domains in an Active Directory provider named ad.isilon.com:

```
isi auth ads trusts controllers list ad.isilon.com
```

isi auth ads trusts list

Displays a list of trusted domains.

Syntax

```
isi auth ads trusts list <provider-name>
```


Options

`<provider-name>`

Specifies an Active Directory provider.

isi auth ads view

Displays the properties of an Active Directory provider.

Syntax

```
isi auth ads view <provider-name>
[--verbose]
```

Options

`<provider-name>`

Specifies the name of the provider to view.

`{--verbose | -v}`

Displays more detailed information.

isi auth duo modify

Modify the Duo provider settings.

Syntax

```
isi auth duo modify
  [--autopush <boolean>]
  [--enabled <boolean>]
  [--failmode <string>]
  [--fallback-local-ip <boolean>]
  [--groups <string>]
  [--http-proxy <string>]
  [--https-timeout <integer>]
  [--prompts <integer>]
  [--pushinfo <boolean>]
  [--host <string>]
  [--ikey <string>]
  [--set-skey]
  [--verbose | -v]
  [--help | -h]
```

Options

`--autopush <boolean>`

Specifies if Duo automatically sends a push login request to the users phone.

`--enabled <boolean>`

Specifies if the Duo provider is enabled.

`--failmode <string>`

Specifies if Duo will fail "safe" (allow access) or "secure" (deny access) on configuration or service errors.

`--fallback-local-ip <boolean>`

Specifies if Duo will report the server IP if the client IP cannot be detected.

`--groups <string>`

Specifies a list of groups Duo is required for (default is all groups).

--http-proxy <string>

Specifies the HTTP proxy to use.

--https-timeout <integer>

Specifies the number of seconds to wait for HTTPS responses from Duo Security.

--prompts <integer>

Specifies the maximum number of authentication prompts to display (1-3 default is 3).

--pushinfo <boolean>

Specifies to include information in the Duo Push message.

--host <string>

Specifies the API hostname.

--ikey <string>

Specifies the integration key.

--set-skey

Specify the secret key.

--verbose | **-v**

Displays more detailed information.

--help | **-h**

Displays help for this command.

isi auth error

Displays error code definitions from the authentication log files.

Syntax

```
isi auth error <error-code>
```

Options

<error-code>

Specifies the error code to convert.

Examples

To view the definition of error code 4, run the following command:

```
isi auth error 4
```

The system displays output similar to the following example:

```
4 = ERROR_TOO_MANY_OPEN_FILES
```

isi auth file create

Creates a file provider.

Syntax

```
isi auth file create <name>  
[--password-file <path>]
```

```

[--group-file <path>]
[--authentication {yes | no}]
[--create-home-directory {yes | no}]
[--enabled {yes | no}]
[--enumerate-groups {yes | no}]
[--enumerate-users {yes | no}]
[--findable-groups <string>]
[--findable-users <string>]
[--group-domain <string>]
[--home-directory-template <path>]
[--listable-groups <string>]
[--listable-users <string>]
[--login-shell <path>]
[--modifiable-groups <string>]
[--modifiable-users <string>]
[--netgroup-file <path>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--ntlm-support {all | v2only | none}]
[--provider-domain <string>]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--restrict-modifiable {yes | no}]
[--unfindable-groups <string>]
[--unfindable-users <string>]
[--unlistable-groups <string>]
[--unlistable-users <string>]
[--unmodifiable-groups <string>]
[--unmodifiable-users <string>]
[--user-domain <string>]
[--verbose]

```

Options

<name>

Sets the file provider name.

--password-file <path>

Specifies the path to a `passwd.db` replacement file.

--group-file <path>

Specifies the path to a `group` replacement file.

--authentication {yes | no}

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--enabled {yes | no}

Enables or disables the provider.

--findable-groups <string>

Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable group. If populated, groups that are not included in this list cannot be resolved.

--findable-users <string>

Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable user. If populated, users that are not included in this list cannot be resolved.

--group-domain <string>

Specifies the domain that this provider will use to qualify groups. The default group domain is `FILE_GROUPS`.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time.

that represent specific variables. For example, %U, %D, and %Z are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--listable-groups <string>

Specifies a group that can be listed if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be listed.

--listable-users <string>

Specifies a user that can be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be listed.

--login-shell <path>

Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

--modifiable-groups <string>

Specifies a group that can be modified in this provider if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be modified.

--modifiable-users <string>

Specifies a user that can be modified in this provider if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be modified.

--netgroup-file <path>

Specifies the path to a `netgroup` replacement file.

--normalize-groups {yes | no}

Normalizes group names to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user names to lowercase before lookup.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM.

--provider-domain <string>

Specifies the domain that the provider will use to qualify user and group names.

--restrict-findable {yes | no}

Specifies whether to check the provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check the provider for filtered lists of listable and unlistable users and groups.

--restrict-modifiable {yes | no}

Specifies whether to check the provider for filtered lists of modifiable and unmodifiable users and groups.

--unfindable-groups <string>

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a group that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unfindable-users <string>

If `--restrict-findable` is enabled and the findable users list is empty, specifies a user that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unlistable-groups <string>

If `--restrict-listable` is enabled and the listable groups list is empty, specifies a group that cannot be listed by this provider. Repeat this option to specify multiple list items.

--unlistable-users <string>

If `--restrict-listable` is enabled and the listable users list is empty, specifies a user that cannot be listed by this provider. Repeat this option to specify multiple list items.

--unmodifiable-groups <string>

If `--restrict-modifiable` is enabled and the modifiable groups list is empty, specifies a group that cannot be modified. Repeat this option to specify multiple list items.

--unmodifiable-users <string>

If `--restrict-modifiable` is enabled and the modifiable users list is empty, specifies a user that cannot be modified. Repeat this option to specify multiple list items.

`--user-domain <string>`

Specifies the domain that this provider will use to qualify users. The default user domain is `FILE_USERS`.

`{--verbose | -v}`

Displays more detailed information.

isi auth file delete

Deletes a file provider.

Syntax

```
isi auth file delete <provider-name>
  [--force]
  [--verbose]
```

Options

`<provider-name>`

Specifies the name of the provider to delete.

`{--force | -f}`

Suppresses command-line prompts and messages.

`{--verbose | -v}`

Displays more detailed information.

isi auth file list

Displays a list of file providers.

Syntax

```
isi auth file list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`{--limit | -l} <integer>`

Displays no more than the specified number of items.

`--format {table | json | csv | list}`

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

`{--no-header | -a}`

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi auth file modify

Modifies a file provider.

Syntax

```
isi auth file modify <provider-name>
  [--name <string>]
  [--password-file <path>]
  [--group-file <path>]
  [--authentication {yes | no}]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
  [--findable-users <string>]
  [--clear-findable-users]
  [--add-findable-users <string>]
  [--remove-findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <path>]
  [--listable-groups <string>]
  [--clear-listable-groups]
  [--add-listable-groups <string>]
  [--remove-listable-groups <string>]
  [--listable-users <string>]
  [--clear-listable-users]
  [--add-listable-users <string>]
  [--remove-listable-users <string>]
  [--login-shell <path>]
  [--modifiable-groups <string>]
  [--clear-modifiable-groups]
  [--add-modifiable-groups <string>]
  [--remove-modifiable-groups <string>]
  [--modifiable-users <string>]
  [--clear-modifiable-users]
  [--add-modifiable-users <string>]
  [--remove-modifiable-users <string>]
  [--netgroup-file <path>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--ntlm-support {all | v2only | none}]
  [--provider-domain <string>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--restrict-modifiable {yes | no}]
  [--unfindable-groups <string>]
  [--clear-unfindable-groups]
  [--add-unfindable-groups <string>]
  [--remove-unfindable-groups <string>]
  [--unfindable-users <string>]
  [--clear-unfindable-users]
  [--add-unfindable-users <string>]
  [--remove-unfindable-users <string>]
  [--unlistable-groups <string>]
  [--clear-unlistable-groups]
  [--add-unlistable-groups <string>]
  [--remove-unlistable-groups <string>]
  [--unlistable-users <string>]
  [--clear-unlistable-users]
  [--add-unlistable-users <string>]
  [--remove-unlistable-users <string>]
  [--unmodifiable-groups <string>]
```

```

[--clear-unmodifiable-groups]
[--add-unmodifiable-groups <string>]
[--remove-unmodifiable-groups <string>]
[--unmodifiable-users <string>]
[--clear-unmodifiable-users]
[--add-unmodifiable-users <string>]
[--remove-unmodifiable-users <string>]
[--user-domain <string>]
[--verbose]

```

Options

<provider-name>

Specifies the name of the file provider to modify. This setting cannot be modified.

--name <string>

Specifies an new name for the authentication provider.

--password-file <path>

Specifies the path to a `passwd.db` replacement file.

--group-file <path>

Specifies the path to a `group` replacement file.

--authentication {yes | no}

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

--cache-entry-expiry <duration>

Specifies the length of time after which the cache entry will expire, in the format `<integer>[Y | M | W | D | H | m | s]`. To turn off cache expiration, set this value to `off`.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--enabled {yes | no}

Enables or disables the provider.

--enumerate-groups {yes | no}

Specifies whether to allow the provider to enumerate groups.

--enumerate-users {yes | no}

Specifies whether to allow the provider to enumerate users.

--findable-groups <string>

Specifies a group that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be resolved. This option overwrites any existing entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

--clear-findable-groups

Removes all entries from the list of findable groups.

--add-findable-groups <string>

Adds an entry to the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-groups <string>

Removes an entry from the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--findable-users <string>

Specifies a user that can be found in the provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be resolved. This option overwrites any existing entries in the findable users list; to add or remove users without affecting current entries, use `--add-findable-users` or `--remove-findable-users`.

--clear-findable-users

Removes all entries from the list of findable users.

--add-findable-users <string>

Adds an entry to the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-users <string>

Removes an entry from the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--group-domain <string>

Specifies the domain that the provider will use to qualify groups. The default group domain is `FILE_GROUPS`.

--group-file <path>

Specifies the path to a group replacement file.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--listable-groups <string>

Specifies a group that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be viewed. This option overwrites any existing entries in the listable groups list; to add or remove groups without affecting current entries, use `--add-listable-groups` or `--remove-listable-groups`.

--clear-listable-groups

Removes all entries from the list of viewable groups.

--add-listable-groups <string>

Adds an entry to the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-groups <string>

Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--listable-users <string>

Specifies a user that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be viewed. This option overwrites any existing entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

--clear-listable-users

Removes all entries from the list of viewable users.

--add-listable-users <string>

Adds an entry to the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-users <string>

Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--login-shell <path>

Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

--modifiable-groups <string>

Specifies a group that can be modified if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any groups that are not included in this list cannot be modified. This option overwrites any existing entries in the modifiable groups list; to add or remove groups without affecting current entries, use `--add-modifiable-groups` or `--remove-modifiable-groups`.

--clear-modifiable-groups

Removes all entries from the list of modifiable groups.

--add-modifiable-groups <string>

Adds an entry to the list of modifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--remove-modifiable-groups <string>

Removes an entry from the list of modifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--modifiable-users <string>

Specifies a user that can be modified if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items. If populated, any users that are not included in this list cannot be modified. This option overwrites any existing entries in the modifiable users list; to add or remove users without affecting current entries, use `--add-modifiable-users` or `--remove-modifiable-users`.

--clear-modifiable-users

Removes all entries from the list of modifiable users.

--add-modifiable-users <string>

Adds an entry to the list of modifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--remove-modifiable-users <string>

Removes an entry from the list of modifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--netgroup-file <path>

Specifies the path to a `netgroup` replacement file.

--normalize-groups {yes | no}

Normalizes group names to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user names to lowercase before lookup.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM and is recommended.

--password-file <path>

Specifies the path to a `passwd.db` replacement file.

--provider-domain <string>

Specifies the domain that this provider will use to qualify user and group names.

--restrict-findable {yes | no}

Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

--restrict-modifiable {yes | no}

Specifies whether to check this provider for filtered lists of modifiable and unmodifiable users and groups.

--unfindable-groups <string>

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a group that cannot be resolved by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unfindable groups list; to add or remove groups without affecting current entries, use `--add-unfindable-groups` or `--remove-unfindable-groups`.

--clear-unfindable-groups

Removes all entries from the list of unfindable groups.

--add-unfindable-groups <string>

Adds an entry to the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-groups <string>

Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unfindable-users <string>

If `--restrict-findable` is enabled and the findable users list is empty, specifies a user that cannot be resolved by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

--clear-unfindable-users

Removes all entries from the list of unfindable groups.

--add-unfindable-users <string>

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-users <string>

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unlistable-groups <string>

If `--restrict-listable` is enabled and the viewable groups list is empty, specifies a group that cannot be listed by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

--clear-unlistable-groups

Removes all entries from the list of unviewable groups.

--add-unlistable-groups <string>

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-groups <string>

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--unlistable-users <string>

If `--restrict-listable` is enabled and the viewable users list is empty, specifies a user that cannot be listed by this provider. Repeat this option to specify multiple list items. This option overwrites any existing entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

--clear-unlistable-users

Removes all entries from the list of unviewable users.

--add-unlistable-users <string>

Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-users <string>

Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--unmodifiable-groups <string>

If `--restrict-modifiable` is enabled and the modifiable groups list is empty, specifies a group that cannot be modified. Repeat this option to specify multiple list items. This option overwrites any existing entries in this provider's unmodifiable groups list; to add or remove groups without affecting current entries, use `--add-unmodifiable-groups` or `--remove-unmodifiable-groups`.

--clear-unmodifiable-groups

Removes all entries from the list of unmodifiable groups.

--add-unmodifiable-groups <string>

Adds an entry to the list of unmodifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--remove-unmodifiable-groups <string>

Removes an entry from the list of unmodifiable groups that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--unmodifiable-users <string>

If `--restrict-modifiable` is enabled and the modifiable users list is empty, specifies a user that cannot be modified. Repeat this option to specify multiple list items. This option overwrites any existing entries in this provider's unmodifiable users list; to add or remove users without affecting current entries, use `--add-unmodifiable-users` or `--remove-unmodifiable-users`.

--clear-unmodifiable-users

Removes all entries from the list of unmodifiable users.

--add-unmodifiable-users <string>

Adds an entry to the list of unmodifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--remove-unmodifiable-users <string>

Removes an entry from the list of unmodifiable users that is checked if `--restrict-modifiable` is enabled. Repeat this option to specify multiple list items.

--user-domain <string>

Specifies the domain that this provider will use to qualify users. The default user domain is `FILE_USERS`.

{--verbose | -v}

Displays detailed information.

isi auth file view

Displays the properties of a file provider.

Syntax

```
isi auth file view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth groups create

Creates a local group.

Syntax

```
isi auth groups create <name>
[--gid <integer>]
[--add-user <name>]
[--add-group <name>]
[--add-gid <integer>]
[--add-uid <integer>]
[--add-sid <string>]
[--add-wellknown <name>]
[--zone <string>]
[--provider <string>]
[--verbose]
[--force]
```

Options

<name>

Specifies the group name.

--gid <integer>

Overrides automatic allocation of the UNIX group identifier (GID) with the specified value. Setting this option is not recommended.

--add-user <name>

Specifies the name of the user to add to the group. Repeat this option to specify multiple users.

--add-group <name>

Specifies the name of the group to add to this group. Repeat this option to specify multiple groups.

--add-gid <integer>

Specifies the GID of the group to add to this group. Repeat this option to specify multiple groups.

--add-uid <integer>

Specifies the UID of the user to add to the group. Repeat this option to specify multiple users.

--add-sid <string>

Specifies the Windows security identifier (SID) of the persona to add to the group, for example S-1-5-21-13. Repeat this option to specify multiple SIDs.

--add-wellknown <name>

Specifies a wellknown persona name to add to the group. Repeat this option to specify multiple personas.

--zone <string>

Specifies the access zone in which to create the group.

--provider <string>

Specifies a local authentication provider in the specified access zone.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Suppresses command-line prompts and messages.

isi auth groups delete

Removes a local group from the system. Members of a group are removed before the group is deleted.

Syntax

```
isi auth groups delete {<group> | --gid <integer> | --sid <string>}
  [--zone <string>]
  [--provider <string>]
  [--force]
  [--verbose]
```

Options

This command requires <group>, --gid <integer>, or --sid <string>.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GID.

<group>

--sid <string>

Specifies the group by SID.

--zone <string>

Specifies the name of the access zone that contains the group.

--provider <string>

Specifies the group's authentication provider.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays the results of running the command.

isi auth groups flush

Flushes cached group information.

Syntax

```
isi auth groups flush
```

Options

There are no options for this command.

Examples

To flush all cached group information, run the following command:

```
isi auth groups flush
```

isi auth groups list

Displays a list of groups.

Syntax

```
isi auth groups list  
  [--domain <string>]  
  [--zone <string>]  
  [--provider <string>]  
  [--limit <integer>]  
  [--format {table | json | csv | list}]  
  [--no-header]  
  [--no-footer]  
  [--verbose]
```

Options

--domain <string>

Specifies the provider domain.

--zone <string>

Specifies an access zone.

--provider <string>

Specifies an authentication provider.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth groups members list

Displays a list of members that are associated with a group.

Syntax

```
isi auth groups members list {<group> | --gid <integer> | --sid <string>}
  [--zone <string>]
  [--provider <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

This command requires *<group>*, *--gid <integer>*, or *--sid <string>*.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GID.

--sid <string>

Specifies the group by SID.

--zone <string>

Specifies an access zone.

--provider <string>

Specifies an authentication provider.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth groups modify

Modifies a local group.

Syntax

```
isi auth groups modify {<group> | --gid <integer> | --sid <string>}
  [--new-gid <integer>]
  [--add-uid <integer>]
  [--remove-uid <integer>]
  [--add-user <name>]
  [--add-group <name>]
  [--add-gid <integer>]
  [--remove-user <name>]
  [--remove-group <name>]
  [--remove-gid <integer>]
  [--add-sid <string>]
  [--remove-sid <string>]
  [--add-wellknown <name>]
  [--remove-wellknown <name>]
  [--zone <string>]
  [--provider <string>]
  [--verbose]
  [--force]
```

Options

This command requires *<group>*, *--gid <integer>*, or *--sid <string>*.

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GID.

--sid <string>

Specifies the group by security identifier (SID).

--new-gid <integer>

Specifies a new GID for the group. Setting this option is not recommended.

--add-uid <integer>

Specifies the UID of a user to add to the group. Repeat this option to specify multiple list items.

--remove-uid <integer>

Specifies the UID of a user to remove from the group. Repeat this option to specify multiple list items.

--add-user <name>

Specifies the name of a user to add to the group. Repeat this option to specify multiple list items.

--remove-user <name>

Specifies the name of a user to remove from the group. Repeat this option to specify multiple list items.

--add-group <name>

Specifies the name of the group to add to this group. Repeat this option to specify multiple groups to add.

--remove-group <name>

Specifies the name of the group to remove from this group. Repeat this option to specify multiple groups to remove.

--add-gid <integer>

Specifies the GID of the group to add to this group. Repeat this option to specify multiple GIDs to add.

--remove-gid <integer>

Specifies the GID of the group to remove from this group. Repeat this option to specify multiple GIDs to remove.

--add-sid <string>

Specifies the SID of an object to add to the group, for example S-1-5-21-13. Repeat this option to specify multiple list items.

--remove-sid <string>

Specifies the SID of an object to remove from the group. Repeat this option to specify multiple list items.

--add-wellknown <name>

Specifies a well-known SID to add to the group. Repeat this option to specify multiple list items.

--remove-wellknown <name>

Specifies a well-known SID to remove from the group. Repeat this option to specify multiple list items.

--zone <string>

Specifies the group's access zone.

--provider <string>

Specifies the group's authentication provider in the format `type:instance`. Valid provider types are as follows.

- ads
- file
- ldap
- local
- nis

For example, `ldap:auth1` specifies an LDAP provider named `auth1`.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Suppresses command-line prompts and messages.

isi auth groups view

Displays the properties of a group, including historical security identifier (SID) information.

Syntax

```
isi auth groups view {<group> | --gid <integer> | --sid <string>}
  [--cached]
  [--provider <string>]
  [--show-groups]
  [--zone <string>]
```

Options

<group>

Specifies the group by name.

--gid <integer>

Specifies the group by GID.

--sid <string>

Specifies the group by SID.

--cached

Displays cached information.

--provider <string>

Specifies the name of an authentication provider.

--show-groups

Displays groups that include this group as a member.

--zone <string>

Specifies an access zone.

isi auth id

Displays your access token.

Syntax

```
isi auth id
```

Options

There are no options for this command.

isi auth krb5 create

Creates an MIT Kerberos provider and joins a user to an MIT Kerberos realm.

Syntax

```
isi auth krb5 create <realm> {<user> | --keytab-file <string> }  
    [--password <string>]  
    [--spn <string>]  
    [--groupnet <groupnet>]  
    [--is-default-realm {yes | no}]  
    [--kdc <string>]  
    [--admin-server <string>]  
    [--default-domain <string>]  
    [--verbose]
```

Options

<realm>

Specifies the Kerberos realm name.

<user>

Specifies the name of a user with permission to create service principal names (SPNs) in the Kerberos realm.

--keytab-file <string>

Specifies the keytab file to import.

--password <string>

Specifies the password used for joining a Kerberos realm.

--spn <string>

Specifies the SPNs to register. Specify **--spn** for each additional SPN that you want to register.

--groupnet <groupnet>

Specifies the groupnet referenced by the Kerberos provider. The groupnet is a top-level networking container that manages hostname resolution against DNS nameservers and contains subnets and IP address pools. The groupnet specifies which networking properties the Kerberos provider will use when communicating with external servers.

--is-default-realm {yes | no}

Specifies whether the Kerberos realm is the default.

--kdc <string>

Specifies the hostname, IPv4 address, or IPv6 address of the Key Distribution Center (KDC). Specify **--kdc** for each additional KDC you want to add to the realm.

--admin-server <string>

Specifies the hostname, IPv4 address, or IPv6 address of the administrative server (master KDC).

--default-domain<string>

Specifies the default Kerberos domain for the Kerberos realm used for translating Kerberos v4 principal names.

{**--verbose** | **-v**}

Displays detailed information.

isi auth krb5 delete

Deletes an MIT Kerberos authentication provider and removes the user from an MIT Kerberos realm.

Syntax

```
isi auth krb5 delete <provider-name>
                    [--force]
```

Options

<provider-name>

Specifies the Kerberos provider name.

{**--force** | **-f**}

Specifies not to ask for a confirmation.

isi auth krb5 domain create

Creates an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain create <domain> <realm>
```

Options

<domain>

Specifies the name of the Kerberos domain.

<realm>

Specifies the name of the Kerberos realm.

isi auth krb5 domain delete

Deletes an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain delete <domain>
                    [--force]
```

Options

`<domain>`

Specifies the name of the Kerberos domain.

`{--force | -f}`

Specifies not to ask for a confirmation.

isi auth krb5 domain list

Displays a list of MIT Kerberos domain mappings.

Syntax

```
isi auth krb5 domain list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

`{--limit | -l} <integer>`

Specifies the number of Kerberos domain mappings to display.

`--format {table | json | csv | list}`

Specifies whether to display the Kerberos domain mappings in a tabular, JSON, CSV, or list formats.

`{--no-header | -a}`

Specifies not to display the headers in the CSV or tabular formats.

`{--no-footer | -z}`

Specifies not to display the table summary footer information.

isi auth krb5 domain modify

Modifies an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain modify <domain>
  [--realm <string>]
```

Options

`<domain>`

Specifies the Kerberos domain name.

`--realm <string>`

Specifies the Kerberos realm name.

isi auth krb5 domain view

Displays the properties of an MIT Kerberos domain mapping.

Syntax

```
isi auth krb5 domain view <domain>
```

Options

<domain>

Specifies the Kerberos domain name.

isi auth krb5 list

Displays a list of MIT Kerberos authentication providers.

Syntax

```
isi auth krb5 list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

{--limit | -l} <integer>

Specifies the number of Kerberos providers to display.

--format {table | json | csv | list}

Specifies to display the Kerberos providers in a tabular, JSON, CSV, or list format.

{--no-header | -a}

Specifies not to display the headers in the CSV or tabular formats.

{--no-footer | -z}

Specifies not to display the table summary footer information.

isi auth krb5 realm create

Creates an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm create <realm>
  [--is-default-realm {yes | no}]
  [--kdc <string>]
  [--admin-server <string>]
  [--default-domain <string>]
```

Options

<realm>

Specifies the name of the Kerberos realm.

--is-default-realm {yes | no}

Specifies whether the Kerberos realm is the default realm.

--kdc <string>

Specifies the hostname, IPv4 address, or IPv6 address of the Key Distribution Center (KDC). Specify **--kdc** for each additional KDC you want to add to the realm.

--admin-server <string>

Specifies the hostname, IPv4 address, or IPv6 address of the administrative server (master KDC).

--default-domain <string>

Specifies the default domain for the realm used for translating the v4 principal names.

isi auth krb5 realm delete

Deletes an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm delete <realm>
    [--force]
```

Options

<realm>

Specifies the Kerberos realm name.

{--force | -f}

Specifies not to ask for a confirmation.

isi auth krb5 realm list

Displays a list of MIT Kerberos realms.

Syntax

```
isi auth krb5 realm list
    [--limit <integer>]
    [--format {table | json | csv | list}]
    [--no-header]
    [--no-footer]
```

Options

{--limit | -l} <integer>

Specifies the number of Kerberos realms to display.

--format {table | json | csv | list}

Specifies whether to display the Kerberos realms in a tabular, JSON, CSV, or list format.

{--no-header | -a}

Specifies not to display the headers in the CSV or tabular formats.

{--no-footer | -z}

Specifies not to display the table summary footer information.

isi auth krb5 realm modify

Modifies an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm modify <realm>
  [--is-default-realm {yes | no}]
  [--kdc <string>]
  [--admin-server <string>]
  [--default-domain <string>]
```

Options

<realm>

Specifies the Kerberos realm name.

--is-default-realm {yes | no}

Specifies whether the Kerberos realm is the default.

--kdc <string>

Specifies the hostname, IPv4 address, or IPv6 address of the Key Distribution Center (KDC). Specify **--kdc** for each additional KDC you want to add to the realm.

--admin-server <string>

Specifies the hostname, IPv4 address, or IPv6 address of the administrative server (master KDC).

--default-domain <string>

Specifies the default domain for the Kerberos realm used for translating v4 principal names.

isi auth krb5 realm view

Displays the properties of an MIT Kerberos realm.

Syntax

```
isi auth krb5 realm view <realm>
```

Options

<realm>

Specifies the Kerberos realm name.

isi auth krb5 spn check

Checks for missing service principal names (SPNs) for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn check <provider-name>
```

Options

<provider-name>

Specifies the Kerberos provider name.

isi auth krb5 spn create

Creates or updates keys for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn create <provider-name> <user> <spn>
  [--password <string>]
```

Options

<provider-name>

Specifies the Kerberos provider name.

<user>

Specifies a user name with permissions to create the service principal names (SPNs) in the Kerberos realm.

<spn>

Specifies the SPN.

--password <string>

Specifies the password used during the modification of a Kerberos realm.

isi auth krb5 spn delete

Deletes keys from an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn delete <provider-name> <spn> {<kvno> | --all}
```

Options

<provider-name>

Specifies the Kerberos provider name.

<spn>

Specifies the service principal name (SPN).

<kvno>

Specifies the key version number.

--all

Deletes all the key versions.

isi auth krb5 spn fix

Adds the missing service principal names (SPNs) for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn fix <provider-name> <user>
    [--password <string>]
    [--force]
```

Options

<provider-name>

Specifies the Kerberos provider name.

<user>

Specifies a user name with permissions to join clients to the given Kerberos domain.

--password <string>

Specifies the password that was used when modifying the Kerberos realm.

{--force | -f}

Specifies not to ask for a confirmation.

isi auth krb5 spn import

Imports keys from a keytab file for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn import <provider-name> <keytab-file>
```

Options

<provider-name>

Specifies the Kerberos provider name.

<keytab-file>

Specifies the keytab file to import.

isi auth krb5 spn list

Lists the service principal names (SPNs) and keys registered for an MIT Kerberos provider.

Syntax

```
isi auth krb5 spn list <provider-name>
    [--limit <integer>]
    [--format {table | json | csv | list}]
    [--no-header]
    [--no-footer]
```


Options

<provider-name>

Specifies the Kerberos provider name.

{--limit | -l} <integer>

Specifies the number of SPNs and keys to display.

--format {table | json | csv | list}

Specifies to display the SPNs and keys in a tabular, JSON, CSV, or list format.

{--no-header | -a}

Specifies not to display the headers in the CSV or tabular formats.

{--no-footer | -z}

Specifies not to display the table summary footer information.

isi auth krb5 view

Displays the properties of an MIT Kerberos authentication provider.

Syntax

```
isi auth krb5 view <provider-name>
```

Options

<provider-name>

Specifies the Kerberos provider name.

isi auth ldap create

Creates an LDAP provider.

Syntax

```
isi auth ldap create <name>
  [--base-dn <string>]
  [--server-uris <string>]
  [--alternate-security-identities-attribute <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--bind-dn <string>]
  [--bind-timeout <integer>]
  [--certificate-authority-file <string>]
  [--check-online-interval <duration>]
  [--cn-attribute <string>]
  [--create-home-directory {yes | no}]
  [--crypt-password-attribute <string>]
  [--email-attribute <string>]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--findable-users <string>]
  [--gecos-attribute <string>]
  [--gid-attribute <string>]
  [--group-base-dn <string>]
  [--group-domain <string>]
  [--group-filter <string>]
  [--group-members-attribute <string>]
```

```

[--group-search-scope <scope>]
[--home-directory-template <string>]
[--homedir-attribute <string>]
[--ignore-tls-errors {yes | no}]
[--listable-groups <string>]
[--listable-users <string>]
[--login-shell <string>]
[--member-lookup-method {default | rfc2307bis}]
[--member-of-attribute <string>]
[--name-attribute <string>]
[--netgroup-base-dn <string>]
[--netgroup-filter <string>]
[--netgroup-members-attribute <string>]
[--netgroup-search-scope <scope>]
[--netgroup-triple-attribute <string>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--nt-password-attribute <string>]
[--ntlm-support {all | v2only | none}]
[--provider-domain <string>]
[--require-secure-connection {yes | no}]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--search-scope <scope>]
[--search-timeout <integer>]
[--shadow-user-filter <string>]
[--shadow-expire-attribute <string>]
[--shadow-flag-attribute <string>]
[--shadow-inactive-attribute <string>]
[--shadow-last-change-attribute <string>]
[--shadow-max-attribute <string>]
[--shadow-min-attribute <string>]
[--shadow-warning-attribute <string>]
[--shell-attribute <string>]
[--ssh-public-key-attribute <string>]
[--uid-attribute <string>]
[--unfindable-groups <string>]
[--unfindable-users <string>]
[--unique-group-members-attribute <string>]
[--unlistable-groups <string>]
[--unlistable-users <string>]
[--user-base-dn <string>]
[--user-domain <string>]
[--user-filter <string>]
[--user-search-scope <scope>]
[--groupnet <groupnet>]
[--template {default | rfc2307 | ad-idmu | ldapsam}]
[--bind-password <string>]
[--set-bind-password]
[--force | -f]
[--verbose] | -v

```

Options

<name>

Sets the LDAP provider name.

--base-dn <string>

Sets the root of the tree in which to search for identities. For example, CN=Users, DC=mycompany, DC=com.

--server-uris <string>

Specifies a list of LDAP server URIs to be used when accessing the server. Repeat this option to specify multiple list items.

Specify the LDAP server URI in the format ldaps://<server>:<port> for secure LDAP or ldap://<server>:<port> for non-secure LDAP.

The server can be specified as an IPv4 address, an IPv6 address, or a hostname.

If you do not specify a port number, the default port is used; 389 for secure LDAP or 636 for non-secure LDAP.



NOTE: If you specify non-secure LDAP, the bind password is transmitted to the server in clear text.

--alternate-security-identities-attribute <string>

Specifies the name to be used when searching for alternate security identities. This name is used when OneFS attempts to resolve a Kerberos principal to a user.

--authentication {yes | no}

Enables or disables the use of the provider for authentication as well as identity. The default value is `yes`.

--balance-servers {yes | no}

Makes the provider connect to a random server on each request.

--bind-dn <string>

Specifies the distinguished name to use when binding to the LDAP server. For example, `CN=myuser, CN=Users, DC=mycompany, DC=com`.

--bind-timeout <integer>

Specifies the timeout in seconds when binding to the LDAP server.

--certificate-authority-file <path>

Specifies the path to the root certificates file.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format `<integer>[Y | M | W | D | H | m | s]`.

--cn-attribute <string>

Specifies the LDAP attribute that contains common names. The default value is `cn`.

--create-home-directory {yes | no}

Specifies whether to automatically create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--crypt-password-attribute <string>

Specifies the LDAP attribute that contains UNIX passwords. This setting has no default value.

--email-attribute <string>

Specifies the LDAP attribute that contains email addresses. The default value is `mail`.

--enabled {yes | no}

Enables or disables the provider.

--enumerate-groups {yes | no}

Specifies whether to allow the provider to enumerate groups.

--enumerate-users {yes | no}

Specifies whether to allow the provider to enumerate users.

--findable-groups <string>

Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable group. If populated, groups that are not included in this list cannot be resolved.

--findable-users <string>

Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify each additional findable user. If populated, users that are not included in this list cannot be resolved.

--gecos-attribute <string>

Specifies the LDAP attribute that contains GECOS fields. The default value is `gecos`.

--gid-attribute <string>

Specifies the LDAP attribute that contains GIDs. The default value is `gidNumber`.

--group-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for groups.

--group-domain <string>

Specifies the domain that the provider will use to qualify groups. The default group domain is `LDAP_GROUPS`.

--group-filter <string>

Sets the LDAP filter for group objects.

--group-members-attribute <string>

Specifies the LDAP attribute that contains group members. The default value is `memberUid`.

--group-search-scope <scope>

Defines the default depth from the base distinguished name (DN) to perform LDAP searches for groups.

The following values are valid:

default

Applies the setting in `--search-scope`.



NOTE: You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree

Searches the base DN and all entries below it.

children

Searches all entries below the base DN, excluding the base DN.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information about home directory variables, see Home directories.

--homedir-attribute <string>

Specifies the LDAP attribute that contains home directories. The default value is `homeDirectory`.

--ignore-tls-errors {yes | no}

Continues over a secure connection even if identity checks fail.

--listable-groups <string>

Specifies a list of groups that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed.

--listable-users <string>

Specifies a list of users that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed.

--login-shell <path>

Specifies the pathname of the user's login shell for users who access the file system through SSH.

--member-lookup-method {default | rfc2307bis}

Sets the method by which group member lookups are performed. Use caution when changing this option directly.

--member-of-attribute <string>

Sets the attribute to be used when searching LDAP for reverse memberships. This LDAP value should be an attribute of the user type `posixAccount` that describes the groups in which the POSIX user is a member.

--name-attribute <string>

Specifies the LDAP attribute that contains UIDs, which are used as login names. The default value is `uid`.

--netgroup-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for netgroups.

--netgroup-filter <string>

Sets the LDAP filter for netgroup objects.

--netgroup-members-attribute <string>

Specifies the LDAP attribute that contains netgroup members. The default value is `memberNisNetgroup`.

--netgroup-search-scope <scope>

Defines the depth from the base distinguished name (DN) to perform LDAP searches for netgroups.

The following values are valid:

default Applies the setting in `--search-scope`.

i **NOTE: You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.**

base Searches only the entry at the base DN.

onelevel Searches all entries exactly one level below the base DN.

subtree Searches the base DN and all entries below it.

children Searches all entries below the base DN, excluding the base DN.

--netgroup-triple-attribute <string>

Specifies the LDAP attribute that contains netgroup triples. The default value is `nisNetgroupTriple`.

--normalize-groups {yes | no}

Normalizes group names to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user names to lowercase before lookup.

--nt-password-attribute <string>

Specifies the LDAP attribute that contains Windows passwords. A commonly used value is `ntpasswdhash`.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support.

--provider-domain <string>

Specifies the domain that the provider will use to qualify user and group names.

--require-secure-connection {yes | no}

Specifies whether to require a TLS connection.

--restrict-findable {yes | no}

Specifies whether to check the provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check the provider for filtered lists of listable and unlistable users and groups.

--search-scope <scope>

Defines the default depth from the base distinguished name (DN) to perform LDAP searches.

The following values are valid:

base Searches only the entry at the base DN.

onelevel Searches all entries exactly one level below the base DN.

subtree Searches the base DN and all entries below it.

children Searches all entries below the base DN, excluding the base DN itself.

--search-timeout <integer>

Specifies the number of seconds after which to stop retrying and fail a search. The default value is 100.

--shadow-user-filter <string>

Sets the LDAP filter for shadow user objects.

--shadow-expire-attribute <string>

Sets the attribute name that indicates the absolute date to expire the account.

--shadow-flag-attribute <string>

Sets the attribute name that indicates the section of the shadow map that is used to store the flag value.

--shadow-inactive-attribute <string>

Sets the attribute name that indicates the number of days of inactivity that is allowed for the user.

--shadow-last-change-attribute <string>

Sets the attribute name that indicates the last change of the shadow information.

--shadow-max-attribute <string>

Sets the attribute name that indicates the maximum number of days that a password can be valid.

--shadow-min-attribute <string>

Sets the attribute name that indicates the minimum number of days between shadow changes.

--shadow-warning-attribute <string>

Sets the attribute name that indicates the number of days before the password expires to warn the user.

--shell-attribute <string>

Specifies the LDAP attribute that contains a user's UNIX login shell. The default value is `loginShell`.

--ssh-public-key-attribute <string>

Sets the attribute name that contains the user's SSH Public Key.

--uid-attribute <string>

Specifies the LDAP attribute that contains UID numbers. The default value is `uidNumber`.

--unfindable-groups <string>

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a list of groups that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unfindable-users <string>

If `--restrict-findable` is enabled and the findable users list is empty, specifies a list of users that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unique-group-members-attribute <string>

Specifies the LDAP attribute that contains unique group members. This attribute is used to determine which groups a user belongs to if the LDAP server is queried by the user's DN instead of the user's name. This setting has no default value.

--unlistable-groups <string>

If `--restrict-listable` is enabled and the listable groups list is empty, specifies a list of groups that cannot be listed by this provider that cannot be viewed. Repeat this option to specify multiple list items.

--unlistable-users <string>

If `--restrict-listable` is enabled and the listable users list is empty, specifies a list of users that cannot be listed by this provider that cannot be viewed. Repeat this option to specify multiple list items.

--user-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for users.

--user-domain <string>

Specifies the domain that the provider will use to qualify users. The default user domain is `LDAP_USERS`.

--user-filter <string>

Sets the LDAP filter for user objects.

--user-search-scope <scope>

Defines the depth from the base distinguished name (DN) to perform LDAP searches for users.

The following values are valid:

default	Applies the search scope that is defined in the default query settings.
base	Searches only the entry at the base DN.
onelevel	Searches all entries exactly one level below the base DN.
subtree	Searches the base DN and all entries below it.
children	Searches all entries below the base DN, excluding the base DN itself.

--groupnet <groupnet>

Specifies the groupnet referenced by the LDAP provider. The groupnet is a top-level networking container that manages hostname resolution against DNS nameservers and contains subnets and IP address pools. The groupnet specifies which networking properties the LDAP provider will use when communicating with external servers.

--template {**default** | **rfc-2307** | **ad-idmu** | **ldapsam**}

Specifies a template to be used to configure the LDAP provider. The templates provide pre-selected attributes. The templates are: RFC 2307, Active Directory Identity Management for UNIX (ad-idmu), and LDAP for Samba (ldapsam).

--bind-password <*string*>

Sets the password for the distinguished name that is used when binding to the LDAP server. To set the password interactively, use the **--set-bind-password** option instead.

--set-bind-password

Interactively sets the password for the distinguished name that is used when binding to the LDAP server. This option cannot be used with **--bind-password**.

[**--force** | **-f**]

Specifies to ignore warnings when creating or modifying an LDAP provider.

{**--verbose** | **-v**}

Displays the results of running the command.

isi auth ldap delete

Deletes an LDAP provider.

Syntax

```
isi auth ldap delete <provider-name>
  [--force]
  [--verbose]
```

Options

<*provider-name*>

Specifies the name of the provider to delete.

{**--force** | **-f**}

Suppresses command-line prompts and messages.

<*provider-name*>

Specifies the name of the provider to delete.

{**--verbose** | **-v**}

Displays more detailed information.

isi auth ldap list

Displays a list of LDAP providers.

Syntax

```
isi auth ldap list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth ldap modify

Modifies an LDAP provider.

Syntax

```
isi auth ldap modify <provider-name>
  [--name <string>]
  [--base-dn <string>]
  [--server-uris <string>]
  [--add-server-uris <string>]
  [--remove-server-uris <string>]
  [--alternate-security-identities-attribute <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--bind-dn <string>]
  [--bind-timeout <integer>]
  [--certificate-authority-file <string>]
  [--check-online-interval <duration>]
  [--cn-attribute <string>]
  [--create-home-directory {yes | no}]
  [--crypt-password-attribute <string>]
  [--email-attribute <string>]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
  [--findable-users <string>]
  [--clear-findable-users]
  [--add-findable-users <string>]
  [--remove-findable-users <string>]
  [--gecos-attribute <string>]
  [--gid-attribute <string>]
  [--group-base-dn <string>]
  [--group-domain <string>]
  [--group-filter <string>]
  [--group-members-attribute <string>]
  [--group-search-scope <scope>]
  [--homedir-attribute <string>]
  [--home-directory-template <string>]
  [--ignore-tls-errors {yes | no}]
  [--listable-groups <string>]
  [--clear-listable-groups]
  [--add-listable-groups <string>]
  [--remove-listable-groups <string>]
  [--listable-users <string>]
```



```

[--clear-listable-users]
[--add-listable-users <string>]
[--remove-listable-users <string>]
[--login-shell <string>]
[--member-lookup-method {default | rfc2307bis}]
[--member-of-attribute <string>]
[--name-attribute <string>]
[--netgroup-base-dn <string>]
[--netgroup-filter <string>]
[--netgroup-members-attribute <string>]
[--netgroup-search-scope <scope>]
[--netgroup-triple-attribute <string>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--nt-password-attribute <string>]
[--ntlm-support {all | v2only | none}]
[--provider-domain <string>]
[--require-secure-connection {yes | no}]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--search-scope <scope>]
[--search-timeout <integer>]
[--shadow-user-filter <string>]
[--shadow-expire-attribute <string>]
[--shadow-flag-attribute <string>]
[--shadow-inactive-attribute <string>]
[--shadow-last-change-attribute <string>]
[--shadow-max-attribute <string>]
[--shadow-min-attribute <string>]
[--shadow-warning-attribute <string>]
[--shell-attribute <string>]
[--ssh-public-key-attribute <string>]
[--uid-attribute <string>]
[--unfindable-groups <string>]
[--clear-unfindable-groups]
[--add-unfindable-groups <string>]
[--remove-unfindable-groups <string>]
[--unfindable-users <string>]
[--clear-unfindable-users]
[--add-unfindable-users <string>]
[--remove-unfindable-users <string>]
[--unique-group-members-attribute <string>]
[--unlistable-groups <string>]
[--clear-unlistable-groups]
[--add-unlistable-groups <string>]
[--remove-unlistable-groups <string>]
[--unlistable-users <string>]
[--clear-unlistable-users]
[--add-unlistable-users <string>]
[--remove-unlistable-users <string>]
[--user-base-dn <string>]
[--user-domain <string>]
[--user-filter <string>]
[--user-search-scope <scope>]
[--template {default | rfc2307 | ad-idmu | ldapsam}]
[--bind-password <string>]
[--set-bind-password]
[--force | -f]
[--verbose | -v]

```

Options

<provider-name>

Specifies the name of the LDAP provider to modify.

--name <string>

Specifies an new name for the authentication provider.

--base-dn <string>

Sets the root of the tree in which to search for identities. For example, CN=Users, DC=mycompany, DC=com.

--server-uris <string>

Specifies a list of LDAP server URIs to be used when accessing the server. Repeat this option to specify multiple list items.

Specify the LDAP server URI in the format `ldaps://<server>:<port>` for secure LDAP or `ldap://<server>:<port>` for non-secure LDAP.

The server can be specified as an IPv4 address, an IPv6 address, or a hostname.

If you do not specify a port number, the default port is used; 389 for secure LDAP or 636 for non-secure LDAP.



NOTE: If you specify non-secure LDAP, the bind password is transmitted to the server in clear text.

--add-server-uris <string>

Adds an entry to the list of server URIs. Repeat this option to specify multiple list items.

The server to be added can be specified as an IPv4 address, an IPv6 address, or a hostname.

--remove-server-uris <string>

Removes an entry from the list of server URIs. Repeat this option to specify multiple list items.

The server to be removed can be specified as an IPv4 address, an IPv6 address, or a hostname.

--alternate-security-identities-attribute <string>

Specifies the name to be used when searching for alternate security identities. This name is used when OneFS attempts to resolve a Kerberos principal to a user.

--authentication {yes | no}

Enables or disables the use of this provider for authentication as well as identity. The default value is `yes`.

--balance-servers {yes | no}

Makes this provider connect to a random server on each request.

--bind-dn <string>

Specifies the distinguished name to use when binding to the LDAP server. For example, `CN=myuser,CN=Users,DC=mycompany,DC=com`.

--bind-timeout <integer>

Specifies the timeout in seconds when binding to the LDAP server.

--certificate-authority-file <path>

Specifies the path to the root certificates file.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format `<integer>[Y | M | W | D | H | m | s]`.

--cn-attribute <string>

Specifies the LDAP attribute that contains common names. The default value is `cn`.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user. The directory path is specified in the path template through the `--home-directory-template` command.

--crypt-password-attribute <string>

Specifies the LDAP attribute that contains UNIX passwords. This setting has no default value.

--email-attribute <string>

Specifies the LDAP attribute that contains email addresses. The default value is `mail`.

--enabled {yes | no}

Enables or disables this provider.

--enumerate-groups {yes | no}

Specifies whether to allow the provider to enumerate groups.

--enumerate-users {yes | no}

Specifies whether to allow the provider to enumerate users.

--findable-groups <string>

Specifies a list of groups that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved in this provider. This option overwrites the entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

--clear-findable-groups

Removes the list of findable groups.

--add-findable-groups <string>

Adds an entry to the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-groups <string>

Removes an entry from the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--findable-users <string>

Specifies a list of users that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be resolved in this provider. This option overwrites the entries in the findable users list; to add or remove users without affecting current entries, use `--add-findable-users` or `--remove-findable-users`.

--clear-findable-users

Removes the list of findable users.

--add-findable-users <string>

Adds an entry to the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-users <string>

Removes an entry from the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--gecos-attribute <string>

Specifies the LDAP attribute that contains GECOS fields. The default value is `gecos`.

--gid-attribute <string>

Specifies the LDAP attribute that contains GIDs. The default value is `gidNumber`.

--group-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for groups.

--group-domain <string>

Specifies the domain that this provider will use to qualify groups. The default group domain is `LDAP_GROUPS`.

--group-filter <string>

Sets the LDAP filter for group objects.

--group-members-attribute <string>

Specifies the LDAP attribute that contains group members. The default value is `memberUid`.

--group-search-scope <scope>

Defines the default depth from the base distinguished name (DN) to perform LDAP searches for groups.

The following values are valid:

default Applies the setting in `--search-scope`.



You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base Searches only the entry at the base DN.

onelevel Searches all entries exactly one level below the base DN.

subtree Searches the base DN and all entries below it.

children Searches all entries below the base DN, excluding the base DN.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--homedir-attribute <string>

Specifies the LDAP attribute that is used when searching for the home directory. The default value is `homeDirectory`.

--ignore-tls-errors {yes | no}

Continues over a secure connection even if identity checks fail.

--listable-groups <string>

Specifies a list of groups that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed in this provider. This option overwrites the entries in the listable groups list; to add or remove groups without affecting current entries, use `--add-listable-groups` or `--remove-listable-groups`.

--clear-listable-groups

Removes all entries from the list of viewable groups.

--add-listable-groups <string>

Adds an entry to the list of listable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-groups <string>

Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--listable-users <string>

Specifies a list of users that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed in this provider. This option overwrites the entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

--clear-listable-users

Removes all entries from the list of viewable users.

--add-listable-users <string>

Adds an entry to the list of listable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-users <string>

Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--login-shell <path>

Specifies the pathname to the user's login shell, for users who access the file system through SSH.

--member-lookup-method {default | rfc2307bis}

Sets the method by which group member lookups are performed. Use caution when changing this option directly.

--member-of-attribute <string>

Sets the attribute to be used when searching LDAP for reverse memberships. This LDAP value should be an attribute of the user type `posixAccount` that describes the groups in which the POSIX user is a member.

--name-attribute <string>

Specifies the LDAP attribute that contains UIDs, which are used as login names. The default value is `uid`.

--netgroup-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for netgroups.

--netgroup-filter <string>

Sets the LDAP filter for netgroup objects.

--netgroup-members-attribute <string>

Specifies the LDAP attribute that contains netgroup members. The default value is `memberNisNetgroup`.

--netgroup-search-scope <scope>

Defines the depth from the base distinguished name (DN) to perform LDAP searches for netgroups.

The following values are valid:

default Applies the setting in `--search-scope`.



You cannot specify `--search-scope=default`. For example, if you specify `--group-search-scope=default`, the search scope is set to the value of `--search-scope`.

base Searches only the entry at the base DN.

onelevel Searches all entries exactly one level below the base DN.

subtree Searches the base DN and all entries below it.

children Searches all entries below the base DN, excluding the base DN.

--netgroup-triple-attribute <string>

Specifies the LDAP attribute that contains netgroup triples. The default value is `nisNetgroupTriple`.

--normalize-groups {yes | no}

Normalizes group names to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user names to lowercase before lookup.

--nt-password-attribute <string>

Specifies the LDAP attribute that contains Windows passwords. A commonly used value is `ntpasswdhash`.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support.

The following values are valid:

`all`

`v2only`

`none`

--provider-domain <string>

Specifies the domain that this provider will use to qualify user and group names.

--require-secure-connection {yes | no}

Specifies whether to require a TLS connection.

--restrict-findable {yes | no}

Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

--search-scope <scope>

Defines the default depth from the base distinguished name (DN) to perform LDAP searches.

The following values are valid:

base Searches only the entry at the base DN.

onelevel Searches all entries exactly one level below the base DN.

subtree Searches the base DN and all entries below it.

children Searches all entries below the base DN, excluding the base DN itself.

--search-timeout <integer>

Specifies the number of seconds after which to stop retrying and fail a search. The default value is **100**.

--shadow-user-filter *<string>*

Sets the LDAP filter for shadow user objects.

--shadow-expire-attribute *<string>*

Sets the attribute name that indicates the absolute date to expire the account.

--shadow-flag-attribute *<string>*

Sets the attribute name that indicates the section of the shadow map that is used to store the flag value.

--shadow-inactive-attribute *<string>*

Sets the attribute name that indicates the number of days of inactivity that is allowed for the user.

--shadow-last-change-attribute *<string>*

Sets the attribute name that indicates the last change of the shadow information.

--shadow-max-attribute *<string>*

Sets the attribute name that indicates the maximum number of days a password can be valid.

--shadow-min-attribute *<string>*

Sets the attribute name that indicates the minimum number of days between shadow changes.

--shadow-warning-attribute *<string>*

Sets the attribute name that indicates the number of days before the password expires to warn the user.

--shell-attribute *<string>*

Specifies the LDAP attribute that is used when searching for a user's UNIX login shell. The default value is `loginShell`.

--ssh-public-key-attribute *<string>*

Sets the attribute name used that contains the user's SSH Public Key.

--uid-attribute *<string>*

Specifies the LDAP attribute that contains UID numbers. The default value is `uidNumber`.

--unfindable-groups *<string>*

Specifies a group that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable groups list; to add or remove groups without affecting current entries, use `--add-unfindable-groups` or `--remove-unfindable-groups`.

--clear-unfindable-groups

Removes all entries from the list of unfindable groups.

--add-unfindable-groups *<string>*

Adds an entry to the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-groups *<string>*

Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unfindable-users *<string>*

Specifies a user that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

--clear-unfindable-users

Removes all entries from the list of unfindable groups.

--add-unfindable-users *<string>*

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-users *<string>*

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unique-group-members-attribute <string>

Specifies the LDAP attribute that contains unique group members. This attribute is used to determine which groups a user belongs to if the LDAP server is queried by the user's DN instead of the user's name. This setting has no default value.

--unlistable-groups <string>

Specifies a group that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

--clear-unlistable-groups

Removes all entries from the list of unviewable groups.

--add-unlistable-groups <string>

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-groups <string>

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--unlistable-users <string>

Specifies a user that cannot be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

--clear-unlistable-users

Removes all entries from the list of unviewable users.

--add-unlistable-users <string>

Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-users <string>

Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--user-base-dn <string>

Specifies the distinguished name of the entry at which to start LDAP searches for users.

--user-domain <string>

Specifies the domain that this provider will use to qualify users. The default user domain is `LDAP_USERS`.

--user-filter <string>

Sets the LDAP filter for user objects.

--user-search-scope <scope>

Defines the depth from the base distinguished name (DN) to perform LDAP searches for users. The valid values are as follows:

The following values are valid:

default

Applies the setting in `--search-scope`.



You cannot specify `--search-scope=default`. For example, if you specify `--user-search-scope=default`, the search scope is set to the value of `--search-scope`.

base

Searches only the entry at the base DN.

onelevel

Searches all entries exactly one level below the base DN.

subtree	Searches the base DN and all entries below it.
children	Searches all entries below the base DN, excluding the base DN.

--template {default | rfc-2307 | ad-idmu | ldapsam}

Specifies a template to be used to configure the LDAP provider. The templates provide pre-selected attributes. The templates provide pre-selected attributes. The templates are: RFC 2307, Active Directory Identity Management for UNIX (ad-idmu), and LDAP for Samba (ldapsam).

--bind-password <string>

Sets the password for the distinguished name that is used when binding to the LDAP server. To set the password interactively, use the `--set-bind-password` option instead.

--set-bind-password

Interactively sets the password for the distinguished name that is used when binding to the LDAP server. This option cannot be used with `--bind-password`.

{--force | -f}

Specifies to ignore warnings when creating or modifying an LDAP provider.

{--verbose | -v}

Displays detailed information.

isi auth ldap view

Displays the properties of an LDAP provider.

Syntax

```
isi auth ldap view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth local list

Displays a list of local providers.

Syntax

```
isi auth local list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth local modify

Modifies a local provider.

Syntax

```
isi auth local modify <provider-name>
  [--authentication {yes | no}]
  [--create-home-directory {yes | no}]
  [--home-directory-template <string>]
  [--lockout-duration <duration>]
  [--lockout-threshold <integer>]
  [--lockout-window <duration>]
  [--login-shell <string>]
  [--machine-name <string>]
  [--min-password-age <duration>]
  [--max-password-age <duration>]
  [--min-password-length <integer>]
  [--password-prompt-time <duration>]
  [--password-complexity {lowercase | uppercase |
    numeric | symbol}]
  [--clear-password-complexity]
  [--add-password-complexity {lowercase | uppercase |
    numeric | symbol}]
  [--remove-password-complexity <string>]
  [--password-history-length <integer>]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the local provider to modify.

--authentication {yes | no}

Uses the provider for authentication as well as identity. The default setting is yes.

--create-home-directory {yes | no}

Creates a home directory the first time a user logs in.

--home-directory-template <string>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--lockout-duration <duration>

Sets the length of time that an account will be inaccessible after multiple failed login attempts.

--lockout-threshold <integer>

Specifies the number of failed login attempts after which an account will be locked out.

--lockout-window <duration>

Sets the time in which the number of failed attempts specified by the `--lockout-threshold` option must be made for an account to be locked out. Duration is specified in the format `<integer>[{Y | M | W | D | H | m | s}]`.

--login-shell <string>

Specifies the path to the UNIX login shell.

--machine-name <string>

Specifies the domain to use to qualify user and group names for the provider.

--min-password-age <duration>

Sets the minimum password age, in the format <integer>[*{Y | M | W | D | H | m | s}*].

--max-password-age <duration>

Sets the maximum password age, in the format <integer>[*{Y | M | W | D | H | m | s}*].

--min-password-length <integer>

Sets the minimum password length.

--password-prompt-time <duration>

Sets the remaining time until a user is prompted for a password change, in the format <integer>[*{Y | M | W | D | H | m | s}*].

--password-complexity {lowercase | uppercase | numeric | symbol}

Specifies the conditions that a password is required to meet. A password must contain at least one character from each specified option to be valid. For example, if `lowercase` and `numeric` are specified, a password must contain at least one lowercase character and one digit to be valid. Symbols are valid, excluding `#` and `@`.

--clear-password-complexity

Clears the list of parameters against which to validate new passwords.

--add-password-complexity {lowercase | uppercase | numeric | symbol}

Adds items to the list of parameters against which to validate new passwords. Repeat this command to specify additional password-complexity options.

--remove-password-complexity <string>

Removes items from the list of parameters against which to validate new passwords. Repeat this command to specify each password-complexity option that you want to remove.

--password-history-length <integer>

Specifies the number of previous passwords to store to prevent reuse of a previous password. The maximum password history length is 24.

{--verbose | -v}

Displays more detailed information.

isi auth local view

Displays the properties of a local provider.

Syntax

```
isi auth local view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth log-level modify

Specifies the logging level for the authentication service on the node.

Syntax

```
isi auth log-level modify <level>
[--verbose]
```

Options

<level>

Sets the log level for the current node. The log level determines how much information is logged.

The following values are valid and are organized from least to most information:

- always
- error
- warning
- info
- verbose
- debug
- trace

i **NOTE:** Levels **verbose**, **debug**, and **trace** may cause performance issues. Levels **debug** and **trace** log information that likely will be useful only when consulting EMC Isilon Technical Support.

{--verbose | -v}

Displays detailed information.

isi auth log-level view

Displays the logging level for the authentication service on the node.

Syntax

```
isi auth log-level view
```

Options

There are no options for this command.

isi auth mapping create

Creates a manual mapping between a source identity and target identity or automatically generates a mapping for a source identity.

Syntax

```
isi auth mapping create {<source>| --source-uid <integer>
| --source-gid <integer> | --source-sid <string>}
[{{--uid | --gid | --sid}}]
[--on-disk]
[--2way]
```

```
[{--target <string> | --target-uid <integer>
 | --target-gid <string> | --target-sid <string>}]
[--zone<string>]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--source-uid <integer>

Specifies the mapping source by UID.

--source-gid <integer>

Specifies the mapping source by GID.

--source-sid <string>

Specifies the mapping source by SID.

--uid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped UID.

--gid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped GID.

--sid

Generates a mapping if one does not exist for the identity; otherwise, retrieves the mapped SID.

--on-disk

Specifies that the source on-disk identity should be represented by the target identity.

--2way

Specifies a two-way, or reverse, mapping.

--target <string>

Specifies the mapping target by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--target-uid <integer>

Specifies the mapping target by UID.

--target-gid <integer>

Specifies the mapping target by GID.

--target-sid <string>

Specifies the mapping target by SID.

--zone<string>

Specifies the access zone that the ID mapping is applied to. If no access zone is specified, the mapping is applied to the default System zone.

isi auth mapping delete

Deletes one or more identity mappings.

Syntax

```
isi auth mapping delete {<source>| --source-uid <integer>
 | --source-gid <integer> | --source-sid <string> | --all}
[{-only-generated | --only-external | --2way | --target <string>
 | --target-uid <integer> | --target-gid <integer> | --target-sid <string>}]
[--zone<string>]
[{-force | -f}]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--source-uid <integer>

Specifies the mapping source by UID.

--source-gid <integer>

Specifies the mapping source by GID.

--source-sid <string>

Specifies the mapping source by SID.

--all

Deletes all identity mappings in the specified access zone. Can be used in conjunction with **--only-generated** and **--only-external** for additional filtering.

--only-generated

Only deletes identity mappings that were created automatically and that include a generated UID or GID from the internal range of user and group IDs. Must be used in conjunction with **--all**.

--only-external

Only deletes identity mappings that were created automatically and that include a UID or GID from an external authentication source. Must be used in conjunction with **--all**.

--2way

Specifies or deletes a two-way, or reverse, mapping.

--target <string>

Specifies the mapping target by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--target-uid <integer>

Specifies the mapping target by UID.

--target-gid <integer>

Specifies the mapping target by GID.

--target-sid <string>

Specifies the mapping target by SID.

--zone <string>

Deletes identity mappings in the specified access zone. If no access zone is specified, mappings are deleted from the default System zone.

{--force | -f}

Does not prompt you for confirmation that you want to delete the mapping.

isi auth mapping dump

Displays or prints the kernel mapping database.

Syntax

```
isi auth mapping dump <file>
  [--zone <string>]
  [--verbose | -v]
```

Options

If no option is specified, the full kernel mapping database is displayed.

<file>

Dumps the database to the specified local file, where `<file>` is an absolute path within the `/ifs` file system. If the file does not exist, it will be created. If the file does exist, it will be overwritten.

`--zone <string>`

Displays the database from the specified access zone. If no access zone is specified, displays all mappings.

Examples

To view the kernel mapping database, run the following command:

```
isi auth mapping dump
```

The system displays output similar to the following example:

```
["ZID:1", "UID:6299", [{"SID:S-1-5-21-1195855716-1407", 128}]]
["ZID:1", "GID:1000000", [{"SID:S-1-5-21-1195855716-513", 48}]]
["ZID:1", "SID:S-1-5-21-1195855716-1407", [{"UID:6299", 144}]]
["ZID:1", "SID:S-1-5-21-1195855716-513", [{"GID:1000000", 32}]]
```

isi auth mapping flush

Flushes the cache for one or all identity mappings. Flushing the cache might be useful if the ID mapping rules have been modified.

Syntax

```
isi auth mapping flush [--all | --source <string>
  | --source-uid <integer> | --source-gid <integer>
  | --source-sid <string>]
[--zone<string>]
```

Options

You must specify either `--all` or one of the source options.

`--all`

Flushes all identity mappings on the cluster.

`--source <string>`

Specifies the mapping source by identity type, in the format `<type>:<value>`—for example, **UID:2002**.

`--source-uid <integer>`

Specifies the source identity by UID.

`--source-gid <integer>`

Specifies the source identity by GID.

`--source-sid <string>`

Specifies the source identity by SID.

`--zone<string>`

Specifies the access zone of the source identity. If no access zone is specified, any mapping for the specified source identity is flushed from the default System zone.

isi auth mapping import

Imports mappings from a source file to the ID mapping database.

Syntax

```
isi auth mapping import <file>
[--replace | -r]
[--verbose | -v]
```

Options

<file>

Specifies the full path to the file to import. File content must be in the same format as the output that is displayed by running the `isi auth mapping dump` command. File must specify an absolute path within the `/ifs` file system.

{--replace | -r}

Overwrites existing entries in the mapping database file with the file content.

{--verbose | -v}

Displays detailed information.

isi auth mapping list

Displays the ID mapping database for an access zone.

Syntax

```
isi auth mapping list
[--zone <string>]
```

Options

--zone <string>

Specifies an access zone.

isi auth mapping modify

Sets or modifies a mapping between two identities.

Syntax

```
isi auth mapping modify {<source>| --source-uid <integer>
| --source-gid <integer> | --source-sid <string> | --target <string>
| --target-uid <integer> | --target-gid <string> | --target-sid <string>}
[--on-disk]
[--2way]
[--zone<string>]
```

Options

<source>

Specifies the mapping source by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--source-uid *<integer>*

Specifies the mapping source by UID.

--source-gid *<integer>*

Specifies the mapping source by GID.

--source-sid *<string>*

Specifies the mapping source by SID.

--target *<string>*

Specifies the mapping target by identity type, in the format *<type>:<value>*—for example, **UID:2002**.

--target-uid *<integer>*

Specifies the mapping target by UID.

--target-gid *<integer>*

Specifies the mapping target by GID.

--target-sid *<string>*

Specifies the mapping target by SID.

--on-disk

Specifies that the source on-disk identity should be represented by the target identity.

--2way

Specifies a two-way, or reverse, mapping.

--zone *<string>*

Specifies the access zone that the ID mapping is applied to. If no access zone is specified, the mapping is applied to the default System zone.

isi auth mapping token

Displays the access token that is calculated for a user during authentication, including the user's historical groups.

Syntax

```
isi auth mapping token {<user> | --uid <integer>
  | --kerberos-principal <string>}
  [--zone <string>]
  [--primary-gid <integer>]
  [--gid <integer>]
```

Options

This command requires *<user>* or **--uid** *<integer>* or **--kerberos-principal** *<string>*.

<user>

Specifies the user by name.

--uid *<integer>*

Specifies the user by UID.

--kerberos-principal *<string>*

Specifies the Kerberos principal by name. For example, `user@realm.com`.

--zone *<string>*

Specifies the name of the access zone that contains the mapping.

--primary-gid *<integer>*

Specifies the primary GID.

--gid *<integer>*

Specifies a token GID. Repeat this option to specify multiple GIDs.

isi auth mapping view

Displays mappings for an identity.

Syntax

```
isi auth mapping view {<id>| --uid <integer>
| --gid <integer> | --sid <string>}
[--nocreate]
[--zone <string>]
```

Options

- <id>**
Specifies the ID of the source identity type in the format *<type>:<value>*—for example, **UID:2002**.
- uid <integer>**
Specifies the mapping source by UID.
- gid <integer>**
Specifies the mapping source by GID.
- sid <string>**
Specifies the mapping source by SID.
- nocreate**
Specifies that nonexistent mappings should not be created.
- zone**
Specifies the access zone of the source identity. If no access zone is specified, OneFS displays mappings from the default System zone.

Examples

The following command displays mappings for a user whose UID is 2002 in the zone3 access zone:

```
isi auth mapping view uid:2002 --zone=zone3
```

The system displays output similar to the following example:

Type	Mapping
Name	test1
On-disk	UID:2002
Unix UID	2002
Unix GID	None
SMB	S-1-5-21-1776575851-2890035977-2418728619-1004
NFSv4	test1

isi auth netgroups view

Displays information about a netgroup.

Syntax

```
isi auth netgroups view <netgroup>
[--zone <string>]
[--provider <string>]
```

```
[--recursive {true | false}]
[--ignore-errors {true | false}]
```

Options

<netgroup>

Specifies the netgroup name.

--zone <string>

Specifies the access zone.

--provider <string>

Specifies the authentication provider.

--recursive {true | false}

Specifies whether to recursively resolve nested netgroups. The default value is `true`.

--ignore {true | false}

Specifies whether to ignore errors and unresolvable netgroups. The default value is `false`.

isi auth nis create

Creates an NIS provider.

Syntax

```
isi auth nis create <name>
  [--nis-domain <string>]
  [--servers <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--findable-users <string>]
  [--group-domain <string>]
  [--home-directory-template <path>]
  [--hostname-lookup {yes | no}]
  [--listable-groups <string>]
  [--listable-users <string>]
  [--login-shell <path>]
  [--normalize-groups {yes | no}]
  [--normalize-users {yes | no}]
  [--provider-domain <string>]
  [--ntlm-support {all | v2only | none}]
  [--request-timeout <integer>]
  [--restrict-findable {yes | no}]
  [--restrict-listable {yes | no}]
  [--retry-time <integer>]
  [--unfindable-groups <string>]
  [--unfindable-users <string>]
  [--unlistable-groups <string>]
  [--unlistable-users <string>]
  [--user-domain <string>]
  [--ypmatch-using-tcp {yes | no}]
  [--groupnet <groupnet>]
  [--verbose]
```

Options

<name>

Sets the name of the NIS provider.

--nis-domain <string>

Specifies the NIS domain name.

--servers <string>

Specifies a list of NIS servers to be used by this provider. Specify the NIS server as an IPv4 address or hostname. Repeat this option to specify multiple list items.

--authentication {yes | no}

Enables or disables the use of the provider for authentication as well as identity. The default value is *yes*.

--balance-servers {yes | no}

Makes the provider connect to a random server on each request.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format *<integer>[{Y | M | W | D | H | m | s}]*.

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--enabled {yes | no}

Enables or disables the provider.

--enumerate-groups {yes | no}

Specifies whether to allow the provider to enumerate groups.

--enumerate-users {yes | no}

Specifies whether to allow the provider to enumerate users.

--findable-groups <string>

Specifies a group that can be found in this provider if *--restrict-findable* is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved.

--findable-users <string>

Specifies a user that can be found in this provider if *--restrict-findable* is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be resolved.

--group-domain <string>

Specifies the domain that this provider will use to qualify groups. The default group domain is *NIS_GROUPS*.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with */ifs* and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, *%U*, *%D*, and *%Z* are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--hostname-lookup {yes | no}

Enables or disables host name lookups.

--listable-groups <string>

Specifies a group that can be viewed in this provider if *--restrict-listable* is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed.

--listable-users <string>

Specifies a user that can be viewed in this provider if *--restrict-listable* is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed.

--login-shell <path>

Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

--normalize-groups {yes | no}

Normalizes group name to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user name to lowercase before lookup.

--provider-domain <string>

Specifies the domain that this provider will use to qualify user and group names.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM.

--request-timeout <integer>

Specifies the request timeout interval in seconds. The default value is 20.

--restrict-findable {yes | no}

Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

--retry-time <integer>

Sets the timeout period in seconds after which a request will be retried. The default value is 5.

--unfindable-groups <string>

If `--restrict-findable` is enabled and the findable groups list is empty, specifies a group that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unfindable-users <string>

If `--restrict-findable` is enabled and the findable users list is empty, specifies a user that cannot be resolved by this provider. Repeat this option to specify multiple list items.

--unlistable-groups <string>

If `--restrict-listable` is enabled and the listable groups list is empty, specifies a group that cannot be viewed by this provider. Repeat this option to specify multiple list items.

--unlistable-users <string>

If `--restrict-listable` is enabled and the listable users list is empty, specifies a user that cannot be viewed by this provider. Repeat this option to specify multiple list items.

--user-domain <string>

Specifies the domain that this provider will use to qualify users. The default user domain is `NIS_USERS`.

--ypmatch-using-tcp {yes | no}

Uses TCP for YP Match operations.

--groupnet <groupnet>

Specifies the groupnet referenced by the NIS provider. The groupnet is a top-level networking container that manages hostname resolution against DNS nameservers and contains subnets and IP address pools. The groupnet specifies which networking properties the NIS provider will use when communicating with external servers.

{--verbose | -v}

Displays the results of running the command.

isi auth nis delete

Deletes an NIS provider.

Syntax

```
isi auth nis delete <provider-name>
  [--force]
  [--verbose]
```

Options

<provider-name>

Specifies the name of the provider to delete.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Returns a success or fail message after running the command.

isi auth nis list

Displays a list of NIS providers and indicates whether a provider is functioning correctly.

Syntax

```
isi auth nis list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth nis modify

Modifies an NIS provider.

Syntax

```
isi auth nis modify <provider-name>
  [--name <string>]
  [--nis-domain <string>]
  [--servers <string>]
  [--add-servers <string>]
  [--remove-servers <string>]
  [--authentication {yes | no}]
  [--balance-servers {yes | no}]
  [--check-online-interval <duration>]
  [--create-home-directory {yes | no}]
  [--enabled {yes | no}]
  [--enumerate-groups {yes | no}]
  [--enumerate-users {yes | no}]
  [--findable-groups <string>]
  [--clear-findable-groups]
  [--add-findable-groups <string>]
  [--remove-findable-groups <string>]
  [--findable-users <string>]
```

```

[--clear-findable-users]
[--add-findable-users <string>]
[--remove-findable-users <string>]
[--group-domain <string>]
[--home-directory-template <string>]
[--hostname-lookup {yes | no}]
[--listable-groups <string>]
[--clear-listable-groups]
[--add-listable-groups <string>]
[--remove-listable-groups <string>]
[--listable-users <string>]
[--clear-listable-users]
[--add-listable-users <string>]
[--remove-listable-users <string>]
[--login-shell <string>]
[--normalize-groups {yes | no}]
[--normalize-users {yes | no}]
[--provider-domain <string>]
[--ntlm-support {all | v2only | none}]
[--request-timeout <integer>]
[--restrict-findable {yes | no}]
[--restrict-listable {yes | no}]
[--retry-time <integer>]
[--unfindable-groups <string>]
[--clear-unfindable-groups]
[--add-unfindable-groups <string>]
[--remove-unfindable-groups <string>]
[--unfindable-users <string>]
[--clear-unfindable-users]
[--add-unfindable-users <string>]
[--remove-unfindable-users <string>]
[--unlistable-groups <string>]
[--clear-unlistable-groups]
[--add-unlistable-groups <string>]
[--remove-unlistable-groups <string>]
[--unlistable-users <string>]
[--clear-unlistable-users]
[--add-unlistable-users <string>]
[--remove-unlistable-users <string>]
[--user-domain <string>]
[--ypmatch-using-tcp {yes | no}]
[--verbose]

```

Options

<provider-name>

Specifies the name of the NIS provider to modify.

--name <string>

Specifies an new name for the authentication provider.

--nis-domain <string>

Specifies the NIS domain name.

--servers <string>

Specifies a list of NIS server to be used by this provider. Repeat this option to specify multiple list items. Specify the NIS server as an IPv4 address or hostname. This option overwrites the entries in the NIS servers list; to add or remove servers without affecting current entries, use `--add-servers` or `--remove-servers`.

--add-servers <string>

Adds an entry to the list of NIS servers. Repeat this option to specify multiple items.

--remove-servers <string>

Removes an entry from the list of NIS servers. Repeat this option to specify multiple items.

--authentication {yes | no}

Enables or disables the use of this provider for authentication as well as identity. The default value is `yes`.

--balance-servers {yes | no}

Makes this provider connect to a random server on each request.

--check-online-interval <duration>

Specifies the time between provider online checks, in the format <integer>[*{Y | M | W | D | H | m | s}*].

--create-home-directory {yes | no}

Specifies whether to create a home directory the first time a user logs in, if a home directory does not already exist for the user.

--enabled {yes | no}

Enables or disables this provider.

--enumerate-groups {yes | no}

Specifies whether to allow this provider to enumerate groups.

--enumerate-users {yes | no}

Specifies whether to allow this provider to enumerate users.

--findable-groups <string>

Specifies a group that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be resolved. This option overwrites the entries in the findable groups list; to add or remove groups without affecting current entries, use `--add-findable-groups` or `--remove-findable-groups`.

--clear-findable-groups

Removes all entries from the list of findable groups.

--add-findable-groups <string>

Adds an entry to the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-groups <string>

Removes an entry from the list of findable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--findable-users <string>

Specifies a user that can be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be resolved. This option overwrites the entries in the findable users list; to add or remove users without affecting current entries, use `--add-findable-users` or `--remove-findable-users`.

--clear-findable-users

Removes all entries from the list of findable users.

--add-findable-users <string>

Adds an entry to the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-findable-users <string>

Removes an entry from the list of findable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--group-domain <string>

Specifies the domain that this provider will use to qualify groups. The default group domain is `NIS_GROUPS`.

--home-directory-template <path>

Specifies the path to use as a template for naming home directories. The path must begin with `/ifs` and can include special character sequences that are dynamically replaced with strings at home directory creation time that represent specific variables. For example, `%U`, `%D`, and `%Z` are replaced with the user name, provider domain name, and zone name, respectively. For more information, see the Home directories section.

--hostname-lookup {yes | no}

Enables or disables host name lookups.

--listable-groups <string>

Specifies a group that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, groups that are not included in this list cannot be viewed. This option overwrites the entries in the listable groups list; to add or remove groups without affecting current entries, use `--add-listable-groups` or `--remove-listable-groups`.

--clear-listable-groups

Removes all entries from the list of viewable groups.

--add-listable-groups <string>

Adds an entry to the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-groups <string>

Removes an entry from the list of viewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--listable-users <string>

Specifies a user that can be viewed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. If populated, users that are not included in this list cannot be viewed. This option overwrites the entries in the listable users list; to add or remove users without affecting current entries, use `--add-listable-users` or `--remove-listable-users`.

--clear-listable-users

Removes all entries from the list of viewable users.

--add-listable-users <string>

Adds an entry to the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-listable-users <string>

Removes an entry from the list of viewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--login-shell <path>

Specifies the path to the user's login shell. This setting applies only to users who access the file system through SSH.

--normalize-groups {yes | no}

Normalizes group names to lowercase before lookup.

--normalize-users {yes | no}

Normalizes user names to lowercase before lookup.

--provider-domain <string>

Specifies the domain that this provider will use to qualify user and group names.

--ntlm-support {all | v2only | none}

For users with NTLM-compatible credentials, specifies which NTLM versions to support. Valid values are `all`, `v2only`, and `none`. NTLMv2 provides additional security over NTLM.

--request-timeout <integer>

Specifies the request timeout interval in seconds. The default value is 20.

--restrict-findable {yes | no}

Specifies whether to check this provider for filtered lists of findable and unfindable users and groups.

--restrict-listable {yes | no}

Specifies whether to check this provider for filtered lists of viewable and unviewable users and groups.

--retry-time <integer>

Sets the timeout period in seconds after which a request will be retried. The default value is 5.

--unfindable-groups <string>

Specifies a group that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable groups list; to add or remove groups without affecting current entries, use `--add-unfindable-groups` or `--remove-unfindable-groups`.

--clear-unfindable-groups

Removes all entries from the list of unfindable groups.

--add-unfindable-groups <string>

Adds an entry to the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-groups <string>

Removes an entry from the list of unfindable groups that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unfindable-users <string>

Specifies a user that cannot be found in this provider if `--restrict-findable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unfindable users list; to add or remove users without affecting current entries, use `--add-unfindable-users` or `--remove-unfindable-users`.

--clear-unfindable-users

Removes all entries from the list of unfindable groups.

--add-unfindable-users <string>

Adds an entry to the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--remove-unfindable-users <string>

Removes an entry from the list of unfindable users that is checked if `--restrict-findable` is enabled. Repeat this option to specify multiple list items.

--unlistable-groups <string>

Specifies a group that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable groups list; to add or remove groups without affecting current entries, use `--add-unlistable-groups` or `--remove-unlistable-groups`.

--clear-unlistable-groups

Removes all entries from the list of unlistable groups.

--add-unlistable-groups <string>

Adds an entry to the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-groups <string>

Removes an entry from the list of unviewable groups that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--unlistable-users <string>

Specifies a user that cannot be listed in this provider if `--restrict-listable` is enabled. Repeat this option to specify multiple list items. This option overwrites the entries in the unlistable users list; to add or remove users without affecting current entries, use `--add-unlistable-users` or `--remove-unlistable-users`.

--clear-unlistable-users

Removes all entries from the list of unviewable users.

--add-unlistable-users <string>

Adds an entry to the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--remove-unlistable-users <string>

Removes an entry from the list of unviewable users that is checked if `--restrict-listable` is enabled. Repeat this option to specify multiple list items.

--user-domain <string>

Specifies the domain that this provider will use to qualify users. The default user domain is `NIS_USERS`.

--ypmatch-using-tcp {yes | no}

Uses TCP for YP Match operations.

{--verbose | -v}

Displays the results of running the command.

isi auth nis view

Displays the properties of an NIS provider.

Syntax

```
isi auth nis view <provider-name>
```

Options

<provider-name>

Specifies the name of the provider to view.

isi auth privileges

Displays a list of system privileges.

Syntax

```
isi auth privileges
  [--zone <string>
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the access zone for which to display system privileges.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}


Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

 **NOTE: When using the --verbose option, the output Read Write: No means that the privileges are read-only.**

isi auth refresh

Refreshes authentication system configuration settings.

Syntax

```
isi auth refresh
```

Options

There are no options for this command.

isi auth roles create

Creates a custom role.

This command creates an empty role. To assign privileges and add members to the role, run the `isi auth roles modify` command.

Syntax

```
isi auth roles create <name>
  [--description <string>]
  [--zone <string>]
  [--verbose]
```

Options

<name>

Specifies the name of the role.

--description <string>

Specifies a description of the role.

--zone <string>

Specifies the name of the access zone in which to create the role.

{--verbose | -v}

Displays the results of running the command.

isi auth roles delete

Deletes a role.

Syntax

```
isi auth roles delete <role>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<role>

Specifies the name of the role to delete.

--zone <string>

Specifies the name of the access zone from which to delete the role.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

isi auth roles list

Displays a list of roles.

Syntax

```
isi auth roles list
  [--zone <string>]
  [--limit <integer>]
  [--sort {id | name | description}]
  [{--descending | -d}]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the access zone for which to list roles.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort {id | name | description}

Specifies the field on which to sort data.

{--descending | -d}

Specifies to sort data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth roles members list

Displays a list of the members of a role.

Syntax

```
isi auth roles members list <role>
  [--zone <string>]
  [--limit <integer>]
  [--sort {id | name | type}]
  [{--descending | -d}]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<role>

Specifies a role by name.

--zone <string>

Specifies the access zone for which to display members of a role.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort {id | name | type}

Sorts data by the specified field.

{--descending | -d}

Sorts data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To view the members of the SystemAdmin role, run the following command:

```
isi auth roles members list systemadmin
```

In the following sample output, the SystemAdmin role currently contains one member, a user named admin:

```
Type Name
-----
user admin
-----
Total: 1
```

isi auth roles modify

Modifies a role.

Syntax

```
isi auth roles modify <role>
  [--zone <string>]
  [--name <string>]
  [--description <string>]
  [--add-group <string>]
  [--remove-group <string>]
  [--add-gid <integer>]
  [--remove-gid <integer>]
  [--add-uid <integer>]
  [--remove-uid <integer>]
  [--add-user <string>]
  [--remove-user <string>]
  [--add-sid <string>]
  [--remove-sid <string>]
  [--add-wellknown <string>]
```

```
[--remove-wellknown <string>]
[--add-priv <string>]
[--add-priv-ro <string>]
[--remove-priv <string>]
[--verbose]
```

Options

<role>

Specifies the name of the role to modify.

--zone <string>

Specifies the name of the zone in which to modify the role.

--name <string>

Specifies a new name for the role. Applies to custom roles only.

--description <string>

Specifies a description of the role.

--add-group <string>

Adds a group with the specified name to the role. Repeat this option for each additional item.

--remove-group <string>

Removes a group with the specified name from the role. Repeat this option for each additional item.

--add-gid <integer>

Adds a group with the specified GID to the role. Repeat this option for each additional item.

--remove-gid <integer>

Removes a group with the specified GID from the role. Repeat this option for each additional item.

--add-uid <integer>

Adds a user with the specified UID to the role. Repeat this option for each additional item.

--remove-uid <integer>

Removes a user with the specified UID from the role. Repeat this option for each additional item.

--add-user <string>

Adds a user with the specified name to the role. Repeat this option for each additional item.

--remove-user <string>

Removes a user with the specified name from the role. Repeat this option for each additional item.

--add-sid <string>

Adds a user or group with the specified SID to the role. Repeat this option for each additional item.

--remove-sid <string>

Removes a user or group with the specified SID from the role. Repeat this option for each additional item.

--add-wellknown <string>

Adds a well-known SID with the specified name—for example, Everyone—to the role. Repeat this option for each additional item.

--remove-wellknown <string>

Removes a well-known SID with the specified name from the role. Repeat this option for each additional item.

--add-priv <string>

Adds a read/write privilege to the role. Applies to custom roles only. Repeat this option for each additional item.

--add-priv-ro <string>

Adds a read-only privilege to the role. Applies to custom roles only. Repeat this option for each additional item.

--remove-priv <string>

Removes a privilege from the role. Applies to custom roles only. Repeat this option for each additional item.

{--verbose | -v}

Displays the results of running the command.

isi auth roles privileges list

Displays a list of privileges that are associated with a role.

Syntax

```
isi auth roles privileges list <role>
  [--zone <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<role>

Specifies a role by name.

--zone <string>

Specifies the role's access zone.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To list the privileges that are associated with the built-in SecurityAdmin role, run the following command:

```
isi auth roles privileges list securityadmin
```

The system displays output similar to the following example:

```
ID
-----
ISI_PRIV_LOGIN_CONSOLE
ISI_PRIV_LOGIN_PAPI
ISI_PRIV_LOGIN_SSH
ISI_PRIV_AUTH
ISI_PRIV_ROLE
-----
Total: 5
```

isi auth roles view

Displays the properties of a role.

Syntax

```
isi auth roles view <role>
[--zone <string>
```

Options

<role>

Specifies the name of the role to view.

--zone <string>

Specifies the role's access zone.

isi auth settings acls modify

Modifies access control list (ACL) settings for OneFS.

Syntax

```
isi auth settings acls modify
  [--create-over-smb {allow | disallow}]
  [--chmod {remove | replace | replace_users_and_groups | merge | deny | ignore}]
  [--chmod-inheritable {yes | no}]
  [--chown {owner_group_and_acl | owner_group_only | ignore}]
  [--access {unix | windows}]
  [--rwx {retain | full_control}]
  [--group-owner-inheritance {native | parent | creator}]
  [--chmod-007 {default | remove}]
  [--calcmode-owner {owner_aces | owner_only}]
  [--calcmode-group {group_aces | group_only}]
  [--synthetic-denies {none | remove}]
  [--utimes {only_owner | owner_and_write}]
  [--dos-attr {deny_smb | deny_smb_and_nfs}]
  [--calcmode {approx | 777}]
  [--calcmode-traverse {ignore | require}]
  [--verbose]
```

Options

--create-over-smb {allow | disallow}

Specifies whether to allow or deny creation of ACLs over SMB.



NOTE: Inheritable ACLs on the system take precedence over this setting. If inheritable ACLs are set on a folder, any new files and folders created in that folder will inherit the folder's ACL. Disabling this setting does not remove ACLs currently set on files. If you want to clear an existing ACL, run the `chmod -b <mode> <file>` command to remove the ACL and set the correct permissions.

--chmod {remove | replace | replace_users_and_groups | merge | deny | ignore}

Specifies how permissions are handled when a `chmod` operation is initiated on a file with an ACL, either locally or over NFS. This setting controls any elements that affect UNIX permissions, including File System Explorer. Enabling this policy setting does not change how `chmod` operations affect files that do not have ACLs. The following values are valid:

remove	For <code>chmod</code> operations, removes any existing ACL and instead sets the <code>chmod</code> permissions. Select this option only if you do not need permissions to be set from Windows.
replace	Removes the existing ACL and creates an ACL equivalent to the UNIX permissions. Select this option only if you want to remove Windows permissions but do not want files to have synthetic ACLs.
replace_users_and_groups	Removes the existing ACL and creates an ACL equivalent to the UNIX permissions for all users/groups referenced in old ACL. Select this option only if you want to remove Windows permissions but do not want files to have synthetic ACLs.
merge	Merges permissions that are applied by <code>chmod</code> with existing ACLs. An ACE for each identity (owner, group, and everyone) is either modified or created, but all other ACEs are unmodified. Inheritable ACEs are also left unmodified to enable Windows users to continue to inherit appropriate permissions. UNIX users can set specific permissions for each of those three standard identities, however.
deny	Prevents users from making NFS and local <code>chmod</code> operations. Enable this setting if you do not want to allow permission sets over NFS.
ignore	Ignores the <code>chmod</code> operation if file has an existing ACL, which prevents an NFS client from making changes to the ACL. Select this option if you defined an inheritable ACL on a directory and want to use that ACL for permissions.

CAUTION: If you attempt to run the `chmod` command on the same permissions that are currently set on a file with an ACL, you may cause the operation to silently fail. The operation appears to be successful, but if you were to examine the permissions on the cluster, you would notice that the `chmod` command had no effect. As an alternative, you can run the `chmod` command away from the current permissions and then perform a second `chmod` command to revert to the original permissions. For example, if your file shows 755 UNIX permissions and you want to confirm this number, you could run `chmod 700 file; chmod 755 file`.

`--chmod-inheritable {yes | no}`

On Windows systems, the ACEs for directories can define detailed inheritance rules. On a UNIX system, the mode bits are not inherited. Making ACLs that are created on directories by the `chmod` command inheritable is more secure for tightly controlled environments but may deny access to some Windows users who would otherwise expect access.

`--chown {owner_group_and_acl | owner_group_only | ignore}`

Changes the user or group that has ownership of a file or folder. The following values are valid:

owner_group_and_acl	Modifies only the owner or group, which enables the <code>chown</code> or <code>chgrp</code> operation to perform as it does in UNIX. Enabling this setting modifies any ACEs in the ACL associated with the old and new owner or group.
owner_group_only	Modifies the owner or group and ACL permissions, which enables the NFS <code>chown</code> or <code>chgrp</code> operation to function as it does in Windows. When a file owner is changed over Windows, no permissions in the ACL are changed.
ignore	Ignores the <code>chown</code> and <code>chgrp</code> operations if file has an existing ACL, which prevents an NFS client from making changes to the owner or group.

NOTE: Over NFS, the `chown` or `chgrp` operation changes the permissions and user or group that has ownership. For example, a file owned by user Joe with `rwX-----` (700) permissions indicates `rwX` permissions for the owner, but no permissions for anyone else. If you run the `chown` command to change ownership of the file to user Bob, the owner permissions are still `rwX` but they now represent the permissions for Bob, rather than for Joe, who lost all of his permissions. This setting does not affect UNIX `chown` or `chgrp` operations performed on files with UNIX permissions, and it does not affect Windows `chown` or `chgrp` operations, which do not change any permissions.

`--access {unix | windows}`

In UNIX environments, only the file owner or superuser has the right to run a `chmod` or `chown` operation on a file. In Windows environments, you can implement this policy setting to give users the right to perform `chmod`

operations that change permissions, or the right to perform `chown` operations that take ownership, but do not give ownership away. The following values are valid:

- | | |
|----------------|---|
| unix | Allows only the file owner to change the mode or owner of the file, which enable <code>chmod</code> and <code>chown</code> access checks to operate with UNIX-like behavior. |
| windows | Allow the file owner and users with <code>WRITE_DAC</code> and <code>WRITE_OWNER</code> permissions to change the mode or owner of the file, which enables <code>chmod</code> and <code>chown</code> access checks to operate with Windows-like behavior. |

--rwx {retain | full_control}

Specifies how to handle `rwx` permissions mapped to windows rights. In UNIX environments, `rwx` permissions indicate that a user or group has read, write, and execute permissions and that a user or group has the maximum level of permissions.

When you assign UNIX permissions to a file, no ACLs are stored for that file. Because a Windows system processes only ACLs, the Isilon cluster must translate the UNIX permissions into an ACL when you view a file's permissions on a Windows system. This type of ACL is called a synthetic ACL. Synthetic ACLs are not stored anywhere; instead, they are dynamically generated and discarded as needed. If a file has UNIX permissions, you may notice synthetic ACLs when you run the `ls file` command to view a file's ACLs.

When you generate a synthetic ACL, the Isilon cluster maps UNIX permissions to Windows rights. Windows supports a more granular permissions model than UNIX does, and it specifies rights that cannot easily be mapped from UNIX permissions. The following values are valid:

- | | |
|---------------------|--|
| retain | Retains <code>rwx</code> permissions and generates an ACE that provides only read, write, and execute permissions. |
| full_control | Treats <code>rwx</code> permissions as full control and generates an ACE that provides the maximum Windows permissions for a user or a group by adding the change permissions right, the take ownership right, and the delete right. |

--group-owner-inheritance {native | parent | creator}

Specifies how to handle inheritance of group ownership and permissions. If you enable a setting that causes the group owner to be inherited from the creator's primary group, you can override it on a per-folder basis by running the `chmod` command to set the `set-gid` bit. This inheritance applies only when the file is created. The following values are valid:

- | | |
|----------------|--|
| native | Specifies that if an ACL exists on a file, the group owner will be inherited from the file creator's primary group. If there is no ACL, the group owner is inherited from the parent folder. |
| parent | Specifies that the group owner be inherited from the file's parent folder. |
| creator | Specifies that the group owner be inherited from the file creator's primary group. |

--chmod-007 {default | remove}

Specifies whether to remove ACLs when running the `chmod (007)` command. The following values are valid:

- | | |
|----------------|--|
| default | Sets 007 UNIX permissions without removing an existing ACL. |
| remove | Removes ACLs from files over UNIX file sharing (NFS) and locally on the cluster through the <code>chmod (007)</code> command. If you enable this setting, be sure to run the <code>chmod</code> command on the file immediately after using <code>chmod (007)</code> to clear an ACL. In most cases, you do not want to leave 007 permissions on the file. |

--calcmode-owner {owner_aces | owner_only}

Specifies how to approximate owner mode bits. The following values are valid:

- | | |
|-------------------|---|
| owner_aces | Approximates owner mode bits using all possible group ACEs. This causes the owner permissions to appear more permissive than the actual permissions on the file. |
| owner_only | Approximates owner mode bits using only the ACE with the owner ID. This causes the owner permissions to appear more accurate, in that you see only the permissions for a particular owner and not the more permissive set. This may cause access-denied problems for UNIX clients, however. |

--calcmode-group {group_aces | group_only}

Specifies how to approximate group mode bits. The following values are valid:


group_aces Approximates group mode bits using all possible group ACEs. This causes the group permissions to appear more permissive than the actual permissions on the file.

group_only Approximates group mode bits using only the ACE with the owner ID. This causes the group permissions to appear more accurate, in that you see only the permissions for a particular group and not the more permissive set. This may cause access-denied problems for UNIX clients, however.

--synthetic-denies {none | remove}

Specifies how to handle synthetic ACLs. The Windows ACL user interface cannot display an ACL if any deny ACEs are out of canonical ACL order. To correctly represent UNIX permissions, deny ACEs may be required to be out of canonical ACL order. The following values are valid:

none Does not modify synthetic ACLs and mode bit approximations, which prevents modifications to synthetic ACL generation and allows “deny” ACEs to be generated when necessary.

 **CAUTION: This option can lead to permissions being reordered, permanently denying access if a Windows user or an application performs an ACL get, an ACL modification, and an ACL set to and from Windows.**

remove Removes deny ACEs when generating synthetic ACLs. This setting can cause ACLs to be more permissive than the equivalent mode bits.

--utimes {only_owner | owner_and_write}

Specifies who can change utimes, which are the access and modification times of a file.

only_owner Allows only owners to change utimes to client-specific times, which complies with the POSIX standard.

owner_and_write Allows owners as well as users with write access to modify utimes to client-specific times, which is less restrictive.

--dos-attr {deny_smb | deny_smb_and_nfs}

Specifies how to handle the read-only DOS attribute for NFS and SMB. The following values are valid:

deny_smb Denies permission to modify files with DOS read-only attribute over SMB only.

deny_smb_nfs Denies permission to modify files with DOS read-only attribute through both NFS and SMB.

--calcmode {approx | 777}

Specifies how to display mode bits. The following values are valid:

approx Specifies to use ACL to approximate mode bits. Displays the approximation of the NFS mode bits that are based on ACL permissions.

777 Specifies to always display 777 if an ACL exists. If the approximated NFS permissions are less permissive than those in the ACL, you may want to use this setting so the NFS client does not stop at the access check before performing its operation. Use this setting when a third-party application may be blocked if the ACL does not provide the proper access.

--calcmode-traverse {ignore | require}

Specifies whether or not traverse rights are required in order to traverse directories with existing ACLs. The following values are valid:

ignore Specifies that traverse rights are not required.

require Specifies that traverse rights are required.

{--verbose | -v}

Displays more detailed information.

isi auth settings acls view

Displays access control list (ACL) settings for OneFS.

Syntax

```
isi auth settings acls view
```

Options

There are no options for this command.

isi auth settings global modify

Modifies the global authentication settings.

Syntax

```
isi auth settings global modify
  [--send-ntlmv2 {yes | no}]
  [--revert-send-ntlmv2]
  [--space-replacement <character>]
  [--revert-space-replacement]
  [--workgroup <string>]
  [--revert-workgroup]
  [--provider-hostname-lookup {dns-first | nis-first | disabled}]
  [--user-object-cache-size <size>]
  [--revert-user-object-cache-size]
  [--on-disk-identity {native | unix | sid}]
  [--revert-on-disk-identity]
  [--rpc-block-time <duration>]
  [--revert-rpc-block-time]
  [--rpc-max-requests <integer>]
  [--revert-rpc-max-requests]
  [--unknown-gid <integer>]
  [--revert-unknown-gid]
  [--unknown-uid <integer>]
  [--revert-unknown-uid]
  [--verbose]
```

Options

--send-ntlmv2 {yes | no}

Specifies whether to send only NTLMv2 responses to an SMB client. The default value is `no`. Valid values are `yes`, `no`. The default value is `no`.

--revert-send-ntlmv2

Reverts the `--send-ntlmv2` setting to the system default value.

--space-replacement <character>

For clients that have difficulty parsing spaces in user and group names, specifies a substitute character. Be careful to choose a character that is not in use.

--revert-space-replacement

Reverts the `--space-replacement` setting to the system default value.

--workgroup <string>

Specifies the NetBIOS workgroup. The default value is `WORKGROUP`.

--revert-workgroup

Reverts the `--workgroup` setting to the system default value.

`--provider-hostname-lookup {dns-first | nis-first | disabled}`

Allows hostname lookup through authentication providers. Applies to NIS only. The default value is `disabled`.

`--user-object-cache-size <size>`

Specifies the maximum size (in bytes) of the security object cache in the authentication service.

`--revert-user-object-cache-size`


Reverts the `--user-object-cache-size` setting to the system default value.

`--on-disk-identity <string>`

Controls the preferred identity to store on disk. If OneFS is unable to convert an identity to the preferred format, it is stored as is. This setting does not affect identities that are already stored on disk.

The accepted values are listed below.

native	Allows OneFS to determine the identity to store on disk. This is the recommended setting.
unix	Always stores incoming UNIX identifiers (UIDs and GIDs) on disk.
sid	Stores incoming Windows security identifiers (SIDs) on disk unless the SID was generated from a UNIX identifier. If the SID was generated from a UNIX identifier, OneFS converts it back to the UNIX identifier and stores it on disk.

 **NOTE: To prevent permission errors after changing the on-disk identity, run the Repair Permissions job with the `convert` mode specified.**

`--revert-on-disk-identity`

Sets the `--on-disk-identity` setting to the system default value.

`--rpc-block-time <integer>`

Specifies the length of time, in milliseconds, before an ID mapper request becomes asynchronous.

`--revert-rpc-block-time`

Sets the `--rpc-block-time` setting to the system default value.

`--rpc-max-requests <integer>`

Specifies the maximum number of simultaneous ID mapper requests allowed. The default value is 64.

`--revert-rpc-max-requests`

Sets the `--rpc-max-requests` setting to the system default value.

`--unknown-gid <integer>`

Specifies the GID to use for the unknown (anonymous) group.

`--revert-unknown-gid`

Sets the `--unknown-gid` setting to the system default value.

`--unknown-uid <integer>`

Specifies the UID to use for the unknown (anonymous) user.

`--revert-unknown-uid`

Sets the `--unknown-uid` setting to the system default value.

`{--verbose | -v}`

Displays more detailed information.

isi auth settings global view

Displays global authentication settings.

Syntax

```
isi auth settings global view
```

Options

There are no options for this command.

Examples

To view the current authentication settings on the cluster, run the following command:

```
isi auth settings global view
```

The system displays output similar to the following example:

```
        Send NTLMv2: No
        Space Replacement:
          Workgroup: WORKGROUP
Provider Hostname Lookup: disabled
        Alloc Retries: 5
Cache Cred Lifetime: 15m
Cache ID Lifetime: 15m
  On Disk Identity: native
        RPC Block Time: 5s
        RPC Max Requests: 16
          RPC Timeout: 30s
System GID Threshold: 80
System UID Threshold: 80
  GID Range Enabled: Yes
    GID Range Min: 1000000
    GID Range Max: 2000000
  UID Range Enabled: Yes
    UID Range Min: 1000000
    UID Range Max: 2000000
  Min Mapped Rid: 2147483648
    Group UID: 4294967292
    Null GID: 4294967293
    Null UID: 4294967293
    Unknown GID: 4294967294
    Unknown UID: 4294967294
```

isi auth settings krb5 modify

Modifies the global settings of an MIT Kerberos authentication provider.

Syntax

```
isi auth settings krb5 modify
  [--allow-weak-crypto <boolean>]
  [--revert-allow-weak-crypto
    [--always-send-preauth <boolean> | --revert-always-send-preauth]
    [--default-realm <string>]
    [--dns-lookup-kdc <boolean> | --revert-dns-lookup-kdc]
    [--dns-lookup-realm <boolean> | --revert-dns-lookup-realm]
  [--verbose]
```

Options

--allow-weak-crypto <boolean>

Enables DES encryption support.

--revert-allow-weak-crypto

Sets the value to the system default for --allow-weak-crypto.

--always-send-preauth <boolean>

Specifies whether to send preauth.

--revert-always-send-preauth
Sets the value of `--always-send-preauth` to the system default.

--default-realm <string>
Specifies the default Kerberos realm name.

--dns-lookup-kdc <boolean>
Allows DNS to find Key Distribution Centers (KDCs).

--revert-dns-lookup-kdc
Sets the value of `--dns-lookup-kdc` to the system default.

--dns-lookup-realm <boolean>
Allows DNS to find the Kerberos realm names.

--revert-dns-lookup-realm
Sets the value of `--dns-lookup-realm` to the system default.

{--verbose | -v}
Displays more detailed information.

isi auth settings krb5 view

Displays MIT Kerberos provider authentication settings.

Syntax

```
isi auth settings krb5 view
```

isi auth settings mapping modify

Modifies identity mapping settings.

Syntax

```
isi auth settings mapping modify
  [--gid-range-enabled {yes | no}]
  [--gid-range-min <integer>]
  [--gid-range-max <integer>]
  [--gid-range-next <integer>]
  [--uid-range-enabled {yes | no}]
  [--uid-range-min <integer>]
  [--uid-range-max <integer>]
  [--uid-range-next <integer>]
  [--zone <string>]
  [--verbose]
```

Options

If no option is specified, the kernel mapping database is displayed.

--gid-range-enabled {yes | no}
Enables automatic allocation of GIDs by the ID mapping service. This setting is enabled by default.

--gid-range-min <integer>
Specifies the lower value in the range of GIDs that are available for allocation. The default value is 1000000.

--gid-range-max <integer>
Specifies the upper value of the range of GIDs that are available for allocation. The default value is 2000000.

--gid-range-next <integer>

Specifies the next GID to allocate.

--uid-range-enabled {yes | no}

Enables automatic allocation of UIDs by the ID mapping service. This setting is enabled by default.

--uid-range-min <integer>

Specifies the lower value in the range of UIDs that are available for allocation. The default value is 1000000.

--uid-range-max <integer>

Specifies the upper value in the range of UIDs that are available for allocation. The default value is 2000000.

--uid-range-next <integer>

Specifies the next UID to allocate.

--zone <string>

Specifies the access zone in which to modify ID mapping settings. If no access zone is specified, settings in the default System zone will be modified.

{--verbose | -v

Displays more detailed information.

isi auth settings mapping view

Displays identity mapping settings in an access zone.

Syntax

```
isi auth settings mapping view  
[--zone <string>]
```

Options

--zone <string>

Displays mapping settings from the specified access zone. If no access zone is specified, displays mappings from the default System zone.

isi auth status

Displays provider status, including available authentication providers and which providers are functioning correctly.

Syntax

```
isi auth status  
  [--zone <string>]  
  [--groupnet <string>]  
  [--limit <integer>]  
  [--format {table | json | csv | list}]  
  [--no-header]  
  [--no-footer]  
  [--verbose]
```

Options

--zone <string>

Specifies an access zone by name.

--groupnet <string>

Specifies a groupnet by name.

--limit [-l | <integer>]

Specifies the number of providers to display.

--format {table | json | csv | list}

Displays providers in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth users create

Creates a user account.

Syntax

```
isi auth users create <name>
  [--enabled {yes | no}]
  [--expiry <timestamp>]
  [--email <string>]
  [--gecos <string>]
  [--home-directory <path>]
  [--password <string>]
  [--password-expires {yes | no}]
  [{--primary-group <name> | --primary-group-gid <integer>
  | --primary-group-sid <string>}]
  [--prompt-password-change {yes | no}]
  [--shell <path>]
  [--uid <integer>]
  [--zone <string>]
  [--provider <string>]
  [--set-password]
  [--verbose]
  [--force]
```

Options

<name>

Specifies the user name.

--enabled {yes | no}

Enables or disables the user.

{--expiry | -x} <timestamp>

Specifies the time at which the user account will expire, using the date format <YYYY>-<MM>-<DD> or the date/time format <YYYY>-<MM>-<DD>T<hh>:<mm>[:<ss>].

--email <string>

Specifies the email address of the user.

--gecos <string>

Specifies the values for the following Gecos field entries in the user's password file:

Full Name:
Office Location:
Office Phone:
Home Phone:
Other information:

Values must be entered as a comma-separated list, and values that contain spaces must be enclosed in quotation marks. For example, the `--gecos="Jane Doe",Seattle,555-5555,, "Temporary worker"` option with these values results in the following entries:

```
Full Name: Jane Doe
Office Location: Seattle
Office Phone: 555-5555
Home Phone:
Other information: Temporary worker
```

--home-directory <path>

Specifies the path to the user's home directory.

--password <string>

Sets the user's password to the specified value. This option cannot be used with the `--set-password` option.

--password-expires {yes | no}

Specifies whether to allow the password to expire.

--primary-group <name>

Specifies the user's primary group by name.

--primary-group-gid <integer>

Specifies the user's primary group by GID.

--primary-group-sid <string>

Specifies the user's primary group by SID.

--prompt-password-change {yes | no}

Prompts the user to change the password during the next login.

--shell <path>

Specifies the path to the UNIX login shell.

--uid <integer>

Overrides automatic allocation of the UNIX user identifier (UID) with the specified value. Setting this option is not recommended.

--zone <string>

Specifies the access zone in which to create the user.

--provider <string>

Specifies a local authentication provider in the specified access zone.

--set-password

Sets the password interactively. This option cannot be used with the `--password` option.

{--verbose | -v}

Displays the results of running the command.

{--force | -f}

Suppresses command-line prompts and messages.

isi auth users delete

Deletes a local user from the system.

Syntax

```
isi auth users delete {<user> | --uid <integer> | --sid <string>}
  [--zone <string>]
  [--provider <string>]
  [--force]
  [--verbose]
```

Options

This command requires `<user>`, `--uid <integer>`, or `--sid <string>`.

`<user>`

Specifies the user by name.

`--uid <integer>`

Specifies the user by UID.

`--sid <string>`

Specifies the user by SID.

`--zone <string>`

Specifies the name of the access zone that contains the user.

`--provider <string>`

Specifies the name of the authentication provider that contains the user.

`{--force | -f}`

Suppresses command-line prompts and messages.

`{--verbose | -v}`

Displays the results of running the command.

isi auth users flush

Flushes cached user information.

Syntax

```
isi auth users flush
```

Options

There are no options for this command.

Examples

To flush all cached user information, run the following command:

```
isi auth user flush
```

isi auth users list

Displays a list of users. If no options are specified, all users in the System access zone are displayed.

 **NOTE:** The `--domain` option must be specified to list Active Directory users.

Syntax

```
isi auth users list
  [--domain <string>]
  [--zone <string>]
  [--provider <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

--domain <string>

Displays only the users in the specified provider domain.

--zone <string>

Specifies the access zone whose users you want to list. The default access zone is `System`.

--provider <string>

Displays only the users in the specified authentication provider. The syntax for specifying providers is *<provider-type>:<provider-name>*, being certain to use the colon separator; for example, `isi auth users list --provider="lsa-ldap-provider:Unix LDAP"`.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi auth users modify

Modifies a local user.

Syntax

```
isi auth users modify {<user> | --uid <integer> | --sid <string>}
  [--enabled {yes | no}]
  [--expiry <timestamp>]
  [--unlock]
  [--email <string>]
  [--gecos <string>]
  [--home-directory <path>]
  [--password <string>]
  [--password-expires {yes | no}]
  [{--primary-group <string> | --primary-group-gid <integer>
  | --primary-group-sid <string>}]
  [--prompt-password-change {yes | no}]
  [--shell <path>]
  [--new-uid <integer>]
  [--zone <string>]
  [--add-group <name>]
  [--add-gid <id>]
  [--remove-group <name>]
  [--remove-gid <id>]
  [--provider <string>]
  [--set-password]
  [--verbose]
  [--force]
```

Options

This command requires `<user>`, `--uid <integer>`, or `--sid <string>`.

<user>

Specifies the user by name.

--uid <integer>

Specifies the user by UID.

--sid <string>

Specifies the user by SID.

--enabled {yes | no}

Enables or disables the user.

{--expiry | -x} <timestamp>

Specifies the time at which the user account will expire, using the date format `<YYYY>-<MM>-<DD>` or the date/time format `<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]`.

--unlock

Unlocks the user account if locked.

--email <string>

Specifies the email address of the user.

--gecos <string>

Specifies the values for the following Gecos field entries in the user's password file:

Full Name:
Office Location:
Office Phone:
Home Phone:
Other information:

Values must be entered as a comma-separated list, and values that contain spaces must be enclosed in quotation marks. For example, the `--gecos="Jane Doe",Seattle,555-5555,, "Temporary worker"` option with these values results in the following entries:

Full Name: Jane Doe
Office Location: Seattle
Office Phone: 555-5555
Home Phone:
Other information: Temporary worker

--home-directory <path>

Specifies the path to the user's home directory.

--password <string>

Sets the user's password to the specified value. This option cannot be used with the `--set-password` option.

--password-expires {yes | no}

Specifies whether to allow the password to expire.

--primary-group <name>

Specifies the user's primary group by name.

--primary-group-gid <integer>

Specifies the user's primary group by GID.

--primary-group-sid <string>

Specifies the user's primary group by SID.

--prompt-password-change {yes | no}

Prompts the user to change the password during the next login.

--shell <path>

Specifies the path to the UNIX login shell.

--new-uid <integer>

Specifies a new UID for the user. Setting this option is not recommended.

--zone <string>

Specifies the name of the access zone that contains the user.

--add-group <name>

Specifies the name of a group to add the user to. Repeat this option to specify multiple list items.

--add-gid <integer>

Specifies the GID of a group to add the user to. Repeat this option to specify multiple list items.

--remove-group <name>

Specifies the name of a group to remove the user from. Repeat this option to specify multiple list items.

--remove-gid <integer>

Specifies the GID of a group to remove the user from. Repeat this option to specify multiple list items.

--provider <string>

Specifies an authentication provider of the format <type>:<instance>. Valid provider types are `ads`, `ldap`, `nis`, `file`, and `local`. For example, an LDAP provider named `auth1` can be specified as `ldap:auth1`.

--set-password

Sets the password interactively. This option cannot be used with the `--password` option.

{--verbose | -v}

Displays the results of running the command.

{--force | -f}

Suppresses command-line prompts and messages.

isi auth users view

Displays the properties of a user, including historical security identifier (SID) history.

Syntax

```
isi auth users view {<user> | --uid <integer> | --sid <string>}
  [--cached]
  [--show-groups]
  [--resolve-names]
  [--show-ssh-keys]
  [--zone <string>]
  [--provider <string>]
```

Options

This command requires <user>, `--uid <integer>`, or `--sid <string>`.

<user>

Specifies the user by name.

--uid <integer>

Specifies the user by UID.

--sid <string>

Specifies the user by SID.

--cached

Returns only cached information.

--show-groups

Displays groups that include the user as a member.

--resolve-names

Resolves the names of all related groups and related identities.

--show-ssh-keys

Displays SSH Public Key information.

--zone <string>

Specifies the name of the access zone that contains the user.

--provider <string>

Specifies the name of the authentication provider that contains the user in the format <type>:<instance>. Valid values for type are `ads`, `ldap`, `nis`, `file`, and `local`. For example an LDAP provider named `auth1` can be specified as `ldap:auth1`, or an Active Directory provider can be specified as `ads:YORK.east.com`.

isi batterystatus list

Displays a list of batteries in the cluster by node, along with the status of each battery.

Syntax

```
isi batterystatus list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ **--no-header** | **-a**}

Displays table and CSV output without headers.

{ **--no-footer** | **-z**}

Displays table output without footers.

{ **--verbose** | **-v**}

Displays more detailed information.

isi batterystatus view

Displays the status of a node's batteries.

Syntax

```
isi batterystatus view
  [--node-lnn <integer>]
```

Options

--node-lnn <integer>

Specifies the node LNN to view. If omitted, battery status for the local node is displayed.

isi certificate authority delete

Delete a TLS certificate authority.

Syntax

```
isi certificate authority delete <id>
  [--force]
  [--verbose]
```

Options

<id>

The system TLS certificate identifier or name.

{--force | -f}

Does not prompt you to confirm that you want to delete the policy.

{--verbose | -v}

Displays a message confirming that the authority was deleted.

isi certificate authority import

Import a TLS certificate authority.

isi certificate authority list

View a list of TLS certificate authorities.

Syntax

```
isi certificate authority list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

The number of certificate authorities to display. You can specify an integer value.

--format {table | json | csv | list}

Display the list of certificate authorities in table, JSON, CSV, or list format.

{--no-header | -a}

Do not display headers in CSV or table output format.

{--no-footer | -z}

Do not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi certificate authority modify

Modify a TLS certificate authority.

Syntax

```
isi certificate authority modify <id>
  [--name <string>]
  [--description <string>]
  [--verbose]
```

Options

<id>

The certificate identifier or name.

--name <string>

An administrator-configured certificate identifier.

--description <string>

A description field provided for administrative convenience, in which you can enter a comment about the certificate.

{--verbose | -v}

Displays more detailed information.

isi certificate authority view

View a TLS certificate authority.

Syntax

```
isi certificate authority view <id>
  [--format {list | json}]
```

Options

<id>

The certificate identifier or name.

--format {list | json}

Display the list of certificate authorities in list or JSON format.

OneFS displays the specified TLS certificate authority.

isi certificate server delete

Delete a Transport Layer Security (TLS) server certificate.

Syntax

```
isi certificate server delete <id>
  [--force]
  [--verbose]
```

Options

- <id>**
The certificate identifier.
- {--force | -f}**
Skips the confirmation prompt for this command.
- {--verbose | -v}**
Displays more detailed information.

isi certificate server import

Import a Transport Layer Security (TLS) server certificate and key.

Syntax

```
isi certificate server import <certificate-path> <certificate-key-path>
  [--name <string>]
  [--description <string>]
  [--certificate-key-password <string>]
  [--verbose]
```

Option

- name <string>**
The administrator-configured certificate identifier.
- <certificate-path>**
The local path to the TLS certificate file, in PEM format. The certificate file is copied into the system certificate store and can be removed after import. This must be an absolute path within the OneFS file system.
- <certificate-key-path>**
The local path to the TLS certificate key file, in PEM format. The certificate key file is copied into the system certificate store, and should be removed after import for security reasons.
- description <string>**
A description field provided for administrative convenience, in which you can enter a comment about the certificate.
- certificate-key-password <string>**
The password for the certificate key if a private key is password-encrypted.
- {--verbose | -v}**
Displays more detailed information.

isi certificate server list

View a list of Transport Layer Security (TLS) server certificates.

Syntax

```
isi certificate server list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l}

The number of certificate servers to display.

--format {table | json | csv | list}

Display the list of certificate servers in table, JSON, CSV, or list format.

{--no-header | -a}

Do not display headers in CSV or table output format.

{--no-footer | -z}

Do not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi certificate server modify

Modify a Transport Layer Security (TLS) server certificate.

Syntax

```
isi certificate server modify <id>
  [--name]
  [--description <string>]
  [--verbose]
```

Options

<id>

The certificate identifier.

--name

Specifies the administrator-configured name of a certificate to use to connect to a TLS-enabled service.

--description <string>

A description field provided for administrative convenience, in which you can enter a comment about the certificate.

{--verbose | -v}

Displays more detailed information.

isi certificate server view

View a Transport Layer Security (TLS) server certificate.

Syntax

```
isi certificate server view <id>
  [--format {list | json}]
```

Options

<id>

The certificate identifier.

--format {list | json}

Display the list of certificate servers in list or JSON format.

isi certificate settings modify

Modify system-wide TLS certificate settings.

Syntax

```
isi certificate settings modify
  [--certificate-monitor-enabled <boolean>]
  [--certificate-pre-expiration-threshold <duration>]
  [--default-https-certificate <string>]
  [--verbose]
```

Options

- certificate-monitor-enabled <boolean>** Enable certificate expiration monitoring.
- certificate-pre-expiration-threshold <duration>** The amount of time before a certificate expires to state a reminder of the upcoming expiry.
- default-https-certificate <string>** Specify the default HTTPS X.509 certificate for the cluster.
- {--verbose | -v}** Displays more detailed information.

isi certificate settings view

View system-wide TLS certificate settings.

Syntax

```
isi certificate settings view
```

Options

None

OneFS displays the system-wide TLS certificate settings.

isi cloud access add

Adds cloud write access to the cluster.

Syntax

```
isi cloud access add <guid>
  [--expiration-date] <timestamp>
  [--force]
  [--verbose]
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cloud account.

--expiration-date <timestamp>

The date and time when unreferenced (stale) data will be removed from the cloud. The timestamp format is *MMDDYY:hh:mm*. For example, **022016:12:00** specifies an expiration date and time of February 20, 2016 at 12:00 PM.

{-- force | -f}

Do not ask for confirmation.

{--verbose | -v}

Displays more detailed information.

Examples

The following example adds cloud write access to a cluster by specifying the cluster GUID and an expiration date:

```
isi cloud access add 000556bf1e82059801563f1ad44a8c155acf
--expiration-date 022016:12:00
```

OneFS displays a message indicating the cloud accounts and file pool policies to which the secondary cluster will have access, and requires confirmation. Type **yes**, and press ENTER to complete the process.

isi cloud access list

Displays a list of clusters on your network that have, or are eligible for, write access to cloud data. Available clusters are the primary cluster and any other clusters to which data has been replicated with SyncIQ or restored with NDMP.

Syntax

```
isi cloud access list
  [--limit] <integer>
  [--sort {name | guid | state}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of eligible clusters displayed in the list.

--sort {name | guid | state}

Sorts the list of eligible clusters according to the specified category.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--descending | -d}

Outputs the list of eligible clusters in descending order according to the specified sort option.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi cloud access remove

Removes cloud write access from the specified cluster.

Syntax

```
isi cloud access remove <guid>
  [--force]
  [--verbose]
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cluster from which you want to remove cloud write access.

{--force | -f}

Execute the command without requiring confirmation.

{--verbose | -v }

Displays more detailed information.

Examples

The following example removes cloud write access from a cluster identified by a specified GUID:

```
isi cloud access remove 000556bf1e82059801563f1ad44a8c155acf
```

OneFS displays a message indicating the cloud accounts and file pool policies to which the cluster will no longer have access, and requires confirmation. Type **yes**, and press ENTER to complete the process.

isi cloud access view

Displays the details of a cluster that either has or is eligible for write access to cloud data.

Syntax

```
isi cloud access view <guid>
```

Options

<guid>

The reference number, or globally unique identifier (GUID), of the cluster.

isi cloud accounts create

Creates a cloud storage account that connects CloudPools to your cloud storage provider.

Syntax

```
isi cloud accounts create <name> <type> <uri>
  [--account-username <string>]
  [--key <string>]
  [--enabled {yes | no}]
  [--account-id <string>]
  [--telemetry-bucket <string>]
  [--storage-region <string>]
  [--skip-ssl-validation {yes | no}]
  [--enable-ocsp {yes | no}]
  [--ocsp-responder-url-required {yes | no}]
  [--proxy <string>]
  [--credential-provider-uri <string>]
  [--credential-provider-agency <string>]
  [--credential-provider-certificate <string>]
  [--credential-provider-mission <string>]
  [--credential-provider-proxy <string>]
  [--credential-provider-role <string>]
  [--force]
  [--verbose]
```

Options

<name>

The name of the cloud storage account.

<type>

The type of cloud storage account. Use one of the following values:

Value	Description
isilon	Dell EMC Isilon
ecs	Dell EMC ECS Appliance
virtustream	Virtustream Storage Cloud
azure	Microsoft Azure
s3	Amazon S3
c2s-s3	Amazon Commercial Cloud Services S3

Value	Description
google	Google Cloud Platform (using interoperability access)
alibaba-cloud	Alibaba Cloud

<uri>

The cloud account URI. This URI must match that provided to the cloud vendor.

{--account-username | -u} <string>

The username for the cloud account. This name must be identical to the user name provided to the cloud vendor.

--key <string>

The cloud account access key or password. This information is provided by the cloud vendor.

--enabled {yes | no}

Enables or disables the account at creation time. By default, a cloud storage account is enabled on creation. To disable an account on creation, use this setting with the `no` option.

--account-id <string>

A required Amazon S3-only setting. The account ID number provided by Amazon when you first establish an account with the vendor.

--telemetry-bucket <string>

A required Amazon S3-only setting. The telemetry bucket name that you specified when you first established an account with the vendor. This is where usage reports are stored.

--storage-region <string>

An optional parameter for Amazon S3 or Google Storage Platform cloud types. The region value must match the storage region that you specified when you first established an account with the cloud provider. For example, `us-west-1`. If you do not specify a region, the cloud provider chooses its default region.

--skip-ssl-validation {yes | no}

Specifies whether to circumvent SSL certificate validation when connecting to a cloud provider's storage repository. Unless you specify this setting with a `yes` instruction, OneFS will attempt to perform SSL certificate validation when connecting. For security purposes, we recommend not enabling this setting. If you are connecting to a cloud provider that is within your corporate network (for example, Isilon or ECS), and you are having trouble connecting, you can skip SSL validation.

--enable-ocsp {yes | no}

Applies only to the C2S-S3 cloud type. It indicates whether to use OCSP to check the revocation status of the authentication certificate.

--ocsp-responder-url-required {yes | no}

Applies only to the C2S-S3 cloud type. It indicates whether a certificate without an OCSP responder URL is considered valid or not.

--proxy <string>

The network proxy through which CloudPools traffic to and from a public cloud provider should be redirected. The specified network proxy must already have been created with the `isi cloud proxies create` command.

--credential-provider-uri <string>

Applies only to the C2S-S3 cloud type. The URI to connect to a credential provider.

--credential-provider-agency <string>

Applies only to the C2S-S3 cloud type. The agency name required to connect to the credential provider.

--credential-provider-certificate <string>

Applies only to the C2S-S3 cloud type. The name or id of a certificate to connect to the credential provider.

--credential-provider-mission <string>

Applies only to the C2S-S3 cloud type. The Mission name required to connect to the credential provider.

--credential-provider-proxy <string>

Applies only to the C2S-S3 cloud type. The name or id of a proxy to connect to the credential provider.

--credential-provider-role <string>

Applies only to the C2S-S3 cloud type. The role name required to connect to the credential provider.

{--force | -f }

Execute the command without requiring confirmation.

{--verbose | -v}

Displays more detailed information.


Examples

The following example creates a Microsoft Azure cloud account:

```
isi cloud accounts create my_azure azure https://myazure.windows.net myuser dhgXJ9OAIahXvYmL
```

isi cloud accounts delete

Deletes a cloud storage account.

 **WARNING: Deleting an account results in the permanent loss of access to the data. In effect, you are deleting the data.**

Syntax

```
isi cloud accounts delete <id>  
[--acknowledge <string>]  
[--verbose]
```

Options

<id>

The name of the cloud account. You can use the `isi cloud accounts list` command to display the names of cloud accounts.

--acknowledge <string>

Enables the account deletion to proceed. This parameter is required. You must include a text string with the parameter, such as **yes**, **proceed**, or other string.

{--verbose | -v}

Displays more detailed information.

Example

The following example deletes a Microsoft Azure cloud account:

```
isi cloud accounts delete my_azure --acknowledge yes
```

When you run the command, OneFS displays the following message and requires confirmation:

```
*****  
WARNING: Deleting an account is extremely dangerous.  
Continuing with this operation will result in a permanent loss of data.  
Type 'confirm delete data' to proceed. Press enter to cancel:
```

To proceed, type **confirm delete data**, and press ENTER.

isi cloud accounts list

Lists cloud accounts.

Syntax

```
isi cloud accounts list
  [--limit <integer>]
  [--sort {id | name | type | account_username | uri | state | bucket}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of cloud accounts displayed in the list.

--sort {id | name | type | account_username | uri | state | bucket}

Sorts the list of cloud accounts according to the specified category. The following values are valid:

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--descending | -d}

Outputs the list of cloud accounts in descending order according to the specified sort option.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi cloud accounts modify

Modifies a cloud account.

Syntax

```
isi cloud accounts modify <id>
  [--name <string>]
  [--account-username <string>]
  [--key <string>]
  [--uri <string>]
  [--enabled {yes | no}]
  [--account-id <string>]
  [--telemetry-bucket <string>]
  [--storage-region <string>]
  [--skip-ssl-validation {yes | no}]
  [--enable-ocsp {yes | no}]
  [--ocsp-responder-url-required {yes | no}]
  [{--proxy <string> | --clear-proxy}]
  [--skip-account-check {yes | no}]
  [--credential-provider-uri <string>]
```

```
[--credential-provider-agency <string>]
[--credential-provider-certificate <string>]
[--credential-provider-mission <string>]
[--credential-provider-proxy <string>]
[--credential-provider-role <string>]
[--verbose]
```

Options

<id>

The ID of the cloud account. In this case, the ID is the same as the cloud account name.

{--name | -n} <string>

The name of the cloud account. In this case, the name is the same as the ID.

{--account-username | -u} <string>

The username for the cloud account. This name must be identical to the user name provided to the cloud vendor.

--key <string>

The cloud account access key or password. This information is provided by the cloud vendor.

--uri <string>

The cloud account URI. This URI must match that provided to the cloud vendor.

--enabled {yes | no}

By default, when you create a cloud storage account, it is enabled. To disable the account on creation, you can use this setting with the `no` option.

--account-id <string>

This is a required Amazon S3-only setting. The account ID number provided by Amazon when you first establish an account with the vendor.

--telemetry-bucket <string>

This is a required Amazon S3-only setting. The telemetry bucket name that you specified when you first established an account with the vendor.

--storage-region <string>

This is a required Amazon S3, Google Cloud Platform, Alibaba Cloud setting. The storage region that you specified when you first established an account with the vendor. For example, `us-west-1`.

--skip-ssl-validation {yes | no}

Specifies whether to circumvent SSL certificate validation when connecting to a cloud provider's storage repository. Unless you specify this setting with a `yes` instruction, OneFS will attempt to perform SSL certificate validation when connecting. For security purposes, we recommend not enabling this setting. If you are connecting to a cloud provider (for example, RAN or ECS) that is inside your corporate network, and you are having trouble connecting, you can skip SSL validation.

--enable-ocsp {yes | no}

Applies only to the C2S-S3 cloud type. It indicates whether to use OCSP to check the revocation status of the authentication certificate.

--ocsp-responder-url-required {yes | no}

Applies only to the C2S-S3 cloud type. It indicates whether a certificate without an OCSP responder URL is considered valid or not.

{--proxy <string> | --clear-proxy}

Use `--proxy` to set or change a network proxy through which CloudPools traffic is redirected, on its way to and from a public cloud provider. The specified network proxy must already have been created with the `isi cloud proxies create` command.

Use `--clear-proxy` to remove a previously set proxy. When you remove a proxy, CloudPools traffic flows directly to the cloud provider.

--skip-account-check {yes | no}

If set to `yes`, CloudPools skips the validation step to determine if the cloud storage account is still accessible. We do not recommend skipping this check.

{--verbose | -v}

Displays more detailed information.

Example

The following example modifies a Microsoft Azure cloud account:

```
isi cloud accounts modify my_azure
--uri https://myazure.windows.net
--account-username myuser --key dhgXJ90AIahXvYmL
```

isi cloud accounts view

Displays the details of a cloud account.

Syntax

```
isi cloud accounts view <id>
```

Options

<id>

Specifies the id of the cloud account to view. You can use the `isi cloud accounts list` command to obtain ids of available cloud accounts.

isi cloud archive

Queues one or more files to be archived to the cloud. For files to be archived, they must match the specified file pool policy, or any file pool policy with a cloud target.

Syntax

```
isi cloud archive <files>
[--recursive {yes | no}]
[--policy <string>]
[--verbose]
[--help]
```

Options

<files>

Specifies the files to archive. Specify `--files` for each additional file to process. Alternatively, you can specify a file matching pattern such as `/ifs/data/archive/images/*.jpg`.

{--recursive | -r} {yes | no}

Specifies whether the operation should apply recursively to nested directories in the file string.

--policy <string>

Specifies the file pool policy to apply to the specified files. If you specify one or more files to be archived and do not specify a policy, OneFS will compare the files with each configured file pool policy.

{--verbose | -v}

Displays more detailed information.

Examples

The following example archives multiple files to the cloud according to a specific file pool policy:

```
isi cloud archive /ifs/data/images/big.jpg --file /ifs/data/huge.jpg
--policy my_policy
```

The following example archives an entire directory to the cloud. The operation must match an existing file pool policy to be successful.

```
isi cloud archive /ifs/data/images/*.* --recursive yes
```

isi cloud jobs cancel

Cancels a CloudPools job initiated manually with `isi cloud archive` or `isi cloud recall`. You cannot cancel CloudPools system jobs (such as `cache-writeback`).

Syntax

```
isi cloud jobs cancel <id>
[--verbose]
```

Options

<id>

The ID for the cloud job. Run `isi cloud jobs list` to see a list of all manual and system jobs and their associated IDs.

{--verbose | -v}

Displays more detailed information.

Example

This following example cancels a CloudPools job with the ID of 21.

```
isi cloud job cancel 21
```

isi cloud jobs create

Creates a CloudPool job to archive or recall files.

Syntax

```
isi cloud jobs create {archive | recall}
[--files <string>]
[--begin-filter] <criteria> [--end-filter]
[--recursive {yes | no}]
[--accounts <string>]
[--expiration-date <timestamp>]
[--policy <string>]
[--verbose]
```

Options

{archive | recall}

The type of job. Valid entries are `archive` and `recall`.

--files <string>...

Specifies files or directory names to which the job applies. Use `--files` for each additional specification. Directory names are valid only for archive jobs.

--begin-filter <criteria> --end-filter

A file matching filter that defines a set of files to act on. For a description of `<criteria>` and valid operators to use in the filter, enter `man isi-file-matching` on the command line.

{--recursive | -r} {yes | no}

Specifies whether the job should apply recursively to nested directories.

--accounts <string>

For use only with guidance from Dell EMC Technical Support.

--expiration-date <timestamp>

The expiration date for orphan objects. Use one of the following formats for `<timestamp>`:

- A date string matching the pattern 'YYYY-MM-DD'
- A date/time string matching the pattern 'YYYY-MM-DD[Thh:mm[:ss]]'

--policy <string>

The policy to use in an archive job.

{--verbose | -v}

Displays more detailed information.

Examples

For archive jobs, you can specify one or more directories to archive. The following command archives a single directory:

```
isi cloud jobs create archive --files /ifs/shares/dir1
```

The following example archives multiple directories:

```
isi cloud jobs create archive --files /ifs/shares/dir1 --files /ifs/shares/dir2
```

isi cloud jobs files list

Displays the list of files matched by the specified CloudPools job.

Syntax

```
isi cloud jobs files list <job-id>
  [--offset <integer>]
  [--page <integer>]
  [--id <boolean>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<job-id>

The ID of the job. To find the list of job IDs in CloudPools, run the `isi cloud jobs list` command.

--offset <integer>

The starting file ID number to display.

--page *<integer>*

Used with limit option. If present, specifies the starting page number to display where page size is specified by limit. This option will be deprecated; please use offset option instead.

--id *<boolean>*

Adds an ID number in the display before each file.

{--limit | -l} *<integer>*

Display no more than the specified number of items.

--format {**table** | **json** | **csv** | **list**}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

The following example displays a list of files associated with a specific cloud job:

```
isi cloud jobs files list 21
```

isi cloud jobs list

Lists the status of CloudPools jobs, including system, archive, and recall jobs.

Syntax

```
isi cloud jobs list  
[--limit <integer>]  
[--sort {id | job_state | operation_state | effective_state | type  
| state_change_time | completion_time | create_time | description}]  
[--descending ]  
[--format {table | json | csv | list}]  
[--no-header]  
[--no-footer]  
[--verbose]
```

Options

{--limit | -d} *<integer>*

Displays no more than the specified number of items.

--sort {**id** | **job_state** | **operation_state** | **effective_state** | **type** | **state_change_time** | **completion_time** | **create_time** | **description**}

Orders results by this field. The default value is *id*. Note that, to sort on other than ID, description, effective state, and type, use the `--verbose` parameter with the command.

{--descending | -d}

Sorts and presents data in descending order.

--format {table | json | csv | list}

Display output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi cloud jobs pause

Pause a cloud job. You can resume a paused job with the `isi cloud jobs resume` command.

Syntax

```
isi cloud jobs pause <id>
[--verbose]
```

Options

id

The ID of the cloud job to pause. Use the `isi cloud jobs list` command to view the IDs of all cloud jobs. Although possible, we recommend that you not pause any of the CloudPools system jobs that run in the background and are critical for proper operation. These include:

ID	Description	Effective State	Type
1	Write updated data to the cloud	running	cache-writeback
2	Expire CloudPools cache	running	cache-invalidation
4	Clean up unreferenced data in the cloud	running	cloud-garbage-collection
5	Write updated snapshot data to the cloud	running	snapshot-writeback
6	Update SmartLink file formats	running	smartlink-upgrade
7	Add data to CloudPools cache	running	cache-pre-populate

{--verbose | -v}

Displays more detailed information.

Example

The following example pauses a cloud job with ID 19.

```
isi cloud jobs pause 19
```


isi cloud jobs resume

Resumes a paused cloud job.

Syntax

```
isi cloud jobs resume <id>  
[--verbose]
```

Options

<id>

The ID for the cloud job to resume. Use the `isi cloud jobs list` command to view a list of jobs and their associated IDs.

{--verbose | -v}

Displays more detailed information.

Example

The following command resumes a paused job with an ID of 26:

```
isi cloud jobs resume 26
```

isi cloud jobs view

Shows the details of a cloud job.

Syntax

```
isi cloud jobs view <id>
```

Options

<id>

Specify the ID of the cloud job. Use the `isi cloud jobs list` command to view all jobs and their associated IDs.

Example

The following command shows the details of a job with the ID of 27:

```
isi cloud jobs view 27
```

isi cloud pools create

Creates a CloudPool, which provides the connection between OneFS and a cloud storage account.

Syntax

```
isi cloud pools create <name> <type> <account>
  [--description <string>]
  [--vendor <string>]
  [--verbose]
```

Options

<name>

The name of the CloudPool.

<type>

The type of account, one of isilon, azure, s3, ecs, virtustream, or google.

<account>

The name of the cloud storage account to which the CloudPool connects. The cloud storage account is required and must match the CloudPool type. Only one cloud storage account can be specified.

--description <string>

A description of the CloudPool.

--vendor <string>

The name of the vendor hosting the cloud storage account.

{--verbose | -v}

Displays more detailed information.

Example

This following command creates a CloudPool containing a Microsoft Azure cloud storage account:

```
isi cloud pools create my_cp azure http://myazure.microsoft.com
--description="Financial records 2013" --vendor=Microsoft
```

isi cloud pools delete

Deletes a CloudPool. Proceed with caution. If you delete a CloudPool, OneFS is no longer able to access the associated cloud storage account. If the CloudPool is referenced by a file pool policy, OneFS does not allow you to delete the CloudPool.

Syntax

```
isi cloud pools delete <id>
  [--force]
  [--verbose]
```

Options

<id>

The name of the CloudPool. You can use the `isi cloud pools list` command to list existing CloudPools and their associated IDs.

{--force | -f}

Deletes the account without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

Example

The following command deletes a CloudPool:

```
isi cloud pool delete my_azure_pool
```

When you press ENTER to run the command, OneFS asks for confirmation. Type **yes**, and then press ENTER.

isi cloud pools list

Displays a list of CloudPools.

Syntax

```
isi cloud pools list
[--limit <integer>]
[--sort {id | name | type | description | vendor}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort {id | name | type | description | vendor}

Orders results by this field. The default value is `id`, which, in this case, is the same as `name`. Unless you use the `--verbose` option, you can only sort on `name` or `type`.

{--descending | -d}

Sorts and presents data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi cloud pools modify

Modifies a CloudPool.

Syntax

```
isi cloud pools modify <id>
  [--name <string>]
  [--accounts <string> | --clear-accounts | --add-accounts <string> | --remove-accounts
<string>]
  [--description <string>]
  [--vendor <string>]
  [--verbose]
```

Options

- <id>**
The ID of the CloudPool. Run `isi cloud pools list` to view the IDs of all CloudPools.
- {--name | -n} <string>**
Specifies a new name for the CloudPool.
- {--accounts <string> | --clear-accounts | --add-accounts <string> | --remove-accounts <string>}**
Adds or removes accounts associated with this cloudpool.
Only one account per CloudPool is allowed. To change the account associated with a CloudPool, we recommend the following:
- Create a new CloudPool using the `isi cloud pools create` command, specifying the correct account.
 - Delete the old CloudPool using the `isi cloud pools delete` command. Proceed with caution. If you delete a CloudPool, OneFS is no longer able to access the associated cloud storage account.
- {--description | -d} <string>**
Provides a description of this cloud pool.
- vendor <string>**
Provides the name of the vendor hosting the cloud pool account.
- {--verbose | -v}**
Displays more detailed information.

Examples

The following command adds a vendor name and description to an existing CloudPool:

```
isi cloud pools modify my_azure --vendor Microsoft
--description "preferred azure account"
```

The following command removes a cloud account from the CloudPool:

```
isi cloud pools modify my_s3 --remove-accounts s3_acct_1
```

isi cloud pools view

View detailed information about a CloudPool.

Syntax

```
isi cloud pools view <id>
```

Options

<id>

The ID of the cloud pool. Run the `isi cloud pool list` command to view all CloudPools and their associated IDs.

Example

The following command displays information about a CloudPool named `my_azure_pool`.

```
isi cloud pools view my_azure_pool
```

isi cloud proxies create

Creates a network proxy through which a cloud storage account can connect to a cloud storage provider.

Syntax

```
isi cloud proxies create <name> <host> <type> <port>  
  [--username <string>]  
  [--password <string>]  
  [--verbose]
```

Options

<name>

The name of the network proxy. This can be any alphanumeric string, but should be a simple, recognizable name.

<host>

The DNS name or IP address of the proxy server. For example, `myproxy1.example.com` or `192.168.107.107`.

<type>

The proxy protocol type, one of `socks_4`, `socks_5`, or `http`.

<port>

The port number to communicate with the proxy server. The correct port number depends on the port opened up on the proxy server for communication with CloudPools.

{--username | -u} <string>

The user name to authenticate with the SOCKS v5 or HTTP proxy server. Note that SOCKS v4 does not support authentication.

{--password | -p} <string>

The password to authenticate with the SOCKS v5 or HTTP proxy server.

{--verbose | -v}

Displays more detailed information.

Examples

The following example creates a network proxy to use with CloudPools:

```
isi cloud proxies create myproxy1 myprox1.example.com socks_5 1080
--username mycloudpools --password dhgXJ90AIahXvYmL
```

isi cloud proxies delete

Deletes a network proxy in CloudPools. Note that CloudPools prevents deletion of a proxy that is attached to a cloud storage account.

Syntax

```
isi cloud proxies delete <name>
[--force]
[--verbose]
```

Options

<name>

The unique id or name of the network proxy. You can use the `isi cloud proxies list` command to display the names of proxies.

{--force | -f}

Enables the proxy deletion to proceed without confirmation.

{--verbose | -v}

Displays more detailed information.

Example

The following example deletes a network proxy named `myproxy1`:

```
isi cloud accounts delete myproxy1
```

When you run the command, OneFS displays the following message and requires confirmation:

```
Are you sure? (yes/[no]):
```

To proceed, type **yes**, and press ENTER. If the proxy is attached to a cloud storage account, OneFS displays the following message:

```
Cannot delete proxy while used by accounts
```

isi cloud proxies list

Displays a list of network proxies created in CloudPools.

Syntax

```
isi cloud proxies list
[--limit <integer>]
[--sort {id | name | host | type | port}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of network proxies displayed in the list.

--sort {id | name | host | type | port}

Sorts the list of cloud proxies according to the specified category.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--descending | -d}

Outputs the list of network proxies in descending order according to the specified sort option.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

The following example creates a network proxy to use with CloudPools:

```
isi cloud proxies create myproxy1 myprox1.example.com socks_5 1080
--username mycloudpools --password dhgXJ90AIahXvYmL
```

isi cloud proxies modify

Modifies the properties of a network proxy.

Syntax

```
isi cloud proxies modify <name>
  [--name <string>]
  [--host <string>]
  [--type {socks_4 | socks_5 | http}]
  [--port <integer>]
  [--username <string>]
  [--clear-username]
  [--password <string>]
  [--clear-password]
  [--verbose]
```

Options

<name>

The current name of the network proxy.

{--name | -n} <string>

The new name of the network proxy. This can be any alphanumeric string, but should be a simple, recognizable name.

--host <string>

The DNS name or IP address of the proxy server. For example, `myproxy1.example.com` or `192.168.107.107`.

{--type | -t} {socks_4 | socks_5 | http}

The network proxy protocol, one of `socks_4`, `socks_5`, or `http`.

--port <integer>

The port number to communicate with the proxy server. The correct port number depends on the port opened up on the proxy server for communication with CloudPools.

{--username | -u} <string>

The user name to authenticate with the SOCKS v5 or HTTP proxy server. Note that SOCKS v4 does not support authentication.

--clear-username

Clear the user name that was previously specified for proxy server authentication.

{--password | -p} <string>

The password to authenticate with the SOCKS v5 or HTTP proxy server.

--clear-password

Clear the password that was previously specified for proxy server authentication.

{--verbose | -v}

Displays more detailed information.

Examples

The following example modifies a network proxy in CloudPools:

```
isi cloud proxies modify myproxy1 --type socks_4 --clear-username --clear-password
```

isi cloud proxies view

View the details of a network proxy created for CloudPools.

Syntax

```
isi cloud proxies view <name>
```

Options

<name>

Specifies the id or name of the network proxy to view. You can use the `isi cloud proxies list` command to display a list of the available proxies.

Example

The following example displays the details of a network proxy named `myproxy1`:

```
isi cloud proxies view myproxy1
```


isi cloud recall

Recalls files from the cloud, restoring the full files to their original directories. To make sure that the specified files are present in the cloud, OneFS scans the cluster for SmartLink files prior to performing the recall.

Syntax

```
isi cloud recall <files>
  [--begin-filter] <criteria> [--end-filter]
  [--recursive {yes | no}]
  [--verbose]
```

You can also provide a file matching filter to specify a set of files to act on. The basic syntax follows. For a full description on file matching criteria, use `man isi-file-matching` on the command line.

Options

<files>

The files to recall. For multiple specifications, use `--files` for each additional file name entry.

[--begin-filter] <criteria> [--end-filter]

A file matching filter that defines a set of files to act on. For a description of `<criteria>` and valid operators to use in the filter, enter `man isi-file-matching` on the command line.

{--recursive | -r} {yes | no}

Specifies whether the recall should apply recursively to nested subdirectories.

{--verbose | -v}

Displays more detailed information about the operation.

Usage

The `isi cloud recall` command restores the full file to its original directory, and overwrites the associated SmartLink file. If the file pool policy that originally archived the file to the cloud is still in effect, the next time the SmartPools job runs, the recalled file is archived to the cloud again. If you do not want the recalled file to be re-archived, you can move the file to a different directory that would not be affected by the file pool policy, or you can modify or delete the policy.

Examples

The following example recalls all files from the cloud for a directory and its subdirectories:

```
isi cloud recall /ifs/data/archives/archives2014/projects/*. *
--recursive yes
```

The command starts a cloud job. If you use the `--verbose` parameter, OneFS reports the job number, as in the following example:

```
Created job [29]
```

You can use the `isi cloud jobs view` command with the job number to see information about the job.

isi cloud restore_coi

Restores the cloud object index (COI) for a cloud storage account on the cluster. The `isi cloud access add` command also restores the COI for a cloud storage account.

Usage

 **WARNING: Do not execute this command unless instructed to do so by Isilon Technical Support.**

A cloud object index (COI) is a persistent mapping between cloud objects, their retention periods, and optionally, the files that use the cloud objects. The cluster uses the COI when performing cleanup (garbage collection), to ensure it considers all versions of files and objects correctly.

The `isi cloud restore coi` command allows a cluster to complete a COI to include all versions of all objects. The command might be used in the following situations:

- To handle COI corruption in cases where COI entries are corrupted or deleted. This command can restore the COI for a specified cloud account.
- To increase the retention time on the cluster where the command is run for objects in the specified cloud account.

Syntax

```
isi cloud restore coi
  [--accounts <string>]
  [--expiration-date <timestamp>]
  [--verbose]
```

Options

`--accounts <string>...`

The name of the cloud storage account whose COI you intend to restore. By restoring the COI, you enable OneFS to not only read data from the cloud, but also to write data to the cloud.

Use an additional `--accounts` parameter for each additional cloud account.

`--expiration-date <timestamp>`

The expiration date for orphaned cloud data objects.

`{--verbose | -v}`

Displays more detailed information about the operation.

Example

The following example restores the COI for a cloud storage account:

```
isi cloud restore_coi --account my_azure_acct
```

isi cloud settings modify

Controls archiving of snapshot files. By default, archiving of snapshots is enabled.

Use [isi cloud settings view](#) on page 164 to see the current settings.

Syntax

```
isi cloud settings modify
  [--default-accessibility {cached | no-cache}]
  [--default-cache-expiration <duration>]
  [--default-compression-enabled {yes | no}]
  [--default-data-retention <duration>]
  [--default-encryption-enabled {yes | no}]
  [--default-full-backup-retention <duration>]
  [--default-incremental-backup-retention <duration>]
  [--default-read-ahead <string>]
  [--default-writeback-frequency <duration>]
  [--verbose]
```

Options

--default-accessibility {cached | no-cache}

Specifies whether, when a SmartLink file is accessed, cloud data is incrementally downloaded (cached) as needed, or fully downloaded (not cached).

--default-cache-expiration <duration>

The minimum amount of time until the cache expires. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-compression-enabled {yes | no}

Specifies whether data is to be compressed when archived to the cloud.

--default-data-retention <duration>

The minimum amount of time that cloud objects associated with a SmartLink file will be retained in the cloud after the SmartLink file is deleted from the cluster. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-encryption-enabled {yes | no}

Specifies whether data is to be encrypted when archived to the cloud.

--default-full-backup-retention <duration>

The length of time that OneFS retains cloud data referenced by a SmartLink file that has been backed up by a full NDMP backup and is subsequently deleted. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--default-incremental-backup-retention <duration>

The length of time that OneFS retains cloud data referenced by a SmartLink file that has been backed up by an incremental NDMP backup, or replicated by a SyncIQ operation, and is subsequently deleted. A number followed by a unit of time is accepted. For example, a setting of 5Y would specify a five-year duration.

--default-read-ahead {partial | full}

Specifies the cache readahead strategy when SmartLink files are accessed. A partial strategy means that only the amount of data needed by the user is cached. A full strategy means that all file data will be cached when the user accesses a SmartLink file.

--default-writeback-frequency <duration>

The minimum amount of time to wait before OneFS updates cloud data with local changes. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

{--verbose | -v}

Displays more information about the operation.

Example

The following example modifies several of the default CloudPools settings:

```
isi cloud settings modify --default-writeback-frequency 12H
--default-cache-expiration 9H --default-accessability no-cache
--default-encryption-enabled yes
```

isi cloud settings regenerate-encryption-key

Generates a new master encryption key for new data that will be archived to the cloud. Previously encrypted archived data continues to require previously generated encryption keys. All previous encryption keys are preserved for use with the existing archived data.

Syntax

```
isi cloud settings regenerate-encryption-key
[--verbose]
```

Option

{--verbose | -v}

Displays more detailed information.

isi cloud settings view

Displays the current default settings in CloudPools. You can use the `isi cloud settings modify` command to change default settings.

Syntax

```
isi cloud settings view
```

Options

There are no options for this command.

Example

The following example shows sample output. Explanations of the displayed properties are included in the descriptions for [isi cloud settings modify](#) on page 162.

```
B248930-PSL-1# isi cloud settings view
  Default Accessibility: cached
  Default Cache Expiration: 1D
  Default Compression Enabled: No
  Default Data Retention: 1W
  Default Encryption Enabled: No
  Default Full Backup Retention: 5Y
  Default Incremental Backup Retention: 5Y
  Default Read Ahead: partial
  Default Writeback Frequency: 9H
```

isi cluster atime modify

Modify the cluster access time grace period.

Syntax

```
isi cluster atime modify
  [--enabled <boolean>]
  [--precision <integer>]
  [--verbose]
```

Options

--enabled {yes | no} Specifies whether the access time grace period is enabled.

--precision <integer> The amount of time between file access time updates, in seconds.

{--verbose | -v} Displays more detailed information.

isi cluster atime view

View file system access time values.

Syntax

```
isi cluster atime view
```

Options

None

OneFS displays whether cluster access time grace periods are enabled, and the prevision interval, in seconds.

isi cluster contact modify

Modify primary and secondary cluster contact information.

Syntax

```
isi cluster contact modify
  [--company <string>]
  [--location <string>]
  [--primary-name <string>]
  [--primary-email <string>]
  [--primary-phone1 <string>]
  [--primary-phone2 <string>]
  [--secondary-name <string>]
  [--secondary-email <string>]
  [--secondary-phone1 <string>]
  [--secondary-phone2 <string>]
  [--verbose]
```

Options

- {--company | -c}** The company name for the cluster contact information.
- {--location | -o}** The company location for the cluster contact information.
- {--primary-name | -n}** The name of the primary administration contact.
- {--primary-email | -e}** The email address for the primary administration contact.
- {--primary-phone1 | -p}** The main phone number for the primary administration contact.
- {--primary-phone2 | -P}** Alternate phone number for the primary administration contact.
- {--secondary-name | -N}** The name of the secondary administration contact.
- {--secondary-email | -E}** The email address for the secondary administration contact.
- {--secondary-phone1 | -s}** The main phone number for the secondary administration contact.

{--secondary-phone2 | -S} Alternate phone number for the secondary administration contact.
{--verbose | -v} Displays more detailed information.

isi cluster contact view

View primary and secondary cluster administration contact information.

Syntax

```
isi cluster contact view  
[--format (list | json)]
```

Options

--format (list | json) Display cluster contacts in list or JSON format.

OneFS displays primary and secondary cluster contact administration information.

isi cluster encoding list

List the supported file system character encoding formats for OneFS.

Syntax

```
isi cluster encoding list
```

Options

None

OneFS displays a list of supported file system character encoding formats.

isi cluster encoding modify

Modify the current file system character encoding selection.

Syntax

```
isi cluster encoding modify <current-encoding>  
[--verbose]
```

Options

<current-encoding> Enter the supported character encoding selection

{--verbose | -v} Displays more detailed information.

isi cluster encoding view

View the selected file system character encoding type.

Syntax

```
isi cluster encoding view
```

Options

None

OneFS displays the selected file system character encoding type.

isi cluster identity modify

Modify the name, description, and message of the day (MOTD) for the cluster.

Syntax

```
isi cluster identity modify
  [--description <string>]
  [--motd <string>]
  [--motd-header <string>]
  [--name <string>]
  [--verbose]
```

Options

- {--description | -d}** Text description of the cluster.
- {--motd | -m}** The cluster message of the day.
- {--motd-header | -M}** The cluster message of the day header text.
- {--name | -n}** The cluster name specified at installation.
- {--verbose | -v}** Displays more detailed information.

isi cluster identity view

View the cluster name, description, and message of the day (MOTD).

Syntax

```
isi cluster identity view
  [--format (list | json)]
```

Options

--format (list | json) Display cluster contacts in list or JSON format.

OneFS displays the cluster name, description, and message of the day.

isi cluster internal-networks modify

Modify the internal network interface connection for the cluster.

Syntax

```
isi cluster internal-networks modify
  [--int-a-ip-addresses <ip_address_range>]
  [--int-b-ip-addresses <ip_address_range>]
  [--failover-ip-addresses <ip_address_range>]
  [--failover-status <boolean>]
  [--int-a-prefix-length <integer>]
  [--int-b-prefix-length <integer>]
```

Options

- int-a-ip-addresses <ip_address_range>**
The IP address range for the int-a network interface. To specify an IP address range, enter two IP addresses separated by a hyphen (-), such as **10.7.0.0-10.7.255.255**.
- int-b-ip-addresses <ip_address_range>**
The IP address range for the int-b network interface. To specify an IP address range, enter two IP addresses separated by a hyphen (-), such as **10.7.0.0-10.7.255.255**.
- failover-ip-addresses <ip_address_range>**
The IP address range for the failover network interface. To specify an IP address range, enter two IP addresses separated by a hyphen (-), such as **10.7.0.0-10.7.255.255**.
- failover-status (enabled | disabled)**
The status of the int-b network interface.
- int-a-prefix-length <integer>**
The prefix size for the int-a network interface.
- int-b-prefix-length <integer>**
The prefix size for the int-b network interface.

isi cluster internal-networks view

View internal network configuration settings for the cluster.

Syntax

```
isi cluster internal-networks view
```


Options

None

OneFS displays cluster internal network configuration settings.

isi cluster join-mode modify

Modify cluster join mode security settings.

Syntax

```
isi cluster join-mode modify <mode>
```

Options

<mode> (**secure** | **manual**) Set the cluster join mode to secure or manual.

isi cluster join-mode view

View cluster join mode security settings.

Syntax

```
isi cluster join-mode view
```

Options

None

OneFS displays the security setting for cluster join mode.

isi cluster lnnset modify

Modify logical node names (LNNs) for the cluster

Syntax

```
isi cluster lnnset modify <lnns>
```

Options

<lnns>

The logical node names (LNNs) to be modified. Specify the current LNN and then the desired new LNN, separated by a hyphen (-). The new LNN must not currently be in use. For example, to rename the node with the current LNN from 1 to 9, enter:

```
isi cluster lnnset modify 1-9
```

isi cluster lnnset view

View all cluster logical node names (LNNs) and their associated device IDs and IP addresses, or view that information for a specific LNN.

Syntax

```
isi cluster lnnset view  
  [--lnn <integer>]
```


Options

--lnn <integer> Specify a LNN to view information about only that LNN.

OneFS displays LNN names and their associated device IDs and IP addresses.

isi cluster reboot

Reboot one or all nodes in a cluster.

 **NOTE:** To perform a rolling reboot, run the `isi upgrade rolling-reboot` command.

Syntax

```
isi cluster reboot  
  [--node-lnn <string>]
```

Options

--node-lnn <string> The LNN of the node to reboot, or `all` to reboot all nodes. If omitted, OneFS selects the local node.

isi cluster shutdown

Shut down one or all nodes of a cluster.

Syntax

```
isi cluster shutdown  
  [--node-lnn <string>]
```

Options

--node-lnn <string> The LNN of the node to shut down, or `all` to shut down all nodes. If omitted, OneFS selects the local node.

isi cluster time modify

Modify the date and time for the cluster.

Syntax

```
isi cluster time modify <time>
```

Options

<time> The date and time in yyyy-mm-ddThh:mm:ss format.

isi cluster time view

View the current cluster date and time.

Syntax

```
isi cluster time view
```

Options

None

OneFS displays the cluster date and time in yyyy-mm-ddThh:mm:ss format.

isi cluster time timezone modify

Modify the cluster time zone.

Syntax

```
isi cluster time timezone modify  
  [--abbreviation <string>]  
  [--path <string>]  
  [--force]  
  [--verbose]
```

Options

 **NOTE:** If you wish to use an interactive time zone selection tool, do not enter values for `--abbreviation` or `--path`.

- {--abbreviation | -a} <string>** An abbreviation for the cluster time zone, such as **PDT** for United States Pacific Daylight Time. Do not enter a value if you wish to use the interactive time zone selector.
- {--path | -p} <string>** A time zone hierarchical path, such as **America/Los_Angeles** for United States Pacific Daylight Time (PDT). Do not enter a value if you wish to use the interactive time zone selector.
- {--force | -f}** Do not prompt for confirmation of any modifications.
- {--verbose | -v}** Display more detailed information.

isi cluster time timezone view

View the cluster time zone.

Syntax

```
isi cluster time timezone view
```

Options

None

OneFS displays the cluster time zone abbreviation and path.

isi compression stats list

View statistics related to data compression in a list of historical five minute intervals beginning with the moment you run the command.

Syntax

```
isi compression stats list
  [--begin <integer>]
  [--end <integer>]
  [--resolution <integer>]
  [--pretty-time]
  [--local]
  [--limit | -l] <integer>]
  [--format (table | json | csv | list)]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose | -v]
  [--help | -h]
```

Options

--begin <time>

Specifies begin time in UNIX Epoch timestamp format.

--end <time>

Specifies end time in UNIX Epoch timestamp format.

--resolution <integer>

Specifies the minimum interval between series data points in seconds.

--pretty-time

Displays the timestamp as a readable string.

--local

Show statistics for the local node only.

{--limit | -l} <integer>

Limits the number of compression statistics to display.

--format (table | json | csv | list)

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--noheader | -a}

Displays data without headers.

{--no-footer | -z}

Displays data without footer information

{--verbose | -v}

Displays more detailed information.

{--help | -h}

Displays help for this command.

isi compression stats view

View statistics related to data compression that were recorded over the last five minutes.

Syntax

```
isi compression stats view  
  [--resolution <integer>]  
  [--local]  
  [--help | -h]
```

Options

--resolution <integer>

Specifies the minimum interval between series data points in seconds.

--local

Shows statistics for the local node only.

{--help | -h}

Displays help for this command.

isi compression settings modify

Modifies the global configuration setting for data compression.

Syntax

```
isi compression settings modify  
  [--enabled <boolean>]  
  [--verbose | -v]  
  [--help | -h]
```

Options

--enabled <boolean>

Determines whether data compression is enabled or disabled.

This setting only applies for data that resides on compatible hardware (i.e. F810).

{--verbose | -v}

Displays more detailed information.

{--help | -h}

Displays help for this command.

isi compression settings view

Displays the global configuration setting for data compression. The output will either be `Enabled: Yes` or `Enabled: No`.

Syntax

```
isi compression settings view
```

isi config

Opens a new prompt where node and cluster settings can be altered.

The command-line prompt changes to indicate that you are in the `isi config` subsystem. While you are in the `isi config` subsystem, other OneFS commands are unavailable and only `isi config` commands are valid.

Syntax

```
isi config
```

NOTE:

- **The following commands are not recognized unless you are currently at the `isi config` command prompt.**
- **Changes are not applied until you run the `commit` command.**
- **Some commands require you to restart the cluster.**

Commands

`changes`

Displays a list of changes to the configuration that have not been committed.

`commit`

Commits configuration settings and then exits `isi config`.

`date <time-and-date>`

Displays or sets the current date and time on the cluster.

<time-and-date> Sets cluster time to the time specified.

Specify **<time-and-date>** in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify **<time>** as one of the following values.

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
h	Specifies hours
s	Specifies seconds


`deliprange [<interface-name> [<ip-range>]]`

Displays a list of internal network IP addresses that can be assigned to nodes or removes specified addresses from the list.

- <interface-name>** Specifies the name of the interface as one of the following values:
 - int-a
 - int-b
 - failover
- <ip-range>** Specifies the range of IP addresses that can no longer be assigned to nodes. Specify in the form <lowest-ip>-<highest-ip>.

encoding [list] [<encoding>]

Sets the default encoding character set for the cluster.

 **CAUTION: Character encoding is typically established during installation of the cluster. Incorrectly modifying character encoding settings may render files unreadable. Modify settings only if necessary and after consultation with Isilon Technical Support.**

- list** Displays the list of supported character sets.

exit

Exits the `isi config` subsystem.

help

Displays a list of all `isi config` commands. For information about specific commands, the syntax is `help [<command>]`.

interface <interface-name> {enable | disable}

The interface command displays the IP ranges, netmask, and MTU and enables or disables internal interfaces. When issued with no argument, this command displays IP ranges, netmask, and MTU of all interfaces. When issued with an <interface-name> argument, this command displays IP ranges, netmask, and MTU for only the specified interface.

- {enable | disable}** Enables or disables the specified interface.
- <interface-name>** Specifies the name of the interface as `int-a` or `int-b`.

iprange [<interface-name> [<lowest-ip>-<highest-ip>]]

Displays a list of internal IP addresses that can be assigned to nodes, or adds addresses to the list.

- <interface-name>** Specifies the name of the interface as `int-a`, `int-b`, or `failover`.
- <lowest-ip>-<highest-ip>** Specifies the range of IP addresses that can be assigned to nodes.

ipset

Obsolete. Use `lnnset` to renumber cluster nodes. The IP address cannot be set manually.

joinmode [<mode>]


Displays the setting for how nodes are added to the current cluster. Options <mode> specifies the cluster add node setting as one of the following values.

- manual** Configures the cluster so that joins can be initiated by either the node or the cluster.
- secure** Configures the cluster so that joins can be initiated by only the cluster.

lnnset [<old-lnn> <new-lnn>]

Displays a table of logical node number (LNN), device ID, and internal IP address for each node in the cluster when run without arguments. Changes the LNN when specified.

- <old lnn>** Specifies the old LNN that is to be changed.
- <new lnn>** Specifies the new LNN that is replacing the old LNN value for that node.

 **NOTE: The new LNN must not be currently assigned to another node. Users logged in to the shell or web administration interface of a node whose LNN is changed must log in again to view the new LNN.**

migrate [<interface-name> [[<old-ip-range>] [<new-ip-range> [-n <netmask>]]]]

Displays a list of IP address ranges that can be assigned to nodes or both adds and removes IP ranges from that list.

- <interface-name>** Specifies the name of the interface as `int-a`, `int-b`, and `failover`.
- <old-ip-range>** Specifies the range of IP addresses that can no longer be assigned to nodes. If unspecified, all existing IP ranges are removed before the new IP range is added. Specify in the form of `<lowest-ip>-<highest-ip>`.
- <new-ip-range>** Specifies the range of IP addresses that can be assigned to nodes. Specify in the form of `<lowest-ip>-<highest-ip>`.
- n <netmask>** Specifies a new netmask for the interface.

NOTE: If more than one node is given a new IP address, the cluster reboots when the change is committed. If only one node is given a new IP address, only that node is rebooted.

mtu [`<value>`]

Displays the size of the maximum transmission unit (MTU) that the cluster uses for internal network communications when run with no arguments. Sets a new size of the MTU value, when specified. This command is for the internal network only.

NOTE: This command is not valid for clusters with an InfiniBand back end.

- <value>** Specifies the new size of the MTU value. Any value is valid, but not all values may be compatible with your network. The most common settings are 1500 for standard frames and 9000 for jumbo frames.

name [`<new_name>`]

Displays the names currently assigned to clusters when run with no arguments. Assigns new names to clusters, as specified.

- <new name>** Specifies a new name for the cluster.

netmask [`<interface-name>`] [`<ip-mask>`]

Displays the subnet IP mask that the cluster is currently using or sets new subnet IP masks, as specified. Specifies the interface name as `int-a` and `int-b`.

- <interface-name>** Specifies the name of the interface. Valid names are `int-a` and `int-b`.
- <ip-mask>** Specifies the new IP mask for the interface.

quit

Exits the `isi config` subsystem.

reboot [{`<node_lnn>` | `a11`}]

Reboots one or more nodes, specified by LNN. If no nodes are specified, reboots the node from which the command is run. To reboot the cluster, specify `a11`.

NOTE: If run on an unconfigured node, this command does not accept any arguments.

remove

Deprecated. Instead, run the `isi devices -a smartfail` command.

shutdown [{`<node_lnn>` | `a11`}]

Shuts down one or more nodes, specified by LNN. If no nodes are specified, shuts down the node from which the command is run. To shut down the cluster, specify `a11`.

NOTE: If run on an unconfigured node, this command does not accept any arguments.

status [`advanced`]

Displays current information on the status of the cluster. To display additional information, including device health, specify `advanced`.

timezone [`<timezone identifier>`]

Displays the current time zone or specifies new time zones. Specifies the new timezone for the cluster as one of the following values:

***<timezone
identifier>***

Specifies the new time zone for the cluster as one of the following values:

Greenwich Mean Time

Eastern Time Zone

Central Time Zone

Mountain Time Zone

Pacific Time Zone

Arizona

Alaska

Hawaii

Japan

Advanced. Opens a prompt with more time zone options.

version

Displays information about the current OneFS version.

wizard

Activates a wizard on unconfigured nodes and reactivates the wizard if you exit it during the initial node configuration process. The wizard prompts you through the node-configuration steps.

OneFS isi commands D through L

This chapter contains documentation of the OneFS CLI commands `isi dedupe reports list` through `isi license view`.

Topics:

- `isi dedupe inline settings modify`
- `isi dedupe inline settings view`
- `isi dedupe reports list`
- `isi dedupe reports view`
- `isi dedupe settings modify`
- `isi dedupe settings view`
- `isi dedupe stats`
- `isi devices add`
- `isi devices config modify`
- `isi devices config view`
- `isi devices drive add`
- `isi devices drive firmware list`
- `isi devices drive firmware update list`
- `isi devices drive firmware update start`
- `isi devices drive firmware update view`
- `isi devices drive firmware view`
- `isi devices format`
- `isi devices list`
- `isi devices node add`
- `isi devices node list`
- `isi devices node smartfail`
- `isi devices node stopfail`
- `isi devices purpose`
- `isi devices purposelist`
- `isi devices smartfail`
- `isi devices stopfail`
- `isi devices suspend`
- `isi devices view`
- `isi diagnostics gather settings modify`
- `isi diagnostics gather settings view`
- `isi diagnostics gather start`
- `isi diagnostics gather status`
- `isi diagnostics gather stop`
- `isi diagnostics netlogger settings modify`
- `isi diagnostics netlogger settings view`
- `isi diagnostics netlogger start`
- `isi diagnostics netlogger status`
- `isi diagnostics netlogger stop`
- `isi email settings modify`
- `isi email settings view`
- `isi esrs modify`
- `isi esrs view`
- `isi esrs dataitems list`
- `isi esrs dataitems modify`
- `isi esrs download list`
- `isi esrs download start`

- isi esrs download view
- isi esrs telemetry modify
- isi esrs telemetry view
- isi event alerts create
- isi event alerts delete
- isi event alerts list
- isi event alerts modify
- isi event alerts view
- isi event channels create
- isi event channels delete
- isi event channels list
- isi event channels modify
- isi event channels view
- isi event events list
- isi event events view
- isi event groups bulk
- isi event groups list
- isi event groups modify
- isi event groups view
- isi event settings modify
- isi event settings view
- isi event test create
- isi event types list
- isi fc settings list
- isi fc settings modify
- isi fc settings view
- isi file-filter settings modify
- isi file-filter settings view
- isi filepool apply
- isi filepool default-policy modify
- isi filepool default-policy view
- isi filepool policies create
- isi filepool policies delete
- isi filepool policies list
- isi filepool policies modify
- isi filepool policies view
- isi filepool templates list
- isi filepool templates view
- isi_for_array
- isi ftp settings modify
- isi ftp settings view
- isi_gather_info
- isi get
- isi hardening apply
- isi hardening revert
- isi hardening status
- isi hdfs crypto encryption-zones create
- isi hdfs crypto settings modify
- isi hdfs crypto encryption-zones list
- isi hdfs crypto settings view
- isi hdfs fsimage job settings modify
- isi hdfs fsimage job settings view
- isi hdfs fsimage job view
- isi hdfs fsimage latest delete
- isi hdfs fsimage latest view
- isi hdfs fsimage settings modify
- isi hdfs fsimage settings view

- isi hdfs inotify settings modify
- isi hdfs inotify settings view
- isi hdfs inotify stream reset
- isi hdfs inotify stream view
- isi hdfs log-level modify
- isi hdfs log-level view
- isi hdfs proxyusers create
- isi hdfs proxyusers delete
- isi hdfs proxyusers list
- isi hdfs proxyusers members list
- isi hdfs proxyusers modify
- isi hdfs proxyusers view
- isi hdfs racks create
- isi hdfs racks delete
- isi hdfs racks list
- isi hdfs racks modify
- isi hdfs racks view
- isi hdfs ranger-plugin settings modify
- isi hdfs ranger-plugin settings view
- isi hdfs settings modify
- isi hdfs settings view
- isi http settings modify
- isi http settings view
- isi job events list
- isi job jobs cancel
- isi job jobs list
- isi job jobs modify
- isi job jobs pause
- isi job jobs resume
- isi job jobs start
- isi job jobs view
- isi job policies create
- isi job policies delete
- isi job policies list
- isi job policies modify
- isi job policies view
- isi job reports list
- isi job reports view
- isi job statistics view
- isi job status
- isi job types list
- isi job types modify
- isi job types view
- isi license add
- isi license generate
- isi license list
- isi license view

isi dedupe inline settings modify

Globally control inline data deduplication. This setting applies only to data that resides on compatible hardware, such as F810 nodes.

Syntax

```
isi dedupe inline settings modify
  [--mode ] enabled | assess | paused | disabled]
  [--verbose | -v]
```

Options

mode

Set the specified inline data deduplication mode. Mode settings are cluster-wide.

- `enabled` turns on inline data deduplication.
- `assess` allows you to assess the potential space savings from inline data deduplication.
- `paused` deactivates inline data deduplication but leaves the index table intact.
- `disabled` deactivate inline data deduplication and de-allocates the index table.

{--verbose | -v}

Displays more detailed information.

Examples

The following command enables inline data deduplication in assessment mode:

```
isi dedupe inline settings modify --assess
```

The following command enables inline data deduplication:

```
isi dedupe inline settings modify --mode enabled
```

isi dedupe inline settings view

Displays current inline data deduplication settings.

Syntax

```
isi dedupe settings view
```

Options

There are no options for this command.

isi dedupe reports list

Displays a list of deduplication reports.

Syntax

```
isi dedupe reports list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

Examples

To view a list of deduplication reports, run the following command:

```
isi dedupe reports list
```

The system displays output similar to the following example:

Time	Job ID	Job Type
2013-05-09T11:03:37	4	Dedupe
2013-05-10T00:02:27	8	Dedupe
2013-05-15T13:03:47	12	Dedupe
2013-05-16T00:02:32	16	Dedupe
2013-05-17T00:02:32	19	Dedupe
2013-05-09T16:14:04	5	DedupeAssessment

Total: 6

isi dedupe reports view

Displays a deduplication report.

Syntax

```
isi dedupe reports view <job-id>
```

Options

<job-id>

Displays the deduplication report for the deduplication job of the specified ID.

Examples

The following command displays a deduplication job:

```
isi dedupe reports view 12
```

The system displays output similar to the following example:

```
Time: 2013-10-14T09:39:22
Job ID: 52
Job Type: Dedupe
Reports
  Time : 2013-10-14T09:39:22
  Results :
Dedupe job report:{
  Start time = 2013-Oct-14:09:33:34
  End time = 2013-Oct-14:09:39:22
  Iteration count = 1
  Scanned blocks = 1716
  Sampled blocks = 78
  Deduped blocks = 1425
  Dedupe percent = 83.042
  Created dedupe requests = 65
  Successful dedupe requests = 65
  Failed dedupe requests = 0
  Skipped files = 0
  Index entries = 38
  Index lookup attempts = 38
  Index lookup hits = 0
}
Elapsed time:          347 seconds
Aborts:               0
Errors:               0
Scanned files:        6
Directories:          2
2 paths:
/ifs/data/dir2,
/ifs/data/dir1
CPU usage:             max 29% (dev 2), min 0% (dev 1), avg 6%
Virtual memory size:  max 128388K (dev 1), min 106628K (dev 1), avg 107617K
Resident memory size: max 27396K (dev 1), min 9980K (dev 2), avg 11585K
Read:                 2160 ops, 124437504 bytes (118.7M)
Write:                30570 ops, 222851584 bytes (212.5M)
```

isi dedupe settings modify

Modifies the settings of deduplication jobs.

Syntax

```
isi dedupe settings modify
  [{[--paths <path>]... | --clear-paths]}
  [--add-paths <path>]...
  [--remove-paths <path>]...
  [{[--assess-paths <path>]... | --clear-assess-paths]
  [--add-assess-paths <path>]...
  [--remove-assess-paths <path>]...
  [--verbose]
```

Options

--paths <path>

Deduplicates files located under the specified root directories.

--clear-paths

Stops deduplication for all previously specified root directories. If you run the `isi dedupe settings modify` command with this option, you must run the command again with either `--paths` or `--add-path` to resume deduplication.

--add-paths <path>

Deduplicates files located under the specified root directory in addition to directories that are already being deduplicated.

--remove-paths <path>

Stops deduplicating the specified root directory.

--assess-paths <path>

Assesses how much space will be saved if files located under the specified root directories are deduplicated.

--clear-assess-paths

Stops assessing how much space will be saved if previously specified root directories are deduplicated. If you run the `isi dedupe settings modify` command with this option, you must run the command again with either `--paths` or `--add-path` to resume deduplication.

--add-assess-paths <path>

Assesses how much space will be saved if the specified root directories are deduplicated in addition to directories that are already being assessed.

--remove-assess-paths <path>

Stops assessing how much space will be saved if the specified root directories are deduplicated.

{--verbose | -v}

Displays more detailed information.

Examples

The following command starts deduplicating `/ifs/data/active` and `/ifs/data/media`:

```
isi dedupe settings modify --add-paths /ifs/data/active,/ifs/data/media
```

The following command stops deduplicating `/ifs/data/active` and `/ifs/data/media`:

```
isi dedupe settings modify --remove-paths /ifs/data/active,/ifs/data/media
```

isi dedupe settings view

Displays current deduplication settings.

Syntax

```
isi dedupe settings view
```

Options

There are no options for this command.

isi dedupe stats

Displays information about how much data is being saved by deduplication.

Syntax

```
isi dedupe stats
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

Examples

To view information about deduplication space savings, run the following command:


```
isi dedupe stats
```

The system displays output similar to the following example:

```
Cluster Physical Size: 17.019G
Cluster Used Size: 4.994G
Logical Size Deduplicated: 13.36M
Logical Saving: 11.13M
Estimated Size Deduplicated: 30.28M
Estimated Physical Saving: 25.23M
```

isi devices add

Defaults to `isi devices drive add`. Scans for available drives and adds the drives to the node.

 **NOTE:** You can add available nodes to a cluster by running the command `isi devices node add`.

Syntax

```
isi devices add {<bay> | all}
  [--sled <string>]
  [--node-lnn <integer>]
```

```
[--force]
[--verbose]
```

Options

{<bay> | all }

Specifies the bay number that contains the drive to be added to the node. You can specify `all` to scan the entire node.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

--node-lnn <integer>

Specifies the node number to scan for new drives. If omitted, the local node will be scanned.

{--force | -f}

Adds the drive or drives without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices config modify

Modifies a node's Automatic Replacement Recognition (ARR) status.

Syntax

```
isi devices config modify
  [--automatic-replacement-recognition {yes | no}]
  [--instant-secure-erase {yes | no}]
  [--node-lnn {all | <string>}]
  [--verbose]
```

Options

--automatic-replacement-recognition {yes | no}

Changes the ARR status for a cluster or specific node. A value of `yes` will enable ARR, a value of `no` will disable ARR.

--instant-secure-erase {yes | no}

Enable or disable the Instant Secure Erase (ISE) feature on the device. A value of `yes` will enable ISE, a value of `no` will disable ISE.

--node-lnn {all | <string>}

Specifies the node LNN for which to modify the device configuration. You can specify `all` to modify configuration of all nodes. If omitted, all nodes will be modified.

{--verbose | -v}

Displays more detailed information.

isi devices config view

Displays the Automatic Replacement Recognition (ARR) status of a node.

Syntax

```
isi devices config view
  [--node-lnn {all | <string>}]
```

Options

--node-lnn {all | <string>}

Specifies the node you want to view. You may specify all nodes. If omitted, ARR status for the local node is displayed.

isi devices drive add

Scans one or more bays for available drives and adds the drives to the node.

Syntax

```
isi devices drive add {<bay> | --sled <string>}
  [--node-lnn <integer>]
  [--force]
  [--verbose]
```

Options

{<bay> | all}

Specifies the bay number that contains the drive to be added to the node. You can specify `all` to scan the entire node.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

--node-lnn <integer>

Specifies the node number to scan for new drives. If omitted, the local node will be scanned.

{**--force** | **-f**}

Adds the drive or drives without asking for confirmation.

{**--verbose** | **-v**}

Displays more detailed information.

isi devices drive firmware list

Displays a list of firmware details for the data drives in a node.

Syntax

```
isi devices drive firmware list
  [--node-lnn <string>]
  [--sled <string> ]
  [{--summary | -s}]
```

```
[--override | -v]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--node-lnn {all | <integer>}

Specifies the node number of the drives you would like to display firmware information for. You may specify all nodes. If omitted, only the drive firmware information for the local node will be displayed.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

{--override | -v}

Uses the legacy node designations instead of the grid XY values.

{--summary | -s}

Displays a summary of drive firmware counts by model and revision.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi devices drive firmware update list

Displays the status of firmware updates on the cluster.

Syntax

```
isi devices drive firmware update list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi devices drive firmware update start

Updates firmware on one or more drives in a node.

Syntax

```
isi devices drive firmware update start <bay>
  [--node-lnn <integer>]
  [--force]
  [--verbose]
```

Options

{<bay> | all}

Specifies the bay number that contains the drive to be updated. You can specify `all` to update every drive in the node.

--node-lnn <integer>

Specifies the node number on which to update drives. If omitted, drives will be updated in the local node.

{--force | -f}

Updates the drive or drives without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices drive firmware update view

Displays information about a drive firmware update for a node.

Syntax

```
isi devices drive firmware update view
  [--node-lnn <integer>]
```

Options

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that is running the firmware update you want to view. If omitted, firmware update status for the local node will be displayed.

isi devices drive firmware view

Displays information about the firmware on a single drive.

Syntax

```
isi devices drive firmware view {<bay> | --lnum <integer>}
  [--node-lnn <integer>]
```

Options

{<bay> | --lnum <integer>}

Specifies the bay number or LNUM (logical drive number) of the drive to view.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to view. If omitted, the drive in the local node will be displayed.

isi devices format

Defaults to `isi devices drive format`. Formats a drive so you can add it to a node.

Syntax

```
isi devices format {<bay> | --sled <string>}
  [--node-lnn <integer>]
  [--purpose <string>]
  [--force | -f]
```

Options

<bay>

Specifies the bay number that contains the drive to be formatted.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to format. If omitted, the specified drive in the local node will be formatted.

--purpose <string>

Specifies the purpose to assign to the new drive. You can view a list of the possible drive purposes by running `isi devices drive purposelist`. If omitted, OneFS will automatically assign the drive purpose.

{--force | -f}

Formats the drive without asking for confirmation.

isi devices list

Defaults to `isi devices drive list`. Displays a list of data drives in a node.

 **NOTE:** You can display nodes that are available to join the cluster by running the command `isi devices node list`.

Syntax

```
isi devices list
  [--node-lnn <string>]
  [--sled <string>]
  [--override]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--node-lnn {all | <integer>}

Specifies the node number of the drives you would like to display. You may specify all nodes. If omitted, only the drives in the local node will be displayed.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

{ --override | -v }

Displays legacy bay numbers instead of grid values.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

isi devices node add

Joins an available node to the cluster.

Syntax

```
isi devices node add <serial-number>
  [--force]
```

Options

<serial-number>

Specifies the serial number of the node you want to add to the cluster.

{ --force | -f }

Adds the node to the cluster without asking for confirmation.

isi devices node list

Displays a list of nodes that are available to join the cluster.

Syntax

```
isi devices node list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ **--no-footer** | **-z** }

Displays table output without footers.

{ **--verbose** | **-v** }

Displays more detailed information.

isi devices node smartfail

Smartfails a node and removes it from the cluster.

Syntax

```
isi devices node smartfail
  [--node-lnn <integer>]
  [--force]
  [--verbose]
```

Options

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that you want to smartfail. If omitted, the local node will be smartfailed.

{ **--force** | **-f** }

Smartfails the drive without asking for confirmation.

{ **--verbose** | **-v** }

Displays more detailed information.

isi devices node stopfail

Discontinues the smartfail process on a node.

Syntax

```
isi devices node stopfail
  [--node-lnn <integer>]
  [--force]
  [--verbose]
```

Options

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that you want to discontinue smartfailing. If omitted, the local node will discontinue smartfailing.

{ **--force** | **-f** }

Discontinues smartfailing the drive without asking for confirmation.

{ **--verbose** | **-v** }

Displays more detailed information.

isi devices purpose

Defaults to `isi devices drive purpose`. Assigns a use case to a drive. For example, you can designate a drive for normal data storage operations, or you can designate the drive for L3 caching instead of storage.

Syntax

```
isi devices purpose {<bay> | --lnum <integer>}
  [--node-lnn <integer>]
  [--purpose <string>]
  [--force | -f]
  [--verbose | -v]
```

Options

{<bay> | --lnum <integer>}

Specifies the bay number or LNUM (logical drive number) of the drive to assign.

--purpose <string>

Specifies the purpose to assign to the drive. You can view a list of the possible drive purposes by running `isi devices drive purposelist`.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to assign. If omitted, the specified drive in the local node will be assigned.

{--force | -f}

Formats the drive without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices purposelist

Defaults to `isi devices drive purposelist`. Displays a list of possible use cases for drives. For example, you may be able to designate a drive for normal data storage operations, or you can designate the drive for L3 caching instead of storage.

Syntax

```
isi devices purposelist
  [--node-lnn <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that you want to view the purpose list for. If omitted, the purpose list of the local node will display.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

isi devices smartfail

Defaults to `isi devices drive smartfail`. Smartfails a drive so you can remove it from a node.

 **NOTE:** You can smartfail a node by running the command `isi devices node smartfail`.

Syntax

```
isi devices smartfail {<bay> | --lnum <integer> | --sled <string>}
    [--node-lnn <integer>]
    [{--force | -f}]
    [{--verbose | -v}]
```

Options

{<bay> | --lnum <integer>}

Specifies the bay number or LNUM (logical drive number) of the drive to smartfail.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to smartfail. If omitted, the specified drive in the local node will be smartfailed.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

{--force | -f}

Smartfails the drive without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices stopfail

Defaults to `isi devices drive stopfail`. Discontinues the smartfail process on a drive.

 **NOTE:** You can discontinue the smartfail process on a node by running the command `isi devices node stopfail`.

Syntax

```
isi devices stopfail {<bay> | --lnum <integer> | --sled <string>}
    [--node-lnn <integer>]
    [{--force | -f}]
    [{--verbose | -v}]
```

Options

{<bay> | --lnum <integer>}

Specifies the bay number or LNUM (logical drive number) of the drive to discontinue smartfailing.

--sled <string>

Specifies the prefix for the location of the sled of drives on which to operate. Must be one of **[A | B | C | D | E]**.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to discontinue smartfailing. If omitted, the specified drive in the local node will be discontinue smartfailing.

{--force | -f}

Discontinues smartfailing the drive without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices suspend

Defaults to `isi devices drive suspend`. Temporarily suspends all activities for a drive.

Syntax

```
isi devices suspend {<bay> | --lnum <integer> | --sled <string>}
  [--node-lnn <integer>]
  [{--force | -f}]
  [{--verbose | -v}]
```

Options

{<bay> | --lnum <integer> | --sled <string>}

Specifies the bay number or LNUM (logical drive number), or prefix for the location of the sled drive(s) to suspend. The `--sled<string>` must be one of **[A | B | C | D | E]**.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to suspend. If omitted, the specified drive in the local node will be suspended.

{--force | -f}

Smartfails the drive without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi devices view

Defaults to `isi devices drive view`. Displays information about a single drive.

Syntax

```
isi devices view {<bay> | --lnum <integer>}
  [--node-lnn <integer>]
```

Options

{<bay> | --lnum <integer>}

Specifies the bay number or LNUM (logical drive number) of the drive to view.

--node-lnn <integer>

Specifies the LNN (logical node number) of the node that contains the drive you want to view. If omitted, the drive in the local node will be displayed.

isi diagnostics gather settings modify

Modifies collection and upload settings for cluster log information.

Syntax

```
isi diagnostics gather settings modify
  [--upload {enable | disable}]
  [--esrs {enable | disable}]
  [--gather-mode {incremental | full}]
  [--http-upload {enable | disable}]
  [--http-upload-host <host>]
  [--http-upload-path <path>]
  [--http-upload-proxy <host>]
  [--http-upload-proxy-port <port>]
  [--ftp-upload {enable | disable}]
  [--ftp-upload-host <host>]
  [--ftp-upload-path <path>]
  [--ftp-upload-proxy <host>]
  [--ftp-upload-proxy-port <port>]
  [--ftp-upload-user <username>]
  [--ftp-upload-pass <password>]
  [--set-ftp-upload-pass]
  [--verbose]
```

Options

--upload {enable | disable}

Enables the upload of gathered logs

--esrs {enable | disable}

Specifies EMC Secure Remote Services (ESRS) for log uploads.

--gather-mode {incremental | full}

Specifies whether you will start an incremental or full gather of logs.

--http-upload {enable | disable}

Specifies HTTP for log uploads.

--http-upload-host <host>

Specifies the HTTP site for upload.

--http-upload-path <path>

Specifies the HTTP upload directory.

--http-upload-proxy <host>

Specifies an HTTP proxy server.

--http-upload-proxy-port <port>

Specifies the HTTP proxy server port.

--ftp-upload {enable | disable}

Specifies FTP for log uploads.

--ftp-upload-host <host>

Specifies the FTP site for upload.

--ftp-upload-path <path>

Specifies the FTP upload directory.

--ftp-upload-proxy <host>

Specifies an FTP proxy server.

--ftp-upload-proxy-port <port>

Specifies the FTP proxy server port.

- ftp-upload-user <username>**
Specifies the FTP site username. The default user is anonymous.
- ftp-upload-pass <password>**
Specifies the FTP site password.
- set-ftp-upload-pass <password>**
Sets the FTP site password.
- {--verbose | -v}**
Displays more detailed information.

isi diagnostics gather settings view

Displays settings for log gathering.

Syntax

```
isi diagnostics gather settings view
```

Options

This command has no options.

isi diagnostics gather start

Starts the process to collect and upload the most recent cluster log information.

Gathered cluster logs are saved under `/ifs/data/Isilon_Support/pkg`.

Syntax

```
isi diagnostics gather start
  [--upload {enable | disable}]
  [--esrs {enable | disable}]
  [--gather-mode {incremental | full}]
  [--http-upload {enable | disable}]
  [--http-upload-host <host>]
  [--http-upload-path <path>]
  [--http-upload-proxy <host>]
  [--http-upload-proxy-port <port>]
  [--ftp-upload {enable | disable}]
  [--ftp-upload-host <host>]
  [--ftp-upload-path <path>]
  [--ftp-upload-proxy <host>]
  [--ftp-upload-proxy-port <port>]
  [--ftp-upload-user <username>]
  [--ftp-upload-pass <password>]
  [--set-ftp-upload-pass]
```

Options

- upload {enable | disable}**
Enables the upload of gathered logs
- esrs {enable | disable}**
Specifies EMC Secure Remote Services (ESRS) for log uploads.
- gather-mode {incremental | full}**

Specifies whether you will start an incremental or full gather of logs.

--http-upload {enable | disable}

Specifies HTTP for log uploads.

--http-upload-host <host>

Specifies the HTTP site for upload.

--http-upload-path <path>

Specifies the HTTP upload directory.

--http-upload-proxy <host>

Specifies an HTTP proxy server.

--http-upload-proxy-port <port>

Specifies the HTTP proxy server port.

--ftp-upload {enable | disable}

Specifies FTP for log uploads.

--ftp-upload-host <host>

Specifies the FTP site for upload.

--ftp-upload-path <path>

Specifies the FTP upload directory.

--ftp-upload-proxy <host>

Specifies an FTP proxy server.

--ftp-upload-proxy-port <port>

Specifies the FTP proxy server port.

--ftp-upload-user <username>

Specifies the FTP site username. The default user is anonymous.

--ftp-upload-pass <password>

Specifies the FTP site password.

--set-ftp-upload-pass <password>

Sets the FTP site password.

isi diagnostics gather status

Displays the current status of a log gather operation.

Syntax

```
isi diagnostics gather status
```

Options

This command has no options.

isi diagnostics gather stop

Stops active log gather operations.

Syntax

```
isi diagnostics gather stop
```

Options

This command has no options.

isi diagnostics netlogger settings modify

Modifies collection settings for IP traffic information.

Syntax

```
isi diagnostics netlogger settings modify
  [--interfaces <interface>]
  [--count <integer>]
  [--duration <duration>]
  [--snaplength <bytes>]
  [--nodelist <LNN>]
  [--clients <IP>]
  [--ports <string>]
  [--protocols {ip | ip6 | arp | tcp | udp}]
  [--verbose]
```

Options

--interfaces <interface>

Specifies the network interface on which to capture traffic.

--count <integer>

Specifies the number of capture files that you will keep after the capture finishes. The default value is three files.

--duration <duration>

Specifies how long you will capture IP traffic for each capture file, in the format <integer>{Y|M|W|D|H|m|s}

--snaplength <bytes>

The snap length for the capture. Default is 320 bytes. Valid range for this value is 64-9100.

--nodelist <nodes>

Specifies nodes to report statistics on. Specify nodes by Logical Node Number (LNN). Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`.

--clients <clients>

Specifies client IPs to report statistics on. Multiple IP addresses can be specified in a comma-separated list. The default value is `all`.

--ports <port>

Specifies TCP or UDP ports to report statistics on. Multiple ports can be specified in a comma-separated list. The default value is `all`.

--protocols {ip | ip6 | arp | tcp | udp}

Specifies a protocol to report statistics on.

{--verbose | -v}

Displays more detailed information.

isi diagnostics netlogger settings view

Displays settings for the capture of IP traffic logs.

Syntax

```
isi diagnostics netlogger settings view
```

Options

This command has no options.

isi diagnostics netlogger start

Starts the process to collect and upload the most recent IP traffic log information.

Gathered cluster logs are saved under `/ifs/data/Isilon_Support/pkg`.

Syntax

```
isi diagnostics netlogger start
[--interfaces <interface>]
[--count <integer>]
[--duration <duration>]
[--snaplength <bytes>]
[--nodelist <LNN>]
[--clients <IP>]
[--ports <string>]
[--protocols {ip | ip6 | arp | tcp | udp}]
```

Options

--interfaces <interface>

Specifies the network interface on which to capture traffic.

--count <integer>

Specifies the number of capture files that you will keep after the capture finishes. The default value is three files.

--duration <duration>

Specifies how long you will capture IP traffic for each capture file, in the format `<integer>{Y|M|W|D|H|m|s}`

--snaplength <bytes>

The snap length for the capture. Default is 320 bytes. Valid range for this value is 64-9100.

--nodelist <nodes>

Specifies nodes to report statistics on. Specify nodes by Logical Node Number (LNN). Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`.

--clients <clients>

Specifies client IPs to report statistics on. Multiple IP addresses can be specified in a comma-separated list. The default value is `all`.

--ports <port>

Specifies TCP or UDP ports to report statistics on. Multiple ports can be specified in a comma-separated list. The default value is `all`.

--protocols {ip | ip6 | arp | tcp | udp}

Specifies a protocol to report statistics on.

isi diagnostics netlogger status

Displays the current status of an IP traffic capture operation.

Syntax

```
isi diagnostics netlogger status
```

Options

This command has no options.

isi diagnostics netlogger stop

Stops active IP traffic capture operations.

Syntax

```
isi diagnostics netlogger stop
```

Options

This command has no options.

isi email settings modify

Modify email settings for the cluster.

Syntax

```
isi email settings modify
  [--mail-relay <string>]
  [--smtp-port <integer>]
  [--mail-sender <string>]
  [--mail-subject <string>]
  [--use-smtp-auth {yes | no}]
  [--smtp-auth-username <string>]
  [--use-encryption {yes | no}]
  [--batch-mode {all | severity | category | none}]
  [--user-template <string>]
  [--clear-user-template]
  [--smtp-auth-passwd <string>]
  [--clear-smtp-auth-passwd]
  [--set-smtp-auth-passwd]
  [--verbose]
```

Options

--mail-relay <string>

Sets the SMTP relay address.

--smtp-port <integer>

Sets the SMTP port. The default is 25.

- mail-sender <string>**
Sets the originator email address.
- mail-subject <string>**
Set the prefix string for the email subject.
- use-smtp-auth {yes | no}**
Use SMTP authentication.
- {--smtp-auth-username | -u} <string>**
Sets the SMTP user name.
- use-encryption {yes | no}**
Use encryption (TLS) for SMTP authentication.
- batch-mode {all | severity | category | none}**
Sets the method that notifications are batched together to be sent by email.
- user-template <string>**
Specifies the path to access a custom email template.
- clear-user-template**
Clears the path specified to access a custom email template.
- {--smtp-auth-passwd | -p} <string>**
Sets the SMTP authentication password.
- clear-smtp-auth-passwd**
Clears the specified SMTP authentication password.
- set-smtp-auth-passwd**
Specifies --smtp-auth-passwd interactively.
- {--verbose | -v}**
Displays more detailed information.

isi email settings view

View cluster email settings.

Syntax

```
isi email settings view
```

Example

To view the currently-configured email settings, run the following command:

```
isi email settings view
```

The system displays output similar to the following example:

```
Mail Relay: -
SMTP Port: 25
Mail Sender: -
Mail Subject: -
Use SMTP Auth: No
SMTP Auth Username: -
Use Encryption: No
Batch Mode: none
User Template: -
SMTP Auth Password Set: False
```

isi esrs modify

Enable/disable secure remote services (ESRS) and modify configuration settings. You must have a valid Dell EMC support contract to use ESRS.

Syntax

```
isi esrs modify
  [--enabled <boolean>]
  [--username <string>]
  [--password <string>]
  [--force]
  [--primary-esrs-gateway <string>]
  [--secondary-esrs-gateway <string>]
  [--alert-on-disconnect <boolean>]
  [--gateway-access-pools <string> (--clear-gateway-access-pools | --add-gateway-access-pools
  <string> | --remove-gateway-access-pools <string>)]
  [--connectivity-check-period <integer>]
  [--ui-reporting-period <integer>]
  [--download-enabled <boolean>]
  [--download-timeout-period <integer>]
  [--download-error-retries <integer>]
  [--download-chunk-size <integer>]
  [--download-filesystem-limit <integer>
```

Options

- enabled (yes | no)** Enables or disables ESRS. The default is no.
- username <string>** The Dell EMC user name for technical assistance.
- password <string>** The password associated with the Dell EMC user name.
- force** Forces ESRS de-provisioning when communication with ESRS or the ESRS gateway are not available.
- primary-esrs-gateway <string>** The host name or IP address of the primary ESRS gateway.
- secondary-esrs-gateway <string>** The host name or IP address of the secondary ESRS gateway.
- alert-on-disconnect (yes | no)** Generate a CELOG alert if on-cluster ESRS is disconnected from the ESRS gateway.
- gateway-access-pools <string> (--clear-gateway-access-pools | --add-gateway-access-pools <string> | --remove-gateway-access-pools <string>)** A list of the pools available to access ESRS gateways.
 - clear-gateway-access-pools** Clear the list of pool IDs.
 - add-gateway-access-pools <string>** Add items to the list of pool identifiers. Use this option multiple times for additional pools.
 - remove-gateway-access-pools <string>** Remove items from the list of pool identifiers. Use this option multiple times for additional pools.

<code>--connectivity-check-period <integer></code>	Specify how often to re-check for ESRS gateway connectivity, in seconds.
<code>--ui-reporting-period <integer></code>	Specify how often to display license usage intelligence reports, in seconds.
<code>--download-enabled (yes no)</code>	Enables or disables the ESRS file download feature.
<code>--download-timeout-period <integer></code>	Specify the timeout period for ESRS file downloads, in seconds. The default is 50 seconds per chunk.
<code>--download-error-retries <integer></code>	Specify the number of times to re-try after experiencing an error or timeouts during a download. The default is three requests per chunk.
<code>--download-chunk-size <integer></code>	Specify the maximum number of bytes per chunk for a download request. The default is one million (1000000).
<code>--download-filesystem-limit <integer></code>	Set a threshold after which downloads are disallowed, based on a percentage of file system resources in use. The default threshold is 80 (80% of total file system resources).

isi esrs view

View secure remote services (ESRS) configuration and gateway settings.

Syntax

```
isi esrs view
```

Options

None

OneFS displays ESRS configuration and gateway setting information for the cluster.

isi esrs dataitems list

List configuration information for secure remote services (ESRS) data items.

Syntax

```
isi esrs dataitems list
  [--limit <integer>]
  [--format (table | json | csv | list)]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- {--limit | -l}** *<integer>* The number of ESRS data items to display.
- format (table | json | csv | list)** Displays ESRS data items in table, JSON, CSV, or list format.
- {--no-header | -a}** Do not display headers in table or CSV formats.
- {--no-footer | -z}** Do not display table summary footer information.
- {--verbose | -v}** Display more detailed information.

OneFS displays a list of ESRS data items, whether items are enabled, the frequency items are generated, and most recent timestamp.

isi esrs dataitems modify

Modify secure remote services (ESRS) data item settings and configure individual items.

Syntax

```
isi esrs dataitems modify
  [--dataitems-enabled <boolean>]
  [--name <string>]
  [--enabled <boolean>]
  [--frequency <duration>]
  [--timestamp-delete]
```

Options

- dataitems-enabled <boolean>** Enables or disables ESRS data items. The default is enabled.
- name <string>** The name of an ESRS data item.
- enabled <boolean>** Enables or disables an ESRS data item.
- frequency <duration>** Establishes how often an ESRS data item is generated.
- timestamp-delete** Resets the last run timestamp for a selected ESRS data item to zero.

isi esrs download list

List the files available in the secure remote services (ESRS) file locker.

Syntax

```
isi esrs download list
  [--limit <integer>]
  [--format (table | json | csv | list)]
  [--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

- {--limit | -l}** *<integer>* The number of ESRS downloads to display.
- format (table | json | csv | list)** Displays ESRS data items in table, JSON, CSV, or list format.
- {--no-header | -a}** Do not display headers in table or CSV formats.
- {--no-footer | -z}** Do not display table summary footer information.
- {--verbose | -v}** Display more detailed information.

OneFS displays a list of ESRS file locker downloads.

isi esrs download start

Schedule a OneFS job engine job to download a file from the secure remote services (ESRS) file locker.

Syntax

```
isi esrs download start <file> <path>
```

Options

- <file>** The file name as displayed in the ESRS file locker.
- <path>** The path in which the file will be placed after successful download.

isi esrs download view

View details about a specific secure remote services (ESRS) downloadable file.

Syntax

```
isi esrs download view <file>
```

Options

- <file>** The name of the file as displayed in the ESRS file locker.

OneFS displays information about the specified ESRS downloadable file.

isi esrs telemetry modify

Enable or disable secure remote services (ESRS) telemetry data uploads.

Syntax

```
isi esrs telemetry modify
  [--enabled (yes | no)]
```

Options

--enabled (yes | no) Enables or disables the ESRS telemetry upload feature. This option is enabled by default.

isi esrs telemetry view

View secure remote services (ESRS) telemetry settings.

Syntax

```
isi esrs telemetry view
```

Options

None

OneFS displays whether ESRS telemetry is enabled and the number of pending uploads.

isi event alerts create

Creates a new alert.

Syntax

```
isi event alerts create <name> <condition> channel
  [--category <string>]
  [--eventgroup <string>]
  [--severity {emergency | critical | warning | information}]
  [--limit <integer>]
  [--interval <duration>]
  [--transient <duration>]
  [--verbose]
```

Options

<name>

Specifies the alert name.

<condition>

Specifies the condition under which alert is sent.

Condition values are case sensitive. The following values are valid:

NEW	Reports on event group occurrences that have never been reported on before.
NEW_EVENTS	Reports on event group occurrences that are new since the event group was last reported on.
ONGOING	Provides periodic reports on event group occurrences that have not been resolved.
SEVERITY_INCREASE	Reports on event group occurrences whose severity has increased since the event group was last reported on.
SEVERITY_DECREASE	Reports on event group occurrences whose severity has decreased since the event group was last reported on.
RESOLVED	Reports on event group occurrences that have been resolved since the event group was last reported on.

<channel> ...

Specifies the channel over which to deliver the alert. Use the `--channel` parameter to specify additional channels.

--category <string>...

Specifies the name of one or more event group categories to alert on.

--eventgroup <string>...

Specifies the name of one or more event groups to alert on.

--severity {emergency | critical | warning | information}

Specifies the event severity that the alert will report on. Severity values are case sensitive. Repeat `--severity` to make the alert report on additional severity levels.

--limit <integer>

Sets the maximum number of alerts that can be sent. Applies only to the `NEW_EVENTS` alert condition.

--interval <integer> <time>

Sets the time period between reports for ongoing alerts. Applies only to the `ONGOING` alert condition.

The following `<time>` values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

--transient <integer> <time>

Sets a minimum time that an event group occurrence must exist before it is reported on. Any occurrence lasting less than the time period is considered transient and will not be reported.

The following `<time>` values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

{--verbose | -v}

Displays more detailed information.

isi event alerts delete

Deletes an alert.

Syntax

```
isi event alerts delete <name>
  [--force | -f]
  [--verbose | -v]
```

Options

<name>

Specifies the name of the alert you want to delete.

{--force | -f}

Deletes the alert without asking for confirmation.

{--verbose | -v}

Displays more detailed information.

isi event alerts list

Displays a list of alerts.

Syntax

```
isi event alerts list
  [--channel | -c <string>]
  [--limit | -l <integer>]
  [--sort {name | eventgroup | category | channel | condition | limit |
  interval | transient}]
  [--descending | -d]
  [--format {table | json | csv | list}]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose | -v]
```

Options

{--channel | -c} <string>...

Displays alerts for the specified channel only.

{--limit | -l} <integer>

Sets the maximum number of alerts to display.

--sort {name | eventgroup | category | channel | condition | limit | interval | transient}

Specifies the field to sort items by.

{--descending | -d}

Sorts the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi event alerts modify

Modifies an alert.

Syntax

```
isi event alerts modify <name>
  [--eventgroup <string>]
  [--clear-eventgroup]
  [--add-eventgroup <string>]
  [--remove-eventgroup <string>]
  [--category <string>]
  [--clear-category]
  [--add-category <string>]
  [--remove-category <string>]
  [--channel | -c <string>]
  [--clear-channel]
  [--add-channel <string>]
  [--remove-channel <string>]
  [--severity {emergency | critical | warning | information}]
  [--clear-severity]
  [--add-severity {emergency | critical | warning | information}]
  [--remove-severity {emergency | critical | warning | information}]
  [--condition {NEW | NEW_EVENTS | ONGOING | SEVERITY_INCREASE
  | SEVERITY_DECREASE | RESOLVED}]
  [--limit <integer>]
  [--interval <integer> <time>]
  [--transient <integer> <time>]
  [--verbose]
```

Options

<name>

Specifies the name of the alert you want to modify.

--eventgroup <string>...

Specifies the name of one or more event groups to alert on.

--clear-eventgroup

Clears the value for an event group to alert on.

--add-eventgroup <string>...

Adds the name of one or more event groups to alert on.

--remove-eventgroup <string>...

Removes the name of one or more event groups to alert on.

--category <string>...

Specifies the name of one or more event group categories to alert on.

--clear-category

Clears the value for an event group category to alert on.

--add-category <string>...

Adds the name of one or more event group categories to alert on.

--remove-category <string>...

Removes the name of one or more event group categories to alert on.

{--channel | -c} <string>...

Specifies the name of one or more channels to deliver the alert over.

--clear-channel
Clears the value for a channel to deliver the alert over.

--add-channel <string>...
Adds the name of one or more channels to deliver the alert over.

--remove-channel <string>...
Removes the name of one or more channels to deliver the alert over.

--severity {emergency | critical | warning | information}
Specifies the event severity that the alert will report on. Severity values are case sensitive. Repeat `--severity` to make the alert report on additional severity levels.

--clear-severity
Clears all severity filters for an alert.

--add-severity {emergency | critical | warning | information}
Adds another severity value for an alert to report on. Repeat `--add-severity` to make the alert report on additional severity levels.

--remove-severity {emergency | critical | warning | information}
Removes a severity value that an alert is reporting on. Repeat `--remove-severity` to remove an additional severity level that an alert is reporting on.

--condition {NEW | NEW_EVENTS | ONGOING | SEVERITY_INCREASE | SEVERITY_DECREASE | RESOLVED}
Specifies the condition under which alert is sent.
Condition values are case sensitive. The following values are valid:

NEW	Reports on event group occurrences that have never been reported on before.
NEW_EVENTS	Reports on event group occurrences that are new since the event group was last reported on.
ONGOING	Provides periodic reports on event group occurrences that have not been resolved.
SEVERITY_INCREASE	Reports on event group occurrences whose severity has increased since the event group was last reported on.
SEVERITY_DECREASE	Reports on event group occurrences whose severity has decreased since the event group was last reported on.
RESOLVED	Reports on event group occurrences that have been resolved since the event group was last reported on.

--limit <integer>
Sets the maximum number of alerts that can be sent. Applies only to the `NEW_EVENTS` alert condition.

--interval <integer> <time>
Sets the time period between reports for ongoing alerts. Applies only to the `ONGOING` alert condition.
The following `<time>` values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

--transient <integer> <time>
Sets a minimum time that an event group occurrence must exist before it is reported on. Any occurrence lasting less than the time period is considered transient and will not be reported.
The following `<time>` values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

{--verbose | -v}

Displays more detailed information.

isi event alerts view

Displays the details of an alert.

Syntax

```
isi event alerts view <name>
```

Options

<name>

Specifies the name of the alert condition.

isi event channels create

Creates a new channel.

Syntax

```
isi event channels create <name> <type>
  [--enabled {true | false}]
  [--allowed-nodes <integer>]
  [--excluded-nodes <integer>]
  [--address <string>]
  [--send-as <string>]
  [--subject <string>]
  [--smtp-host <string>]
  [--smtp-port <integer>]
  [--smtp-use-auth <boolean>]
  [--smtp-username <string>]
  [--smtp-password <string>]
  [--smtp-security {STARTTLS | NONE}]
  [--batch {NONE | ALL | CATEGORY | SEVERITY}]
  [--batch-period <integer> <time>]
  [--host <string>]
  [--community <string>]
  [--use-snmp-trap {true | false}]
  [--snmp-use-v3 {true | false}]
  [--snmp-security-name <string>]
  [--snmp-security-level {noAuthNoPriv | authNoPriv | authPriv}]
  [--snmp-auth-protocol {MD5 | SHA}]
  [--snmp-auth-password <string>]
  [--snmp-priv-protocol {AES | DES}]
  [--snmp-priv-password <string>]
  [--snmp-engine-id <string>]
  [--verbose]
```

Options

<name>

Specifies the channel name.

<type>

Specifies the mechanism by which alerts are sent.

Type values are case sensitive. The following values are valid:

smtp	Alerts are sent as emails through an SMTP server.
snmp	Alerts are sent through SNMP.
connectemc	Alerts are sent through ConnectEMC.

--enabled {true | false}

Specifies whether the channel is enabled.

--allowed-nodes <integer>...

Specifies one or more nodes that are allowed to send alerts through the channel. If you do not specify any allowed nodes, all nodes in the cluster will be allowed to send alerts. The value of <integer> is the node number you want to allow.

--excluded-nodes <integer>...

Specifies one or more nodes that are not allowed to send alerts through the channel. The value of <integer> is the node number you want to exclude.

--address <string>...

For SMTP channels only. Specifies one or more email addresses you want to receive alerts on this channel. The value of <string> is an email address.

--send-as <string>

For SMTP channels only. Specifies the email address you want to send alerts from on this channel. The value of <string> is an email address.

--subject <string>

For SMTP channels only. Specifies the subject line for emails sent on this channel.

--smtp-host <string>

For SMTP channels only. Specifies the SMTP relay host.

--smtp-port <integer>

For SMTP channels only. Specifies the SMTP relay port.

--smtp-use-auth {true | false}

For SMTP channels only. Enables or disables SMTP authentication.

--smtp-username <string>

For SMTP channels only. Specifies the username for SMTP authentication.

--smtp-password <string>

For SMTP channels only. Specifies the password for SMTP authentication.

--smtp-security {STARTTLS | NONE}

For SMTP channels only. Enables or disables SMTP encryption.

--batch {NONE | ALL | CATEGORY | SEVERITY}

For SMTP channels only. Specifies how SMTP alerts will be batched.

--batch-period <integer> <time>

For SMTP channels only. Specifies the interval between batched alerts.

The following <time> values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks

D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

- host <string>**
For SNMP channels only. Specifies the host name or address
- community <string>**
For SNMP channels only. Specifies the community string.
- use-snmp-trap {true | false}**
Specifies using `snmptrap` instead of `snmpinform` for this channel.
- snmp-use-v3 {true | false}**
Specifies using version 3 traps for this channel.
- snmp-security-name <string>**
Specifies the name for the version 3 user.
- snmp-security-level {noAuthNoPriv | authNoPriv | authPriv}**
Specifies the SNMPv3 security level for this channel..
- snmp-auth-protocol {MD5 | SHA}**
Specifies the SNMPv3 authentication algorithm to use for this channel.
- snmp-auth-password <string>**
Specifies the SNMPv3 authentication password.
- snmp-priv-protocol {AES | DES}**
Specifies the SNMPv3 encryption algorithm to use for this channel.
- snmp-priv-password <string>**
Specifies the SNMPv3 encryption password.
- snmp-engine-id <string>**
Specifies the SNMPv3 engine ID of the receiving trap daemon.
- {--verbose | -v}**
Displays more detailed information.

isi event channels delete

Deletes a channel.

Syntax

```
isi event channels delete <name>
    [--force]
    [--verbose]
```

Options

- <name>**
Specifies the name of the channel you want to delete.
- {--force | -f}**
Deletes the channel without asking for confirmation.
- {--verbose | -v}**
Displays more detailed information.

isi event channels list

Displays a list of channels.

Syntax

```
isi event channels list
  [--limit <integer>]
  [--sort {id | name | type | enabled | allowed_nodes | excluded_nodes |
  address | send_as | subject | smtp_host | smtp_port |
  smtp_use_auth | smtp_username | smtp_password | smtp_security |
  batch | batch_period | host | community}
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Sets the maximum number of channels to display.

--sort {id | name | type | enabled | allowed_nodes | excluded_nodes | address | send_as | subject | smtp_host | smtp_port | smtp_use_auth | smtp_username | smtp_password | smtp_security | batch | batch_period | host | community}

Specifies the field to sort items by.

{--descending | -d}

Sorts the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi event channels modify

Modifies a channel.

Syntax

```
isi event channels <name>
  [--type {smtp | snmp | connectemc}]
  [--enabled {true | false}]
  [--allowed-nodes <integer>]
  [--clear-allowed-nodes]
  [--add-allowed-nodes <integer>]
  [--remove-allowed-nodes <integer>]
  [--excluded-nodes <integer>]
  [--clear-excluded-nodes]
  [--add-excluded-nodes <integer>]
  [--remove-excluded-nodes <integer>]
  [--address <string>]
```

```

[--clear-address]
[--add-address <string>]
[--remove-address <string>]
[--send-as <string>]
[--subject <string>]
[--smtp-host <string>]
[--smtp-port <integer>]
[--smtp-use-auth <boolean>]
[--smtp-username <string>]
[--smtp-password <string>]
[--smtp-security {STARTTLS | NONE}]
[--batch {NONE | ALL | CATEGORY | SEVERITY}]
[--batch-period <integer> <time>]
[--host <string>]
[--community <string>]
[--verbose]

```

Options

<name>

Specifies the name of the channel you want to modify.

<type>

Specifies the mechanism by which alerts are sent.

Type values are case sensitive. The following values are valid:

smtp	Alerts are sent as emails through an SMTP server.
snmp	Alerts are sent through SNMP.
connectemc	Alerts are sent through ConnectEMC.

--enabled {true | false}

Specifies whether the channel is enabled.

--allowed-nodes <integer>...

Specifies one or more nodes that are allowed to send alerts through the channel. If you do not specify any allowed nodes, all nodes in the cluster will be allowed to send alerts. The value of <integer> is the node number you want to allow.

--clear-allowed-nodes

Clears all values for allowed nodes.

--add-allowed-nodes <integer>...

Adds one or more nodes to the allowed nodes list. The value of <integer> is the node number you want to allow.

--remove-allowed-nodes <integer>...

Removes one or more nodes from the allowed nodes list. The value of <integer> is the node number you want to remove.

--excluded-nodes <integer>...

Specifies one or more nodes that are not allowed to send alerts through the channel. The value of <integer> is the node number you want to exclude.

--clear-excluded-nodes

Clears all values for excluded nodes.

--add-excluded-nodes <integer>...

Adds one or more nodes to the excluded nodes list. The value of <integer> is the node number you want to exclude.

--remove-excluded-nodes <integer>...

Removes one or more nodes from the excluded nodes list. The value of <integer> is the node number you want to remove.

--address <string>...

For SMTP channels only. Specifies one or more email addresses you want to receive alerts on this channel. The value of `<string>` is an email address.

--clear-address

For SMTP channels only. Clears all values for email addresses.

--add-address <string>...

For SMTP channels only. Specifies one or more email addresses you want to add to the alert distribution list for this channel. The value of `<string>` is an email address.

--remove-address <string>...

For SMTP channels only. Specifies one or more email addresses you want to remove from the alert distribution list for this channel. The value of `<string>` is an email address.

--send-as <string>

For SMTP channels only. Specifies the email address you want to send alerts from on this channel. The value of `<string>` is an email address.

--subject <string>

For SMTP channels only. Specifies the subject line for emails sent on this channel.

--smtp-host <string>

For SMTP channels only. Specifies the SMTP relay host.

--smtp-port <integer>

For SMTP channels only. Specifies the SMTP relay port.

--smtp-use-auth {true | false}

For SMTP channels only. Enables or disables SMTP authentication.

--smtp-username <string>

For SMTP channels only. Specifies the username for SMTP authentication.

--smtp-password <string>

For SMTP channels only. Specifies the password for SMTP authentication.

--smtp-security {STARTTLS | NONE}

For SMTP channels only. Enables or disables SMTP encryption.

--batch {NONE | ALL | CATEGORY | SEVERITY}

For SMTP channels only. Specifies how SMTP alerts will be batched.

--batch-period <integer> <time>

For SMTP channels only. Specifies the interval between batched alerts.

The following `<time>` values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

--host <string>

For SNMP channels only. Specifies the host name or address

--community <string>

For SNMP channels only. Specifies the community string.

{--verbose | -v}

Displays more detailed information.

isi event channels view

Displays the details of a channel.

Syntax

```
isi event channels view <name>
```

Options

<name>

Specifies the name of the channel you want to view.

isi event events list

Displays all events.

Syntax

```
isi event events list
  [--eventgroup-id <name>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--eventgroup-id <name>

Displays events that are included in the specified event group.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

isi event events view

Displays the details of an event.

Syntax

```
isi event events view <id>
```

Options

<id>

Specifies the instance ID of the event you want to view.

isi event groups bulk

Changes the status of all event groups.

Syntax

```
isi event groups bulk
  [--ignore {true | false}]
  [--resolved {true | false}]
  [--verbose]
```

Options

--ignore {true | false}

Specifies whether all event groups have a status of ignored.

--resolved {true | false}

Specifies whether all event groups have a status of resolved.

After you resolve an event group, you cannot reverse that action. Any new events that would have been added to the resolved event group will be added to a new event group.

{--verbose | -v}

Displays more detailed information.

isi event groups list

Displays a list of all event groups.

Syntax

```
isi event groups list
  [--begin <timestamp>]
  [--end <timestamp>]
  [--resolved {true | false}]
  [--ignore {true | false}]
  [--events <integer>]
  [--cause <string>]
  [--limit <integer>]
  [--sort {id | started | causes_long | last_event | ignore |
  ignore_time | resolved | ended | events | severity}
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--begin | -b} <timestamp>

Filters the list to only show event groups that were created after the specified date and time.

Specify <timestamp> in the following format:

{--end | -e} <timestamp>

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Filters the list to only show event groups that were created before the specified date and time.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

--resolved {true | false}

Specifies whether the list will show only event groups that are resolved, or not resolved.

--ignore {true | false}

Specifies whether the list will show only event groups that are ignored, or not ignored.

--events <integer>

Filters the list to only show event groups with the specified number of events recorded against the event group.

--cause <string>

Filters the list to only show event groups with the specified cause.

{--limit | -l} <integer>

Sets the maximum number of event groups to display.

--sort {id | started | causes_long | last_event | ignore | ignore_time | resolved | ended | events | severity}

Specifies the field to sort items by.

{--descending | -d}

Sorts the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a}

Displays table and CSV output without headers.

{ --no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi event groups modify

Changes the status of an event group.

Syntax

```
isi event <id>
  [--ignore {true | false}]
  [--resolved {true | false}]
  [--verbose]
```

Options

<id>

Specifies the ID number of the event group you want to modify.

--ignore {true | false}

Specifies whether the event group has a status of ignored.

--resolved {true | false}

Specifies whether the event group has a status of resolved.

After you resolve an event group, you cannot reverse that action. Any new events that would have been added to the resolved event group will be added to a new event group.

{--verbose | -v}

Displays more detailed information.

isi event groups view

View the details of an event group.

Syntax

```
isi event groups view <id>
```

Options

<id>

Specifies the ID number of the event group you want to view.

isi event settings modify

Configures event storage settings.

Syntax

```
isi event settings modify
  [--retention-days <integer>]
  [--storage-limit <integer>]
  [--maintenance-start <timestamp>]
  [--clear-maintenance-start]
  [--maintenance-duration <duration>]
  [--heartbeat-interval <string>]
  [--verbose]
```

Options

{--retention-days | -r} <integer>

Retention of resolved event group data in days.

{--storage-limit | -s} <integer>

Sets the amount of memory that event data can occupy on your cluster. You can set this limit to be between 1 and 100 megabytes of memory. For smaller clusters, the minimum amount of memory that will be set aside is 1 gigabyte.

--maintenance-start <timestamp>

Sets the start date and time of a maintenance window.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

--clear-maintenance-start

Clears the value for the start date and time of a maintenance window.

--maintenance-duration <integer> <time>

Sets the duration of a maintenance window.

The following *<time>* values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

--heartbeat-interval <string>

Sets the interval between heartbeat events.

The following *<time>* values are valid:

- daily
- weekly
- monthly

{--verbose | -v}

Displays more detailed information.

isi event settings view

Displays event storage settings.

Syntax

```
isi event settings view
```

isi event test create

Creates a test alert.

Syntax

```
isi event test create <message>  
    [--verbose]
```

Options

<message>

Specifies the message text of the test alert.

{--verbose | -v}

Displays more detailed information.

isi event types list

View a list of OneFS event categories.

Syntax

```
isi event types list
  [--category <string>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--category <string>

Enter a OneFS event category. If you do not enter a category, OneFS displays all categories.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

isi fc settings list

Lists Fibre Channel port settings.

Syntax

```
isi fc settings list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

--format {table | json | csv | list}

Displays Fibre Channel port settings in table, JSON, CSV, or list format.

{ --no-header | -a }

Does not display headers in table or CSV formats.

{ --no-footer | -z }

Does not display table summary footer information.

Examples

The following command displays all ports on node 5:

```
isi fc settings list
```

The system displays output similar to the following example:

Port	WWNN	WWPN	State	Topology	Rate
5:1	2000001b3214ccc3	2100001b3214ccc3	enabled	auto	auto
5:2	2000001b3234ccc3	2101001b3234ccc3	enabled	auto	auto
5:3	2000001b3254ccc3	2100001b3254ccc3	enabled	auto	auto
5:4	2000001b3234ccc3	2103001b3274ccc3	enabled	auto	auto

isi fc settings modify

Modifies Fibre Channel settings for a specific port.

Syntax

```
isi fc settings modify --port <port>
  [--wwnn <string>]
  [--wvpn <string>]
  [--state {enable | disable}]
  [--topology {auto | loop | ptp}]
  [--rate {auto | 1 | 2 | 4 | 8}]
```

Options

--wwnn <string>

Specifies the world-wide node name (WWNN) of the port as a string of 16 hexadecimal numerals.

--wvpn <string>

Specifies the world-wide port name (WWPN) of the port as a string of 16 hexadecimal numerals.

--state {enable | disable}

Specifies whether the port is enabled or disabled.

--topology {auto | loop | ptp}

Specifies the type of Fibre Channel topology that the port expects. The following settings are valid:

auto	Causes the port to detect the topology automatically. This is the recommended setting. Specify this setting if you are using a fabric topology.
loop	Causes the port to expect an arbitrated loop topology, with multiple backup devices connected to a single port in a circular formation.
ptp	Causes the port to expect a point-to-point topology, with one backup device or Fibre Channel switch directly connected to the port.

--rate {auto | 1 | 2 | 4 | 8}

Specifies the rate that OneFS will attempt to send data through the port. The following rates are valid:

auto	OneFS automatically negotiates with the DMA to determine the rate. This is the recommended setting.
1	Attempts to send data through the port at a speed of 1 Gb per second.
2	Attempts to send data through the port at a speed of 2 Gb per second.
4	Attempts to send data through the port at a speed of 4 Gb per second.
8	Attempts to send data through the port at a speed of 8 Gb per second.

isi fc settings view

Displays settings for a specific Fibre Channel port.

Syntax

```
isi fc settings view port
  [--format {list | json}]
```

Options

<port>

A Fibre Channel port ID in the format <lnn>.<fc port>.

--format {list | json}

Displays the Fibre Channel port settings in list or JSON format.

isi file-filter settings modify

Modifies file filtering settings in an access zone.

Syntax

```
isi file-filter settings modify
  [--file-filtering-enabled {yes | no}]
  [--revert-file-filtering-enabled]
  [--file-filter-extensions <string>...]
  [--clear-file-filter-extensions]
  [--add-file-filter-extensions <string>]
  [--remove-file-filter-extensions <string>]
  [--revert-file-filter-extensions]
  [--file-filter-type {allow | deny}]
  [--revert-file-filter-type]
  [--zone <string>]
  [--verbose]
```

Options

--file-filtering-enabled {yes | no}

Enables or disables file filtering in the access zone. File filtering is disabled by default.

--revert-file-filtering-enabled

Sets the value of `--file-filtering-enabled` to the system default value.

--file-filter-extensions <string>...

Specifies a list of file types by their extensions. Each extension should start with a "." such as `.txt`. You can specify multiple extensions in a comma-separated list or you run `--file-filter-extensions` for each extension.

--clear-file-filter-extensions

Deletes the entire list of file filter extensions.

--add-file-filter-extensions <string>

Adds one or more file filter extensions to the list. Each extension should start with a "." such as `.txt`. You can specify multiple extensions in a comma-separated list or you run `--add-file-filter-extensions` for each extension.

--remove-file-filter-extensions <string>

Removes one or more file filter extensions from the list. Each extension should start with a "." such as `.txt`. You can specify multiple extensions in a comma-separated list or you run `--remove-file-filter-extensions` for each extension.

--revert-file-filter-extensions

Sets the value of `--file-filter-extensions` to the system default value.

--file-filter-type {allow | deny}

Specifies whether the file types in the extensions list will be allowed or denied write access. The default filter type is `deny`.

--revert-file-filter-type

Sets the value of `--revert-file-filter-type` to the system default value.

--zone <string>

Specifies the access zone to which the settings apply. If you do not specify a zone, the settings are applied to the System zone.

{--verbose | -v}

Displays more detailed information.

isi file-filter settings view

Displays file filtering settings for an access zone.

Syntax

```
isi file-filter settings view
  [--zone <string>]
  [--format {list | json}]
```

Options

--zone

Specifies the name of the access zone. If you do not specify an access zone, the system will display the file filtering settings of the System zone.

--format {list | json}

Specifies whether to display the output as a list (default) or in JavaScript Object Notation (JSON).

isi filepool apply

Applies all file pool policies to the specified file or directory path. If no policy matches the file or directory path, OneFS applies the default file pool policy.

Syntax

```
isi filepool apply <path>
  [--path] <path>
  [--dont-restripe]
  [--nop]
  [--stats]
  [--quiet]
  [--recurse]
  [--verbose]
```

Options

--path<path>

Specifies the path to the file to be processed. This parameter is required.

--dont-restripe

Changes the per-file policies without restriping the file.

--nop

Calculates the specified settings without actually applying them. This option is best used with `--verbose` or `--stats`.

--stats

Displays statistics on the files processed.

--quiet

Suppresses warning messages.

--recurse

Specifies recursion through directories.

--verbose

Displays the configuration settings to be applied. We recommend using verbose mode. Otherwise the command would not display any screen output except for error messages.

Examples

These examples show the results of running `isi filepool apply` in verbose mode. In the examples, the output shows the results of comparing the path specified with each file pool policy. The `recurse` option is set so that all files in the `/ifs/data/projects` path are matched against all file pool policies. The first policy listed is always the system default policy. In this example, the second match is to the file pool policy **Technical Data**.

```
isi filepool apply --path=/ifs/data/projects --verbose --recurse
```

```
Processing file /ifs/data/projects
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0
New File Attributes
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0

{'default' :
  {'Policy Number': -2,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':1, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 1,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
'system':
```

```
'Technical Data':
  {'Policy Number': 0,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':0, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 0,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
```

This example shows the result of using the `--nop` option to calculate the results that would be produced by applying the file pool policy.

```
isi filepool apply --path=/ifs/data/projects --nop --verbose
```

```
Processing file /ifs/data/projects
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0
New File Attributes
Protection Level is DiskPool minimum
Layout policy is concurrent access
coalescer_enabled is true
data_disk_pool_policy_id is any pool group ID
data SSD strategy is metadata
snapshot_disk_pool_policy_id is any pool group ID
snapshot SSD strategy is metadata
cloud provider id is 0

{'default' :
  {'Policy Number': -2,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':1, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 1,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
'system':
  {'Policy Number': -1,
   'Files matched': {'head':0, 'snapshot': 0},
   'Directories matched': {'head':0, 'snapshot': 0},
   'ADS containers matched': {'head':0, 'snapshot': 0},
   'ADS streams matched': {'head':0, 'snapshot': 0},
   'Access changes skipped': 0,
   'Protection changes skipped': 0,
   'File creation templates matched': 0,
   'File data placed on HDDs': {'head':0, 'snapshot': 0},
   'File data placed on SSDs': {'head':0, 'snapshot': 0},
  },
```

isi filepool default-policy modify

Modifies default file pool policy settings. The default file pool policy specifies storage settings for all files to which a higher-priority file pool policy does not apply.

Syntax

```
isi filepool default-policy modify
  [--data-access-pattern {random | concurrency | streaming}]
  [--set-requested-protection {default | +1 | +2:1 | +2 | +3:1 |
+3 | +4 | 2x | 3x | 4x | 5x | 6x |
7x | 8x}]
  [--data-storage-target <string>]
  [--data-ssd-strategy {metadata | metadata-write | data | avoid}]
  [--snapshot-storage-target <string>]
  [--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
  [--enable-coalescer {yes | no}]
  [--cloud-pool <string>]
  [--cloud-accessibility {cached | no-cache}]
  [--cloud-cache-expiration <duration>]
  [--cloud-compression-enabled {yes | no}]
  [--cloud-data-retention <duration>]
  [--cloud-encryption-enabled {yes | no}]
  [--cloud-full-backup-retention <duration>]
  [--cloud-incremental-backup-retention <duration>]
  [--cloud-read-ahead <string>]
  [--cloud-writeback-frequency <duration>]
  [--cloud-archive-snapshot-files {yes | no}]
  [--verbose]
```

Options

--data-access-pattern <string>

Specifies the preferred data access pattern, one of random, streaming, or concurrent.

--set-requested-protection <string>

Specifies the requested protection for files that match this filepool policy (for example, +2:1).

--data-storage-target <string>

Specifies the node pool or tier to which the policy moves files on the local cluster.

--data-ssd-strategy <string>

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in addition to the number required to satisfy the requested protection. Enabling global namespace acceleration (GNA) makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly more storage space compared with the other SSD strategy options.

--snapshot-storage-target <integer>

The ID of the node pool or tier chosen for storage of snapshots.

--snapshot-ssd-strategy <string>

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

--enable-coalescer {yes | no}

Enable or disable the coalescer, also referred to as SmartCache. The coalescer protects write-back data through a combination of RAM and stable storage. It is enabled by default, and should be disabled only in cooperation with EMC Isilon Customer Support.

--cloud-pool <string>

Specifies the default CloudPool and, therefore, the cloud storage account where cloud data is to be archived.

--cloud-accessibility {cached | no-cache}

Specifies whether, when a SmartLink file is accessed, cloud data is incrementally downloaded (cached) as needed, or fully downloaded (not cached).

--cloud-cache-expiration <duration>

Specifies the minimum amount of time until the cache expires. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-compression-enabled {yes | no}

Specifies whether data is to be compressed when archived to the cloud.

--cloud-data-retention <duration>

Specifies the minimum amount of time that archived data will be retained in the cloud after a SmartLink file is deleted from the cluster. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-encryption-enabled {yes | no}

Specifies whether data is to be encrypted when archived to the cloud.

--cloud-full-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of a full backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-incremental-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of an incremental backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-read-ahead {partial | full}

Specifies the cache readahead strategy when SmartLink files are accessed. A partial strategy means that only the amount of data needed by the user is cached. A full strategy means that all file data will be cached when the user accesses a SmartLink file.

--cloud-writeback-frequency <duration>

Specifies the minimum amount of time to wait before OneFS updates cloud data with local changes. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-archive-snapshot-files {yes | no}

Whether or not policies should archive files with snapshots.

--verbose

Displays more detailed information.

Example

The command shown in the following example modifies the default file pool policy in several ways. The command sets the `requested-protection-level` to **+2:1**, sets the `data-storage-target` to **anywhere** (the system default), and changes the `data--ssd-strategy` to **metadata-write**.

```
isi filepool default-policy modify --set-requested-protection=+2:1 --data-storage-target=anywhere --data-ssd-strategy=metadata-write
```

isi filepool default-policy view

View default file pool policy settings. The default file pool policy specifies storage settings for all files to which a higher-priority file pool policy does not apply.

Syntax

```
isi filepool default-policy view
```

The following display shows sample output from the command:

```
Apply Order: -
File Matching Pattern: -
Set Requested Protection: default
  Data Access Pattern: concurrency
  Enable Coalescer: True
  Data Storage Target: anywhere
  Data SSD Strategy: metadata
Snapshot Storage Target: anywhere
Snapshot SSD Strategy: metadata
Cloud Pool: -
Cloud Compression Enabled: -
Cloud Encryption Enabled: -
Cloud Data Retention: -
Cloud Incremental Backup Retention: -
Cloud Full Backup Retention: -
Cloud Accessibility: -
Cloud Read Ahead: -
Cloud Cache Expiration: -
Cloud Writeback Frequency: -
Cloud Archive Snapshot Files: -
```

isi filepool policies create

Create a custom file pool policy to identify a specific storage target and perform other actions on matched files and directories.

Syntax

```
isi filepool policies create <name>
[--description <string>]
[--begin-filter{<predicate> <operator> <link>}...--end-filter]
[--apply-order <integer>]
[--data-access-pattern {random | concurrency | streaming}]
[--set-requested-protection {default | +1 | +2:1 | +2 | +3:1 | +3 | +4 | 2x | 3x | 4x | 5x | 6x | 7x | 8x}]
[--data-storage-target <string>]
[--data-ssd-strategy {metadata | metadata-write | data | avoid}]
[--snapshot-storage-target <string>]
[--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
[--enable-coalescer {Yes | No}]
[--enable-packing {Yes | No}]
```

```

[--cloud-pool <string>]
[--cloud-accessibility {cached | no-cache}]
[--cloud-cache-expiration <duration>]
[--cloud-compression-enabled {yes | no}]
[--cloud-data-retention <duration>]
[--cloud-encryption-enabled {yes | no}]
[--cloud-full-backup-retention <duration>]
[--cloud-incremental-backup-retention <duration>]
[--cloud-read-ahead <string>]
[--cloud-writeback-frequency <duration>]
[--cloud-archive-snapshot-files {yes | no}]
[--verbose | -v]

```

Options

<name>

Specifies the name of the file pool policy to create.

--begin-filter {<predicate> <operator> <link>}... **--end-filter**

Specifies the file-matching criteria that determine the files to be managed by the filepool policy.

Each file matching criterion consists of three parts:

- Predicate. Specifies what attribute(s) to filter on. You can filter by path, name, file type, timestamp, or custom attribute, or use a combination of these attributes.
- Operator. Qualifies an attribute (for example, birth time) to describe a relationship to that attribute (for example, before).
- Link - Combines attributes using AND and OR statements.

The following predicates are valid:

--size=<nn>[{B | KB | MB | GB | TB | PB}**]**

Selects files according to the specified size.

--path=<pathname>

Selects files relative to the specified pathname.

--file-type= <value>

Selects only the specified file-system object type.

The following values are valid:

file	Specifies regular files.
directory	Specifies directories.
other	Specifies links.

--name= <value> [--case-sensitive= {true | false}**]**

Selects only files whose names match the specified string. Use **--case-sensitive=true** to enable case-sensitivity.

When forming the name, you can include the following wildcards:

- *
- []
- ?

--birth-time=<timestamp>

Selects files that were created relative to the specified date and time. Timestamp arguments are formed as **YYYY-MM-DDTHH:MM:SS**. For example, **2013-09-01T08:00:00** specifies a timestamp of September 1, 2013 at 8:00 A.M. You can use **--operator=** with an argument of **gt** to mean after the timestamp or **lt** to mean before the timestamp.

--changed-time=<timestamp>

Selects files that were modified relative to the specified date and time.

--metadata-changed-time=<timestamp>

Selects files whose metadata was modified relative to the specified date and time.

--accessed-time=<timestamp>

Selects files that were accessed at the specified time interval.

--custom-attribute=<value>

Selects files based on a custom attribute.

You can use the `operator=` option to specify a qualifier for the file-matching criterion. Specify operators in the following form:

```
--operator=<value>
```

The following operator values are valid:

Value	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

Link arguments can be used to specify multiple file-matching criteria. The following links are valid:

--and

Connects two file-matching criteria where files must match both criteria.

--or

Connects two file-matching criteria where files must match one or the other criteria.

--description <string>

Specifies a description of the filepool policy

--apply-order <integer>

Specifies the order index for execution of this policy.

--data-access-pattern <string>

Data access pattern random, streaming or concurrent.

--set-requested-protection <string>

Specifies a protection level for files that match this filepool policy (e.g., +3, +2:3, 8x).

--data-storage-target <string>

The name of the node pool or tier to which the policy moves files on the local cluster. If you do not specify a data storage target, the default is **anywhere**.

--data-ssd-strategy <string>

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in

addition to the number required to satisfy the requested protection. Enabling GNA makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly increases storage requirements compared with the other SSD strategy options.

--snapshot-storage-target <string>

The name of the node pool or tier chosen for storage of snapshots. If you do not specify a snapshot storage target, the default is **anywhere**.

--snapshot-ssd-strategy <string>

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

--enable-coalescer {Yes | No}

Enables the coalescer.

--enable-packing {Yes | No}

Enables packing.

--cloud-pool <string>

Specifies the default CloudPool and, therefore, the cloud storage account where cloud data is to be archived.

--cloud-accessibility {cached | no-cache}

Specifies whether, when a SmartLink file is accessed, cloud data is incrementally downloaded (cached) as needed, or fully downloaded (not cached).

--cloud-cache-expiration <duration>

Specifies the minimum amount of time until the cache expires. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-compression-enabled {yes | no}

Specifies whether data is to be compressed when archived to the cloud.

--cloud-data-retention <duration>

Specifies the minimum amount of time that archived data will be retained in the cloud after a SmartLink file is deleted from the cluster. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-encryption-enabled {yes | no}

Specifies whether data is to be encrypted when archived to the cloud.

--cloud-full-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of a full backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-incremental-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of an incremental backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-read-ahead {partial | full}

Specifies the cache readahead strategy when SmartLink files are accessed. A partial strategy means that only the amount of data needed by the user is cached. A full strategy means that all file data will be cached when the user accesses a SmartLink file.

--cloud-writeback-frequency <duration>

Specifies the minimum amount of time to wait before OneFS updates cloud data with local changes. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--verbose

Displays more detailed information.

Examples

The following example creates a file pool policy that moves all files in directory `/ifs/data/chemical/arco/finance` to the local storage target named `Archive_2`.

```
isi filepool policies create Save_Fin_Data --begin-filter
--path=/ifs/data/chemical/arco/finance --end-filter
--data-storage-target Archive_2 --data-ssd-strategy=metadata
```

The following example matches older files that have not been accessed or modified later than specified dates, and moves the files to an archival tier of storage.

```
isi filepool policies create archive_old
--data-storage-target ARCHIVE_1 --data-ssd-strategy avoid
--begin-filter --file-type=file --and --birth-time=2013-09-01
--operator=lt --and --accessed-time=2013-12-01 --operator=lt
--and --changed-time=2013-12-01 --operator=lt --end-filter
```

isi filepool policies delete

Delete a custom file pool policy. The default file pool policy cannot be deleted.

To list all file pool policies, run the `isi filepool policies list` command.

Syntax

```
isi filepool policies delete <name>
  [--force]
  [--verbose]
```

Options

<name>

Specifies the name of the file pool policy to be deleted.

--force

Deletes the file pool policy without asking for confirmation.

--verbose

Displays more detailed information.

Example

The following command deletes a file pool policy named `ARCHIVE_OLD`. The `--force` option circumvents the requirement to confirm the deletion.

```
isi filepool policies delete ARCHIVE_OLD --force
```

isi filepool policies list

List all custom file pool policies configured on the system.

Syntax

```
isi filepool policies list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit <integer>

Specifies a limit to the number of policies that are displayed.

--format

Output the list of file pool policies in a variety of formats. The following options are valid: `table` , `json` , `csv` , and `list` .

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following example lists custom file pool policies in .csv format and outputs the list to a file in the OneFS file system.

```
isi filepool policies list --format csv > /ifs/data/policy.csv
```

isi filepool policies modify

Modify a custom file pool policy.

Syntax

```
isi filepool policies modify <id>
[--name <string>]
[--description <string>]
[--begin-filter{<predicate> <operator> <link>}...--end-filter]
[--apply-order <integer>]
[--data-access-pattern {random | concurrency | streaming }]
[--clear-data-access-pattern ]
[--set-requested-protection {{default | +1 | +2:1 | +2 | +3:1 | +3 | +4 | 2x | 3x | 4x | 5x
| 6x | 7x | 8x} | --clear-set-requested-protection}]
[--clear-set-requested-protection ]
[--data-storage-target <string>]
[--clear-data-storage-target]
[--data-ssd-strategy {metadata-write | data | avoid}]
[--snapshot-storage-target <string>]
```

```

[--clear-snapshot-storage-target]
[--snapshot-ssd-strategy {metadata | metadata-write | data | avoid}]
[--enable-coalescer {Yes | No}]
[--clear-enable-coalescer]
[--enable-packing Yes | No]
[--clear-enable-packing]
[--cloud-pool <string>]
[--clear-cloud-pool]
[--cloud-accessibility {cached | no-cache}]
[--cloud-cache-expiration <duration>]
[--cloud-compression-enabled {yes | no}]
[--cloud-data-retention <duration>]
[--cloud-encryption-enabled {yes | no}]
[--cloud-full-backup-retention <duration>]
[--cloud-incremental-backup-retention <duration>]
[--cloud-read-ahead <string>]
[--cloud-writeback-frequency <duration>]
[--cloud-archive-snapshot-files {yes | no}]
[--verbose]

```

Options

<id>

Specifies an ID for the filepool policy.

<name>

Specifies the name of the file pool policy to create.

--begin-filter {<predicate> <operator> <link>}... --end-filter

Specifies the file-matching criteria that determine the files to be managed by the filepool policy.

Each file matching criterion consists of three parts:

- Predicate. Specifies what attribute(s) to filter on. You can filter by path, name, file type, timestamp, or custom attribute, or use a combination of these attributes.
- Operator. Qualifies an attribute (for example, birth time) to describe a relationship to that attribute (for example, before).
- Link - Combines attributes using AND and OR statements.

The following predicates are valid:

--size=<nn>[{B | KB | MB | GB | TB | PB}**]**

Selects files according to the specified size.

--path=<pathname>

Selects files relative to the specified pathname.

--file-type= <value>

Selects only the specified file-system object type.

The following values are valid:

file	Specifies regular files.
directory	Specifies directories.
other	Specifies links.

--name= <value> [--case-sensitive= {true | false}]

Selects only files whose names match the specified string. Use `--case-sensitive=true` to enable case-sensitivity.

When forming the name, you can include the following wildcards:

- *
- []
- ?

--birth-time=<timestamp>

Selects files that were created relative to the specified date and time. Timestamp arguments are formed as **YYYY-MM-DDTHH:MM:SS**. For example, **2013-09-01T08:00:00** specifies a timestamp of September 1, 2013 at 8:00 A.M. You can use **--operator=** with an argument of **gt** to mean after the timestamp or **lt** to mean before the timestamp.

--changed-time=<timestamp>

Selects files that were modified relative to the specified date and time.

--metadata-changed-time=<timestamp>

Selects files whose metadata was modified relative to the specified date and time.

--accessed-time=<timestamp>

Selects files that were accessed at the specified time interval.

--custom-attribute=<value>

Selects files based on a custom attribute.

You can use the **operator=** option to specify a qualifier for the file-matching criterion. Specify operators in the following form:

```
--operator=<value>
```

The following operator values are valid:

Value	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

Link arguments can be used to specify multiple file-matching criteria. The following links are valid:

--and

Connects two file-matching criteria where files must match both criteria.

--or

Connects two file-matching criteria where files must match one or the other criteria.

--description <string>

Specifies a description of the filepool policy

--apply-order <integer>

Specifies the order index for execution of this policy.

--data-access-pattern <string>

Specifies the data access pattern as random, streaming or concurrent.

--clear-data-access-pattern

Removes the action to set the data access pattern on matching files.

--set-requested-protection <string>

Specifies a protection level for files that match this filepool policy (for example, +3, +2:3, 8x).

--clear-set-requested-protection

Removes the action to set the requested protection on matching files.

--data-storage-target <string>

The name of the node pool or tier to which the policy moves files on the local cluster.

--clear-data-storage-target

Removes the action to set the data storage target on matching files.

--data-ssd-strategy <string>

Specifies how to use SSDs to store local data.

avoid

Writes all associated file data and metadata to HDDs only.

metadata

Writes both file data and metadata to HDDs. This is the default setting. An extra mirror of the file metadata is written to SSDs, if SSDs are available. The SSD mirror is in addition to the number required to satisfy the requested protection. Enabling GNA makes read acceleration available to files in node pools that do not contain SSDs.

metadata-write

Writes file data to HDDs and metadata to SSDs, when available. This strategy accelerates metadata writes in addition to reads but requires about four to five times more SSD storage than the **Metadata** setting. Enabling GNA does not affect read/write acceleration.

data

Uses SSD node pools for both data and metadata, regardless of whether global namespace acceleration is enabled. This SSD strategy does not result in the creation of additional mirrors beyond the normal requested protection but requires significantly increases storage requirements compared with the other SSD strategy options.

--snapshot-storage-target <string>

The name of the node pool or tier chosen for storage of snapshots.

--clear-snapshot-storage-target

Removes the action to set the snapshot storage target on matching files.

--snapshot-ssd-strategy <string>

Specifies how to use SSDs to store snapshots. Valid options are `metadata`, `metadata-write`, `data`, `avoid`. The default is `metadata`.

--enable-coalescer {Yes | No}

Enables the coalescer.

--clear-enable-coalescer

Removes the action to change the coalescer setting on matching files.

--enable-packing

Enables packing.

--clear-enable-packing

Removes the action to change the packing setting on matching files.

--cloud-pool <string>

Specifies the default CloudPool and, therefore, the cloud storage account where cloud data is to be archived.

--clear-cloud-pool

Remove the action to set a cloud storage target on matching files.

--cloud-accessibility {cached | no-cache}

Specifies whether, when a SmartLink file is accessed, cloud data is incrementally downloaded (cached) as needed, or fully downloaded (not cached).

--cloud-cache-expiration <duration>

Specifies the minimum amount of time until the cache expires. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-compression-enabled {yes | no}

Specifies whether data is to be compressed when archived to the cloud.

--cloud-data-retention <duration>

Specifies the minimum amount of time that archived data will be retained in the cloud after a SmartLink file is deleted from the cluster. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-encryption-enabled {yes | no}

Specifies whether data is to be encrypted when archived to the cloud.

--cloud-full-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of a full backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-incremental-backup-retention <duration>

Specifies the minimum amount of time that cloud files will be retained after the creation of an incremental backup. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--cloud-read-ahead {partial | full}

Specifies the cache readahead strategy when SmartLink files are accessed. A partial strategy means that only the amount of data needed by the user is cached. A full strategy means that all file data will be cached when the user accesses a SmartLink file.

--cloud-writeback-frequency <duration>

Specifies the minimum amount of time to wait before OneFS updates cloud data with local changes. A number followed by a unit of time is accepted. For example, a setting of 9H would specify a nine-hour duration. Similarly, a setting of 2D would specify a two-day duration.

--verbose

Display more detailed information.

Examples

The following example modifies a file pool policy to move matched files to a different local storage target named `Archive_4`. The next time the SmartPools job runs, matched files would be moved to the new storage target.

```
isi filepool policies modify Save_Fin_Data --begin-filter
--path=/ifs/data/chemical/arco/finance --end-filter
--data-storage-target Archive_4 --data-ssd-strategy=metadata
```

The following example matches older files that have not been accessed or modified later than specified dates, and moves the files to an archival tier of storage.

```
isi filepool policies modify archive_old
--data-storage-target ARCHIVE_1 --data-ssd-strategy avoid
--begin-filter --file-type=file --and --birth-time=2013-06-01
--operator=lt --and --accessed-time=2013-09-01 --operator=lt
--and --changed-time=2013-09-01 --operator=lt --end-filter
```

isi filepool policies view

Displays detailed information about a custom file pool policy.

Syntax

```
isi filepool policies view <name>
```

Options

<name>

Specifies the name of the file pool policy to view. Run the `isi filepool policies list` command to list the names of all custom file pool policies.

Example

The following example displays details about a file pool policy named `my_policy`:

```
isi filepool policies view my_policy
```

Output from the command would look similar to the following display:

```
                Name: my_policy
                Description: Archive older files to the cloud
                State: OK
                State Details:
                Apply Order: 1
                File Matching Pattern: Path == data/old_files (begins with) AND Name == *.*
                Set Requested Protection: -
                Data Access Pattern: -
                Enable Coalescer: -
                Data Storage Target: -
                Data SSD Strategy: -
                Snapshot Storage Target: -
                Snapshot SSD Strategy: -
                Cloud Pool: my_s3_pool
                Cloud Compression Enabled: False
                Cloud Encryption Enabled: False
                Cloud Data Retention: 604800
                Cloud Incremental Backup Retention: 604800
                Cloud Full Backup Retention: 157680000
                Cloud Accessibility: cached
                Cloud Read Ahead: partial
                Cloud Cache Expiration: 86400
                Cloud Writeback Frequency: 14400
                Cloud Archive Snapshot Files: False
```

isi filepool templates list

Lists available file pool policy templates.

Syntax

```
isi filepool templates list
  [--limit <integer>]
  [--sort <string>]
  [--descending <string>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Specifies the number of templates to display.

--sort <string>

Sorts data by the field specified.

--descending | -d

Sorts data in descending order.

--format

Displays file pool templates in the specified format. The following values are valid:

table
json
csv
list

--no-header | **-a**

Displays table and CSV output without headers.

--no-footer | **-z**

Displays table output without footers.

--verbose | **-v**

Displays more detailed information.

isi filepool templates view

View the detailed settings in a file pool policy template.

Syntax

```
isi filepool templates view <name>
```

Options

<name>

The name of the template to view.

isi_for_array

Runs commands on multiple nodes in an array, either in parallel or in serial.

When options conflict, the one specified last takes precedence.

 **NOTE: The `-k`, `-u`, `-p`, and `-q` options are valid only for SSH transport.**

Syntax

```
isi_for_array  
  [--array-name <array>]  
  [--array-file <filename>]  
  [--directory <directory>]  
  [--diskless]  
  [--ignore-errors]  
  [--known-hosts-file <filename>]  
  [--user <user>]  
  [--nodes <nodes>]  
  [--password <password>]  
  [--pre-command <command>]  
  [--query-password]  
  [--quiet]  
  [--serial]  
  [--storage]  
  [--transport <transport-type>]  
  [--throttle <count>]  
  [--exclude-nodes <nodes>]  
  [--exclude-down-nodes]
```

Options

{--array-name | -a} <array>

Uses <array>.

{--array-file | -A} <filename>

Reads array information from <filename>. The default looks first for \$HOME/.array.xml, then for /etc/ifs/array.xml.

{--directory | -d} <directory>

Runs commands from the specified directory on remote computers. The current working directory is the default directory. An empty <directory> results in commands being run in the user's home directory on the remote computer.

{--diskless | -D}

Runs commands from diskless nodes.

{--ignore-errors | -I}

Suppresses the printing of error messages for nodes that return non-zero exit status. Returns the maximum exit status from all nodes.

{--known-hosts-file | -k} <filename>

Uses <filename> for SSH known hosts file instead of the default /dev/null directory.

{--user | -u | -l} <user>

Logs in as <user> instead of as the default root user.

{--nodes | -n} <nodes>

Runs commands on the specified nodes, which can be specified multiple times. Must be a list of either node names or ranges of node IDs; for example, 1, 3-5, *neal8*, 10. If no nodes are explicitly listed, the whole array is used.

{--password | -p | --pw} <password>

Uses the specified password instead of the default password.

{--pre-command | -P} <command>

Runs the specified command before any other commands. This is useful for setting up the environment and it can be specified multiple times. You can specify - to reset the list of pre-commands.

{--query-password | -q}

Prompts the user for a password.

{--quiet | -Q}

Suppresses printing of the host prefix for each output line.

{--serial | -s}

Runs commands in serial instead of parallel.

{--storage | -S}

Run commands from storage nodes.

{--transport | -t} <transport-type>

Specifies the network transport type. The default value is *rpc*. Valid transports values are *rpc* or *ssh*.

{--throttle | -T} <count>

Adjusts throttling. To disable throttling, specify **0**. The default value is 80 for multitasking and 35 for forking.

{--exclude-nodes | -x} <nodes>

Excludes specified nodes from the command. This argument is specified in the same manner as the *-n* option.

{--exclude-down-nodes | -X}

Excludes offline nodes from the command. This command is limited to cluster local use only.

Example

In SmartLock compliance mode, to run `isi_for_array` for a command that requires root privileges, you must specify `sudo` twice. For example, the following command runs `isi statistics client list` on each node in a compliance cluster.

```
sudo isi_for_array -u compadmin sudo isi statistics client list
```

isi ftp settings modify

Modifies cluster FTP settings.

Syntax

```
isi ftp settings modify
  [--accept-timeout <duration>]
  [--revert-accept-timeout]
  [--allow-anon-access {yes | no}]
  [--revert-allow-anon-access]
  [--allow-anon-upload {yes | no}]
  [--revert-allow-anon-upload]
  [--allow-dirlists {yes | no}]
  [--revert-allow-dirlists]
  [--allow-downloads {yes | no}]
  [--revert-allow-downloads]
  [--allow-local-access {yes | no}]
  [--revert-allow-local-access]
  [--allow-writes {yes | no}]
  [--revert-allow-writes]
  [--always-chdir-homedir {yes | no}]
  [--revert-always-chdir-homedir]
  [--anon-chown-username <string>]
  [--revert-anon-chown-username]
  [--anon-password-list <string>...]
  [--clear-anon-password-list]
  [--add-anon-password-list <string>...]
  [--remove-anon-password-list <string>...]
  [--revert-anon-password-list]
  [--anon-root-path <path>]
  [--revert-anon-root-path]
  [--anon-umask <integer-octal>]
  [--revert-anon-umask]
  [--ascii-mode {off | upload | download | both}]
  [--revert-ascii-mode]
  [--chroot-exception-list <string>...]
  [--clear-chroot-exception-list]
  [--add-chroot-exception-list <string>...]
  [--remove-chroot-exception-list <string>...]
  [--revert-chroot-exception-list]
  [--chroot-local-mode {all | none | all-with-exceptions | none-with-exceptions}]
  [--revert-chroot-local-mode]
  [--connect-timeout <duration>]
  [--revert-connect-timeout]
  [--data-timeout <duration>]
  [--revert-data-timeout]
  [--denied-user-list <string>...]
  [--clear-denied-user-list]
  [--add-denied-user-list <string>...]
  [--remove-denied-user-list <string>...]
  [--revert-denied-user-list]
  [--dirlist-localtime {yes | no}]
  [--revert-dirlist-localtime]
  [--dirlist-names {numeric | textual | hide}]
  [--revert-dirlist-names]
  [--file-create-perm <integer-octal>]
  [--revert-file-create-perm]
  [--limit-anon-passwords {yes | no}]
  [--revert-limit-anon-passwords]
```

```

[--local-root-path <path>]
[--revert-local-root-path]
[--local-umask <integer-octal>]
[--revert-local-umask]
[--server-to-server {yes | no}]
[--revert-server-to-server]
[--session-support {yes | no}]
[--revert-session-support]
[--session-timeout <duration>]
[--revert-session-timeout]
[--user-config-dir <path>]
[--revert-user-config-dir]
[--service {yes | no}]

```

Options

--accept-timeout <duration>

Specifies the time, in seconds, that a remote client has to establish a PASV style data connection before timeout. All integers between 30 and 600 are valid values. The default value is 60.

--revert-accept-timeout

Sets the value to the system default for `--accept-timeout`.

--allow-anon-access {yes | no}

Controls whether anonymous logins are permitted. If enabled, both the usernames ftp and anonymous are recognized as anonymous logins. The default value is No.

--revert-allow-anon-access

Sets the value to the system default for `--allow-anon-access`.

--allow-anon-upload {yes | no}

Controls whether anonymous users are able to upload files under certain conditions. For anonymous users to be able to upload, you must set the `--allow-writes` option to Yes, and the anonymous user must have write permission on the desired upload location. The default value is Yes.

--revert-allow-anon-upload

Sets the value to the system default for `--allow-anon-upload`.

--allow-dirlists {yes | no}

Controls whether directory list commands are enabled. The default value is Yes.

--revert-allow-dirlists

Sets the value to the system default for `--allow-dirlists`.

--allow-downloads {yes | no}

Controls whether files can be downloaded. The default value is Yes.

--revert-allow-downloads

Sets the value to the system default for `--allow-downloads`.

--allow-local-access {yes | no}

Controls whether local logins are permitted. If set to Yes, normal user accounts can be used to log in. The default value is Yes.

--revert-allow-local-access

Sets the value to the system default for `--allow-local-access`.

--allow-writes {yes | no}

Sets and displays whether commands that change the file system are permitted. Controls whether any of the following commands are allowed:

- STOR
- DELE
- RNFR
- RNTD
- MKD
- RMD

- APPE
- SITE

The default value is `yes`.

--revert-allow-writes

Sets the value to the system default for `--allow-writes`.

--always-chdir-homedir {yes | no}

Controls whether FTP always initially changes directories to the home directory of the user. If set to `No`, you can set up a `chroot` area in FTP without having a home directory for the user. The default value is `Yes`.

--revert-always-chdir-homedir

Sets the value to the system default for `--always-chdir-homedir`.

--anon-chown-username <string>

Gives ownership of anonymously uploaded files to the specified user. The value must be a local username. The default value is `root`.

--revert-anon-chown-username

Sets the value to the system default for `--anon-chown-username`.

--anon-password-list <string>...

Displays the list of anonymous user passwords.

--clear-anon-password-list

Clears the list of passwords for anonymous users.

--add-anon-password-list <string>...

Adds items to list of passwords for anonymous users. Specify `--add-anon-password-list` for each additional password to add.

--remove-anon-password-list <string>...

Removes items from list of passwords for anonymous users. Specify `--remove-anon-password-list` for each additional password to remove.

--revert-anon-password-list

Sets the value to the system default for `--anon-password-list`.

--anon-root-path <path>

Displays and specifies the root path for anonymous users, which is a directory in `/ifs` that the Very Secure FTP Daemon (VSFTPD) will try to change to after an anonymous login. Valid paths are in `/ifs`. The default value is `/ifs/home/ftp`.

--revert-anon-root-path

Sets the value to the system default for `--anon-root-path`.

--anon-umask <integer-octal>

Specifies the umask for file creation by anonymous users. Valid values are octal umask numbers. The default value is `077`.



NOTE: The value must contain the 0 prefix; otherwise it will be interpreted as a base 10 integer.

--revert-anon-umask

Sets the value to the system default for `--anon-umask`.

--ascii-mode {off | upload | download | both}

Enables ASCII downloads, uploads, or both.

--revert-ascii-mode

Sets the value to the system default for `--ascii-mode`.

--chroot-exception-list <string>

Displays the list of local user `chroot` exceptions.

--clear-chroot-exception-list

Clears the list of local user `chroot` exceptions.

--add-chroot-exception-list <string>

Adds users to the `chroot` exception list.

--remove-chroot-exception-list <string>
Removes users from the `chroot` exception list.

--revert-chroot-exception-list
Sets the value to the system default for `--chroot-exception-list`.

--chroot-local-mode {all | none | all-with-exceptions | none-with-exceptions}
Specifies which users are placed in a `chroot` jail in their home directory after login.

--revert-chroot-local-mode
Sets the value to the system default for `--chroot-local-mode`.

--connect-timeout <duration>
Specifies the timeout in seconds for a remote client to respond to a PORT style data connection. Valid durations are integers between 30 and 600. The default value is 60 (one minute).

--revert-connect-timeout
Sets the value to the system default for `--connect-timeout`.

--data-timeout <duration>
Specifies the maximum time (in seconds) data transfers are allowed to stall with no progress before the remote client is removed. Valid durations are integers between 30 and 600. The default value is 300 (five minutes).

--revert-data-timeout
Sets the value to the system default for `--data-timeout`.

--denied-user-list <string>
Displays the list of denied users.

--clear-denied-user-list
Clears the list of denied users.

--add-denied-user-list <string>
Add users to the list of denied users.

--remove-denied-user-list <string>
Removes users from the list of denied users.

--revert-denied-user-list
Sets the value to the system default for `--denied-user-list` (empty).

--dirlist-localtime {yes | no}
Specifies whether the time displayed in directory listings is in your local time zone. Valid values are `Yes` and `No`. If `No`, time displays on GMT. If `Yes`, the time displays in your local time zone. The default value is `No`.

The last-modified times returned by commands issued inside of the FTP shell are also affected by this parameter.


--revert-dirlist-localtime
Sets the value to the system default for `--dirlist-localtime`.

--dirlist-names {numeric | textual | hide}
Determines what information is displayed about users and groups in directory listings. The following are valid:

numeric	Numeric IDs are shown in the user and group fields of directory listings.
textual	Names are shown in text format in the user and group fields of directory listings.
hide	All user and group information in directory listings is displayed as <code>ftp</code> . This is the default setting.

--revert-dirlist-names
Sets the value to the system default for `--dirlist-names`.

--file-create-perm <integer-octal>
Specifies the permissions with which uploaded files are created. Valid values are octal permission numbers. The default value is `0666`.

 **NOTE: For uploaded files to be executable, set the permissions to 0777.**

--revert-file-create-perm

Sets the value to the system default for `--file-create-perm`.

--limit-anon-passwords {yes | no}

Limits anonymous passwords.

--revert-limit-anon-passwords

Sets the value to the system default for `--limit-anon-passwords`.

--local-root-path <path>

Specifies the initial directory in `/ifs` for a local login. Valid paths are in `/ifs`. The default path is the local user home directory.

--revert-local-root-path

Sets the value to the system default for `--local-root-path`.

--local-umask <integer-octal>

Specifies the umask for file creation by local users. Valid values are octal umask numbers. The default value is 077.



NOTE: The value must contain the 0 prefix; otherwise it will be interpreted as a base 10 integer.

--revert-local-umask

Sets the value to the system default for `--local-umask`.

--server-to-server {yes | no}

Specifies whether to allow server-to-server (FXP) transfers. Valid values are Yes and No. The default value is No.

--revert-server-to-server

Sets the value to the system default for `--server-to-server`.

--session-support {yes | no}

Enables or disables FTP session support. If set to YES, the command maintains login sessions for each user through Pluggable Authentication Modules (PAM). If set to NO, the command prevents automatic home directory creation if that functionality is otherwise available. The default value is YES.

--revert-session-support

Sets the value to the system default for `--session-support`.

--session-timeout <duration>

Specifies the maximum time (in seconds) that a remote client may spend between FTP commands before the remote client is kicked off. Valid values are integers between 30 and 600. The default value is 300 (five minutes).

--revert-session-timeout

Sets the value to the system default for `--session-timeout`.

--user-config-dir <path>

Specifies the directory where user-specific configurations that override global configurations can be found. The default value is the local user home directory.

--revert-user-config-dir

Sets the value to the system default for `--user-config-dir`.

--service {yes | no}

Specifies whether the FTP service is enabled.

isi ftp settings view

Shows the FTP settings for the cluster.

Syntax

```
isi ftp settings view
```


Options

There are no options for this command.

Example

The following is an example of the output generated by this command:

```
Accept Timeout: 1m
Allow Anon Access: No
Allow Anon Upload: Yes
  Allow Dirlists: Yes
  Allow Downloads: Yes
Allow Local Access: Yes
  Allow Writes: Yes
Always Chdir Homedir: Yes
Anon Chown Username: root
Anon Password List: -
  Anon Root Path: /ifs/home/ftp
  Anon Umask: 0077
  Ascii Mode: off
Chroot Exception List: -
  Chroot Local Mode: none
  Connect Timeout: 1m
  Data Timeout: 5m
  Denied User List: -
  Dirlist Localtime: No
  Dirlist Names: hide
  File Create Perm: 0666
Limit Anon Passwords: Yes
  Local Root Path: -
  Local Umask: 0077
  Server To Server: No
  Session Support: Yes
  Session Timeout: 5m
  User Config Dir: -
FTP Service Enabled: No
```

isi_gather_info

Collects and uploads the most recent cluster log information to EMC Secure Remote Services (ESRS).

Multiple instances of `-i`, `-f`, `-s`, `-S`, and `-l` are allowed.

`gather_expr` and `analysis_expr` can be quoted.

The default temporary directory is `/ifs/data/Isilon_Support/` (change with `-L` or `-T`).

Syntax

```
isi_gather_info
[-h]
[-v]
[-u <user>]
[-p <password>]
[-i]
[--incremental]
[-l]
[-f <filename>]
[-n <nodes>]
[--local-only]
[--skip-node-check]
[-s gather-script]
[-S gather-expr]
[-l gather-expr]
[-a analysis-script]
[-A analysis-expr]
```

```

[-t <tarfile>]
[-x exclude_tool]
[-I]
[-L]
[-T <temp-dir>]
[--tardir <dir>]
[--symlinkdir <dir>]
[--varlog_recent]
[--varlog_all]
[--nologs]
[--group <name>]
[--clean-cores]
[--clean-all]
[--no-dumps]
[--dumps]
[--no-cores]
[--cores]
[--upgrade-archive]
[--debug]
[--verbose]
[--noconfig]
[--save-only]
[--save]
[--upload]
[--noupload]
[--re-upload <filename>]
[--verify-upload]
[--http]
[--nohttp]
[--http-host <host>]
[--http-path <dir>]
[--http-proxy <host>]
[--http-proxy-port <port>]
[--ftp]
[--noftp]
[--ftp-user <user>]
[--ftp-pass <password>]
[--ftp-host <host>]
[--ftp-path <dir>]
[--ftp-port <alt-port>]
[--ftp-proxy <host>]
[--ftp-proxy-port <port>]
[--ftp-mode <mode>]
[--esrs]
[--email]
[--noemail]
[--email-addresses]
[--email-from]
[--email-subject]
[--email-body]
[--skip-size-check]

```

Options

- h**
Prints this message and exits.
- v**
Prints version info and exits.
- u <user>**
Specifies the login as <user> instead of as the default root user.
- p <password>**
Uses <password>.
- i**

Includes only the listed utility. See also the `-l` option for a list of utilities to include. The special value `all` may be used to include every known utility.

--incremental

Gathers only those logs that changed since last log upload.

-l

Lists utilities and groups that can be included. See `-i` and `--group`.

-f <filename>

Gathers *<filename>* from each node. The value must be an absolute path.

-n <nodes>

Gathers information from only the specified nodes. Nodes must be a list or range of LNNs, for example, `1, 4-10, 12, 14`. If no nodes are specified, the whole array is used. Note that nodes are automatically excluded if they are down.

--local-only

Gathers information only from only the local node. Run this option when gathering files from the `/ifs` filesystem.

--skip-node-check

Skips the check for node availability.

-s gather-script

Runs *<gather-script>* on every node.

-S gather-expr

Runs *<gather-expr>* on every node.

-l gather-expr

Runs *<gather-expr>* on the local node.

-a analysis-script

Runs *<analysis-script>* on results.

-A analysis-expr

Runs *<analysis-expr>* on every node.

-t <tarfile>

Saves all results to the specified *<tarfile>* rather than to the default tar file.

-x exclude_tool

Excludes the specified tool or tools from being gathered from each node. Multiple tools can be listed as comma-separated values.

-I

Saves results to `/ifs`. This is the default setting.

-L

Save all results to local storage `/var/crash/support/`.

-T <temp-dir>

Saves all results to *<temp-dir>* instead of the default directory. `-T` overrides `-L` and `-I`.

--tardir <dir>

Places the final package directly into the specified directory.

--symlinkdir <dir>

Creates a symlink to the final package in the specified directory.

--varlog_recent

Gathers all logs in `/var/log`, with the exception of the compressed and rotated old logs. The default setting is all logs.

--varlog_all
Gathers all logs in `/var/log`, including compressed and rotated old logs. This is the default setting.

--nologs
Does not gather the required minimum number of logs.

--group <name>
Adds a specific group of utilities to the tar file.

--clean-cores
Deletes cores from `/var/crash` after successful compression of the package.

--clean-dumps
Deletes dumps from `/var/crash` after successful compression of the package.

--clean-all
Deletes cores and dumps from `/var/crash` after successful compression of the package.

--no-dumps
Does not gather hang dumps for the package.

--dumps
Adds cores to the package.

--no-cores
Does not gather cores for the package.

--cores
Adds dumps to the package.

--upgrade-archive
Adds the upgrade archive to the package.

--debug
Displays debugging messages.

--verbose
Displays more detailed information.

--noconfig
Uses built-in default values and bypasses the configuration file.

--save-only
Saves the CLI-specified configuration to file and exits.

--save
Saves the CLI-specified configuration to file and runs it.

--upload
Uploads logs to Isilon Technical Support automatically. This is the default setting.

--noupload
Specifies no automatic upload to Isilon Technical Support.

--re-upload <filename>
Re-uploads the specified `<filename>`.

--verify-upload
Creates a tar file and uploads to test connectivity.

--http
Attempts HTTP upload. This is the default setting.

--nohttp
Specifies no HTTP upload attempt.

--http-host <host>

Specifies an alternate HTTP site for upload.

--http-path <dir>
Specifies an alternate HTTP upload directory.

--http-proxy <host>
Specifies the proxy server to use.

--http-proxy-port <port>
Specifies the proxy port to use.

--ftp
Attempts FTP upload. This setting is the default value.

--noftp
Specifies no FTP upload attempt.

--ftp-user <user>
Specifies an alternate user for FTP (default: anonymous).

--ftp-pass <password>
Specifies an alternate password for FTP.

--ftp-host <host>
Specifies an alternate FTP site for upload.

--ftp-path DIR
Specifies an alternate FTP upload directory.

--ftp-port <alt-port>
Specifies an alternate FTP port for upload.

--ftp-proxy <host>
Specifies the proxy server to use.

--ftp-proxy-port <port>
Specifies the proxy port to use.

--ftp-mode <mode>
Specifies the mode of FTP file transfer. The following values are valid: *both*, *active*, *passive*. The default value is *both*.

--esrs
Attempts ESRS upload.

--email
Attempts SMTP upload. If set, SMTP is tried first.

--noemail
Specifies no SMTP upload attempt. This is the default value.

--email-addresses
Specifies email addresses as comma-separated strings.

--email-from
Specifies the sender's email address.

--email-subject
Specifies an alternative email subject.

--email-body
Specifies alternative email text shown on head of body.

--skip-size-check
Does not check the size of the gathered file.

isi get

Displays information about a set of files, including the requested protection, current actual protection, and whether write-coalescing is enabled.

Requested protection appears in one of three colors: green, yellow, or red. Green indicates full protection. Yellow indicates degraded protection under a mirroring policy. Red indicates a loss of one or more data blocks under a parity policy.

Syntax

```
isi get {[[-a] [-d] [-g] [-s] [{"-D | -DD | -DDC}]] [-R] <path>}  
| {[[-g] [-s] [{"-D | -DDO | -DDC}]] [-R] -L <lin>}}
```

Options

- a**
Displays the hidden "." and ".." entries of each directory.
- d**
Displays the attributes of a directory instead of the contents.
- g**
Displays detailed information, including snapshot governance lists.
- s**
Displays the protection status using words instead of colors.
- D**
Displays more detailed information.
- DDO**
Includes information about protection groups and security descriptor owners and groups.
- DDC**
Includes cyclic redundancy check (CRC) information.
- R**
Displays information about the subdirectories and files of the specified directories.
- <path>**
Displays information about the specified file or directory.
Specify as a file or directory path.
- L <lin>**
Displays information about the specified file or directory.
Specify as a file or directory LIN.

Examples

The following command displays information on `ifs/home/` and all of its subdirectories:

```
isi get -R /ifs/home
```

The system displays output similar to the following example:

```
POLICY    LEVEL PERFORMANCE COAL  FILE  
default  4x/2  concurrency on   ./  
default  8x/3  concurrency on   ../
```

```
default      4x/2 concurrency on   admin/
default      4x/2 concurrency on   ftp/

/ifs/home/admin:
default      4+2/2 concurrency on   .zshrc

/ifs/home/ftp:
default      4x/2 concurrency on   incoming/
default      4x/2 concurrency on   pub/

/ifs/home/ftp/incoming:

/ifs/home/ftp/pub:
```

isi hardening apply

Applies security hardening to the cluster.

Syntax

```
isi hardening apply <profile>
[--report {yes | no}]
[--verbose]
```

Options

<profile>

Specifies the hardening profile that will be applied to the Isilon cluster. Currently, OneFS supports only the DISA (Defense Information Systems Agency) STIG (Security Technology Implementation Guide) profile for security hardening on the cluster.

--report {yes | no}

Specifies whether to check the state of the Isilon cluster and report the results without actually applying the hardening profile. The system displays any issues it finds, which can be resolved by the hardening engine or deferred to be fixed manually.

{--verbose | -v}

Displays more detailed information.

isi hardening revert

Reverts all security hardening that has been applied to the cluster.

Syntax

```
isi hardening revert
[--verbose]
[--force]
```

Options

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Suppresses command-line prompts and messages to revert hardening.

isi hardening status

Displays the status of security hardening for the cluster and each cluster node, and indicates the hardening profile applied to the cluster.

Syntax

```
isi hardening status
```

Options

There are no options for this command.

isi hdfs crypto encryption-zones create

Create and name the HDFS encryption zones for transparent data encryption.

Syntax

```
isi hdfs crypto encryption-zones create <path><keyname>
```

Options

<path> **<keyname>**

Specifies a directory and key name for the encryption zone.

The encryption zone must be somewhere within the HDFS root directory for that zone.

isi hdfs crypto settings modify

Configures HDFS for transparent data encryption.

Syntax

```
isi hdfs crypto settings modify  
  [--kms-url <string>]
```

Options

--kms-url **<string>**

Specifies the URL of the Key Management Server.

isi hdfs crypto encryption-zones list

List the HDFS encryption zones.

Syntax

```
isi hdfs crypto encryption-zones list
```

Options

There are no options for this command.

isi hdfs crypto settings view

View settings for HDFS transparent data encryption.

Syntax

```
isi hdfs crypto settings view
```

Options

There are no options for this command.

isi hdfs fsimage job settings modify

Change the interval between successive FSImages.

Syntax

```
isi hdfs fsimage job settings modify  
  [--generation-interval <string>]  
  [--verbose]  
  [--zone <string>]
```

Options

- generation-interval <string>**
The interval between successive FSImages.
- help <string>**
Display help for this command.
- {--verbose | -v}**
Display more detailed information.
- zone <string>**

The access zone to which the HDFS settings apply.

isi hdfs fsimage job settings view

Review the frequency of an FSImage job.

Syntax

```
isi hdfs fsimage job settings view  
[--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs fsimage job view

Review the status of an FSImage job.

Syntax

```
isi hdfs fsimage job view  
[--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs fsimage latest delete

Delete the latest FSImage.

Syntax

```
isi hdfs fsimage latest delete  
[--zone <string>]  
[{--verbose | -v}]  
[{--force | -f}]
```

Options

--zone <string>

The access zone to which the HDFS settings apply.

{--verbose | -v}

Display more detailed information.

{--force | -f}

Do not prompt for confirmation.

isi hdfs fsimage latest view

Review the latest FSImage.

Syntax

```
isi hdfs fsimage latest view  
[--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs fsimage settings modify

Enable FSImage on the HDFS access zone. For more information, see the Additional Resources - Cloudera Navigator section of this guide.

<http://doc.isilon.com/onefs/hdfs/05-ifs-c-hdfs-admin-guide-overview-chapter.htm>

Syntax

```
isi hdfs fsimage settings modify  
[--enabled {yes | no}]  
[--zone <string>]  
[--verbose]
```

Options

--enabled {yes | no}

Enables or disables the HDFS FSImage service. Allow access to FSImage and start FSImage generation. The HDFS FSImage service is disabled by default. This service should only be enabled on a Hadoop-enabled Access Zone that will use Cloudera Navigator.

--help <string>

Display help for this command.

{--verbose | -v}

Display more detailed information.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs fsimage settings view

Review the status of FSImage.

Syntax

```
isi hdfs fsimage settings view
  [--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs inotify settings modify

Enable INotify on the HDFS access zone.

Syntax

```
isi hdfs inotify settings modify
  [--enabled {yes | no}]
  [--maximum-delay <string>]
  [--retention <string>]
  [--zone <string>]
  [--verbose]
```

Options

--enabled {yes | no}

Allows access to FSImage and starts FSImage generation. The HDFS FSImage service is disabled by default. This service should only be enabled on a Hadoop-enabled access zone that will use Cloudera Navigator.

--help <string>

Display help for this command.

--maximum-delay <string>

The maximum duration until an edit event is reported in INotify.

--retention <string>

The minimum duration edits will be retained.

{--verbose | -v}

Display more detailed information.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs inotify settings view

Review the configuration of the INotify stream.

Syntax

```
isi hdfs inotify settings view  
[--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs inotify stream reset

Reset the INotify stream by deleting collected events.

Syntax

```
isi hdfs inotify stream reset  
[--zone <string>]  
[{--verbose | -v}]  
[{--force | -f}]
```

Options

--zone <string>

The access zone to which the HDFS settings apply.

{--verbose | -v}

Display more detailed information.

{--force | -f}

Do not prompt for confirmation.

isi hdfs inotify stream view

Review the INotify stream.

Syntax

```
isi hdfs inotify stream view  
[--zone <string>]
```

Options

--help <string>

Display help for this command.

--zone <string>

The access zone to which the HDFS settings apply.

isi hdfs log-level modify

Modifies the log level of the HDFS service on the node.

Syntax

```
isi hdfs log-level modify
  [--set {always|error|warning|info|verbose|debug|trace|default} ]
  [--verbose | -v]
```

Options

--set {always | error | warning | info | verbose | debug | trace | default}

Sets the default logging level for the HDFS service on the cluster. The default value is `default`.

--verbose | -v

Displays more detailed information.

isi hdfs log-level view

Displays the current log level of the HDFS service on the node.

Syntax

```
isi hdfs log-level view
```

Options

There are no options for this command.

isi hdfs proxyusers create

Creates a proxy user that can securely impersonate another user or group.

Syntax

```
isi hdfs proxyusers create <proxyuser-name>
  [--zone <zone-name>]
  [--add-group <group-name>...]
  [--add-gid <group-identifier>...]
  [--add-user <user-name>...]
  [--add-uid <user-identifier>...]
  [--add-sid <security-identifier>...]
  [--add-wellknown <well-known-name>...]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the user name of a user currently configured on the cluster to be designated as a proxy user.

--zone <zone-name>

Specifies the access zone the user authenticates through.

--add-group <group-name>...

Adds the group specified by name to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple group names in a comma-separated list.

--add-gid <group-identifier>...

Adds the group by specified by UNIX GID to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple UNIX GIDs in a comma-separated list.

--add-user <user-name>...

Adds the user specified by name to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple user names in a comma-separated list.

--add-uid <user-identifier>...

Adds the user specified by UNIX UID to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple UNIX UIDs in a comma-separated list.

--add-sid <security-identifier>...

Adds the user, group of users, machine or account specified by Windows SID to the list of proxy user members. The object must authenticate to the same access zone as the proxy user. You can specify multiple Windows SIDs in a comma-separated list.

--add-wellknown <well-known-name>...

Adds the well-known user specified by name to the list of members the proxy user can impersonate. The well-known user must authenticate to the same access zone as the proxy user. You can specify multiple well-known user names in a comma-separated list.

{ **--verbose** | **-v**}

Displays more detailed information.

Examples

The following command designates `hadoop-user23` in `zone1` as a new proxy user:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds the group of users named `hadoop-users` to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 \  
--add-group=hadoop-users
```

The following command designates `hadoop-user23` in `zone1` as a new proxy user and adds UID 2155 to the list of members that the proxy user can impersonate:

```
isi hdfs proxyusers create hadoop-user23 --zone=zone1 --add-UID=2155
```

isi hdfs proxyusers delete

Deletes a proxy user.

Syntax

```
isi hdfs proxyusers delete <proxyuser-name>  
  [--zone <zone-name>]  
  [--force]  
  [--verbose]
```

Options

<proxyuser-name>

Specifies the user name of the proxy user to be deleted.

--zone <zone-name>

Specifies the access zone that the proxy user authenticates through.

{ --force | -f }

Deletes the specified proxy user without requesting confirmation.

{ --verbose | -v }

Displays more detailed information.

Examples

The following command deletes `hadoop-user23` in `zone1` from the list of proxy users:

```
isi hdfs proxyusers delete hadoop-user23 --zone=zone1
```

isi hdfs proxyusers list

Displays all proxy users that are configured in an access zone.

Syntax

```
isi hdfs proxyusers list
  [--zone <zone-name>]
  [--format {table | json | csv | list}]
  [--no-header ]
  [--no-footer ]
  [--verbose]
```

Options

--zone <zone-name>

Specifies the name of the access zone.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

Examples

The following command displays a list of all proxy users that are configured in `zone1`:

```
isi hdfs proxyusers list --zone=zone1
```

The system displays output similar to the following example:


```
Name
-----
hadoop-user23
hadoop-user25
hadoop-user28
-----
Total: 3
```

isi hdfs proxyusers members list

Displays the users and groups of users, known as members, that can be impersonated by a proxy user.

Syntax

```
isi hdfs proxyusers members list <proxyuser-name>
  [--zone <zone-name>]
  [--format {table | json | csv | list}]
  [--no-header ]
  [--no-footer ]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the name of the proxy user.

--zone <zone-name>

Specifies the access zone the proxy user authenticates through.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

Examples

The following command displays a detailed list of the users and groups that are members of proxy user hadoop-user23 in zone1:

```
isi hdfs proxyusers members list hadoop-user23 --zone=zone1 -v
```

The system displays output similar to the following example:

```
Type: user
Name: krb_user_005
   ID: UID:1004
-----
Type: group
Name: krb_users
   ID: SID:S-1-22-2-1003
-----
Type: wellknown
Name: LOCAL
   ID: SID:S-1-2-0
```

isi hdfs proxyusers modify

Modifies a proxy user that can securely impersonate another user or group.

Syntax

```
isi hdfs proxyusers modify <proxyuser-name>
  [--zone <zone-name>]
  [--add-group <group-name>...]
  [--add-gid <group-identifier>...]
  [--add-user <user-name>...]
  [--add-uid <user-identifier>...]
  [--add-sid <security-identifier>...]
  [--add-wellknown <well-known-name>...]
  [--remove-group <group-name>...]
  [--remove-gid <group-identifier>...]
  [--remove-user <user-name>...]
  [--remove-uid <user-identifier>...]
  [--remove-sid <security-identifier>...]
  [--remove-wellknown <well-known-name>...]
  [--verbose]
```

Options

<proxyuser-name>

Specifies the user name of the proxy user to be modified.

--zone <zone-name>

Specifies the access zone that the proxy user authenticates through.

--add-group <group-name>...

Adds the group specified by name to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple group names in a comma-separated list.

--add-gid <group-identifier>...

Adds the group specified by UNIX GID to the list of proxy user members. The proxy user can impersonate any user in the group. The users in the group must authenticate to the same access zone as the proxy user. You can specify multiple UNIX GIDs in a comma-separated list.

--add-user <user-name>...

Adds the user specified by name to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple user names in a comma-separated list.

--add-uid <user-identifier>...

Adds the user specified by UNIX UID to the list of members the proxy user can impersonate. The user must authenticate to the same access zone as the proxy user. You can specify multiple UNIX UIDs in a comma-separated list.

--add-sid <security-identifier>...

Adds the user, group of users, machine or account specified by Windows SID to the list of proxy user members. The object must authenticate to the same access zone as the proxy user. You can specify multiple Windows SIDs in a comma-separated list.

--add-wellknown <well-known-name>...

Adds the well-known user specified by name to the list of members the proxy user can impersonate. The well-known user must authenticate to the same access zone as the proxy user. You can specify multiple well-known user names in a comma-separated list.

--remove-group <group-name>...

Removes the group specified by name from the list of proxy user members so that the proxy user can no longer impersonate any user in the group. You can specify multiple group names in a comma-separated list.

--remove-gid <group-identifier>...

Removes the group specified by UNIX GID from the list of proxy user members so that the proxy user can no longer impersonate any user in the group. You can specify multiple UNIX GIDs in a comma-separated list.

--remove-user <user-name>...

Removes the user specified by name from the list of members the proxy user can impersonate. You can specify multiple user names in a comma-separated list.

--remove-uid <user-identifier>...

Removes the user specified by UNIX UID from the list of members the proxy user can impersonate. You can specify multiple UNIX UIDs in a comma-separated list.

--remove-sid <security-identifier>...

Removes the user, group of users, machine or account specified by Windows SID from the list of proxy user members. You can specify multiple Windows SIDs in a comma-separated list.

--remove-wellknown <well-known-name>...

Removes the well-known user specified by name from the list of members the proxy user can impersonate. You can specify multiple well-known user names in a comma-separated list.

{**--verbose** | **-v**}

Displays more detailed information.

Examples

The following command adds the well-known local user to, and removes the user whose UID is 2155 from, the list of members for proxy user `hadoop-user23` in zone1:

```
isi hdfs proxyusers modify hadoop-user23 --zone=zone1 \  
--add-wellknown=local --remove-uid=2155
```

isi hdfs proxyusers view

Displays the configuration details of a specific proxy user.

Syntax

```
isi hdfs proxyusers view <proxyuser-name>  
[--zone <zone-name>]
```

Options

<proxyuser-name>

Specifies the user name of the proxy user.

--zone <zone-name>

Specifies the access zone the proxy user authenticates through.

Examples

The following command displays the configuration details for the `hadoop-user23` proxy user in zone1:

```
isi hdfs proxyusers view hadoop-user23 --zone=zone1
```

The system displays output similar to the following example:

```
Name: hadoop-user23  
Members: krb_users  
         LOCAL  
         krb_user_004
```

isi hdfs racks create

Creates a new virtual HDFS rack.

Syntax

```
isi hdfs racks create <rack-name>
  [--client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--ip-pools <subnet>:<pool>]...
  [--zone <string>]
  [--verbose]
```

Options

<rack-name>

Specifies the name of the virtual HDFS rack. The rack name must begin with a forward slash—for example, **/example-name**.

--client-ip-ranges <low-ip-address>-<high-ip-address>...

Specifies IP address ranges of external Hadoop compute clients assigned to the virtual rack.

--ip-pools <subnet>:<pool>...

Assigns a pool of Isilon cluster IP addresses to the virtual rack.

--zone <string>

Specifies the access zone that will contain the virtual rack.

{--verbose | -v}

Displays more detailed information.

isi hdfs racks delete

Deletes a virtual HDFS rack.

Syntax

```
isi hdfs racks delete <rack-name>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<rack-name>

Deletes the specified virtual HDFS rack. Each rack name begins with a forward slash—for example, **/example-name**.

--zone <string>

Specifies the access zone that contains the virtual rack you want to delete.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

isi hdfs racks list

Lists the HDFS racks in an access zone.

Syntax

```
isi hdfs racks list
  [--zone <string>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the access zone. The system displays all virtual racks in the specified zone.

--format {table | json | csv | list}

Display HDFS racks in table, JSON, CSV, or list format.

{--no-header | -a}

Do not display headers in CSV or table output format.

{--no-footer | -z}

Do not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi hdfs racks modify

Modifies a virtual HDFS rack.

Syntax

```
isi hdfs racks modify <rack-name>
  [--name <rack-name>]
  [--client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--add-client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--remove-client-ip-ranges <low-ip-address>-<high-ip-address>]...
  [--clear-client-ip-ranges]
  [--ip-pools <subnet>:<pool>]...
  [--add-ip-pools <subnet>:<pool>]...
  [--remove-ip-pools <subnet>:<pool>]...
  [--clear-ip-pools]
  [--zone <string>]
  [--verbose]
```

Options

<rack-name>

Specifies the virtual HDFS rack to be modified. Each rack name begins with a forward slash—for example **/example-name**.

--name <rack-name>

Assigns a new name to the specified virtual rack. The rack name must begin with a forward slash—for example **/example-name**.

--client-ip-ranges <low-ip-address>-<high-ip-address>...

Specifies IP address ranges of external Hadoop compute clients assigned to the virtual rack. The value assigned through this option overwrites any existing IP address ranges. You can add a new range through the `--add-client-ip-ranges` option.

--add-client-ip-ranges <low-ip-address>-<high-ip-address>...

Adds a specified IP address range of external Hadoop compute clients to the virtual rack.

--remove-client-ip-ranges <low-ip-address>-<high-ip-address>...

Removes a specified IP address range of external Hadoop compute clients from the virtual rack. You can only remove an entire range; you cannot delete a subset of a range.

--clear-client-ip-ranges

Removes all IP address ranges of external Hadoop compute clients from the virtual rack.

--ip-pools <subnet>:<pool>...

Assigns pools of Isilon node IP addresses to the virtual rack. The value assigned through this option overwrites any existing IP address pools. You can add a new pool through the `--add-ip-pools` option.

--add-ip-pools <subnet>:<pool>...

Adds a specified pool of Isilon cluster IP addresses to the virtual rack.

--remove-ip-pools <subnet>:<pool>...

Removes a specified pool of Isilon cluster IP addresses from the virtual rack.

--clear-ip-pools

Removes all pools of Isilon cluster IP addresses from the virtual rack.

--zone <string>

Specifies the access zone that contains the virtual rack you want to modify.

{**--verbose** | **-v**}

Displays more detailed information.

isi hdfs racks view

Displays information for a specific virtual HDFS rack.

Syntax

```
isi hdfs racks view <rack-name>
  [--zone <string>]
```

Options

<rack-name>

Specifies the name of the virtual HDFS rack to view. Each rack name begins with a forward slash—for example, `/example-name`.

--zone <string>

Specifies the access zone that contains the virtual rack you want to view.

isi hdfs ranger-plugin settings modify

Modify Apache Ranger plug-in settings for HDFS.

Syntax

```
isi hdfs ranger-plugin settings modify
  [--enabled <boolean>]
```

```
[--policy-manager-url <string>]
[--repository-name <string>]
[--zone <string>]
[--verbose]
```

Options

--enabled <boolean>

Enable the HDFS Ranger plug-in.

--policy-manager-url <string>

The scheme, host name, and port of the Apache Ranger server (for example, `http://ranger.com:6080`).

--repository-name <string>

The HDFS repository name hosted on the Apache Ranger server.

--zone <string>

The access zone containing the HDFS repository.

{--verbose | -v}

Display more detailed information.

isi hdfs ranger-plugin settings view

View Apache Ranger plug-in settings for HDFS.

Syntax

```
isi hdfs ranger-plugin settings view
  [--zone <string>]
```

Options

--zone <string>

The access zone containing the HDFS repository.

isi hdfs settings modify

Modifies the HDFS settings for an access zone.

Syntax

```
isi hdfs settings modify
  [--service {yes | no}]
  [--default-block-size <size>]
  [--default-checksum-type {none | crc32 | crc32c}]
  [--authentication-mode {all | simple_only | kerberos_only}]
  [--root-directory <path>]
  [--webhdfs-enabled {yes | no}]
  [--ambari-server <string>]
  [--ambari-namenode <string>]
  [--ambari-metrics-collector <string>]
  [--odp-version <string>]
  [--data-transfer-cipher {none | aes_128_ctr | aes_192_ctr | aes_256_ctr}]
  [--zone <string>]
  [--verbose]
```

Options

--service {yes | no}

Enables or disables the HDFS service in the specified access zone. The HDFS service is enabled by default.

--default-block-size <size>

The block size (in bytes) reported by the HDFS service. K, M, and G; for example, 64M, 512K, 1G, are valid suffixes. The default value is 128 MB.

--default-checksum-type {none | crc32 | crc32c}

The checksum type reported by the HDFS service. The default value is `none`.

--authentication-mode {all | simple_only | kerberos_only}

The authentication method used for HDFS connections through the specified access zone. The default value is `all`.

--root-directory <path>

Root path that contains HDFS data in the access zone that can be accessed by Hadoop compute client connections. The root directory must be within the access zone base directory.

--webhdfs-enabled {yes | no}

Enables or disables the WebHDFS in the specified access zone. WebHDFS is enabled by default.

--ambari-server <string>

The Ambari server that receives communication from an Ambari agent. The value must be a resolvable hostname, FQDN, IPv4 or IPv6 address.

--ambari-namenode <string>

A point of contact in the access zone that Hadoop services managed through the Ambari interface should connect through. The value must be a resolvable IPv4 address or a SmartConnect zone name.

--ambari-metrics-collector <string>

The host name for the metrics collector. The value must be a resolvable hostname, FQDN, IPv4 or IPv6 address.

--odp-version <string>

The version of the Open Data Platform (ODP) stack repository, including build number if one exists, installed by the Ambari server. This is required to support ODP upgrades on other systems that are part of the Hadoop cluster.

--data-transfer-cipher {none | aes_128_ctr | aes_192_ctr | aes_256_ctr}

The Advanced Encryption Standard (AES) cipher to use for wire encryption.

--zone <string>

The access zone to which the HDFS settings apply.

{--verbose | -v}

Display more detailed information.

isi hdfs settings view

Displays the HDFS settings in an access zone.

Syntax

```
isi hdfs settings view
  [--zone <string>]
```

Options

--zone <string>

Specifies the access zone. The system will display the HDFS settings for the specified zone.

isi http settings modify

Modifies HTTP global settings.

Syntax

```
isi http settings modify
  [--access control {yes | no}]
  [--basic-authentication {yes | no}]
  [--dav {yes | no}]
  [--enable-access-log {yes | no}]
  [--integrated-authentication {yes | no}]
  [--server-root <path>]
  [--service {enabled | disabled | redirect}]
  [--verbose]
```

Options

--access control {yes | no}


Enables access control authentication, which allows the Apache web server to perform access checks. Access control authentication requires at least one type of authentication to be enabled.

--basic-authentication {yes | no}

Enables HTTP basic authentication. User credentials are sent in plain text format.

--dav {yes | no}

Enables multiple users to manage and modify files collaboratively across web servers.

 **NOTE: All DAV clients must go through a single node. DAV compliance is not met if you go through SmartConnect, or through two or more node IP addresses.**

--enable-access-log {yes | no}

Enables writing to a log when the HTTP server is accessed.

--integrated-authentication {yes | no}

Enables integrated authentication via NTLM, Kerberos, or both.

--server-root <path>

Specifies the root document directory. This must be a valid directory path within `/ifs`.

--service {enabled | disabled | redirect}

Enables or disables the HTTP service, or redirects to the OneFS web UI.

{--verbose | -v}

Displays more detailed information.

isi http settings view

Displays HTTP global settings.

Syntax

```
isi http settings view
```

Options

There are no options for this command.

Example

The following example shows the output generated by this command:

```
Access Control: No
Basic Authentication: No
    Dav: No
    Enable Access Log: Yes
Integrated Authentication: No
    Server Root: /ifs
    Service: redirect
```

isi job events list

Lists recent job events.

Syntax

```
isi job events list
  [--job-type <string>]
  [--job-id <integer>]
  [--begin <timestamp>]
  [--end <timestamp>]
  [--state {failed | running | cancelled_user | succeeded | paused_user | unknown |
  paused_priority | cancelled_system | paused_policy | paused_system}]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--job-type <string>

Displays all events of all instances of a specific job type (for example, SmartPools).

--job-id <integer>

Displays all events of a specific job instance.

--begin <timestamp>

Specifies the beginning of the time period for which job events should be listed. For example: **--begin "2013-09-17T00:00"**. This means that job events beginning at the first moment of September 17, 2013 should be listed.

--end <timestamp>

Specifies the end of the time period for job events to be listed. For example, **--end "2013-09-17T23:59"** means that job events right up to the last minute of September 17, 2013 should be listed.

--state {failed | running | cancelled_user | succeeded | paused_user | unknown | paused_priority | cancelled_system | paused_policy | paused_system}

Specifies that events of the given state or states should be listed.

{--limit | -1} <integer>

Displays no more than the specified number of job events. If no **timestamp** parameters are specified, the most recent job events of the specified number are listed.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information about job events.

Examples

The following command lists all FSAnalyze events that happened in the month of September.

```
isi job events list --job-type fsanalyze --begin "2013-09-01" --end "2013-09-30"
```

The system displays output similar to the following example.

Time	Message
2013-09-16T22:00:21	FSAnalyze[7] Waiting
2013-09-16T22:00:23	FSAnalyze[7] Running
2013-09-16T22:00:25	FSAnalyze[7] Phase 1: begin scan
2013-09-16T22:01:45	FSAnalyze[7] Phase 1: end scan
2013-09-16T22:01:46	FSAnalyze[7] Phase 2: begin merge
2013-09-16T22:02:30	FSAnalyze[7] Phase 2: end merge
2013-09-16T22:02:31	FSAnalyze[7] Succeeded
2013-09-17T22:00:05	FSAnalyze[9] Waiting
2013-09-17T22:00:08	FSAnalyze[9] Running
2013-09-17T22:00:11	FSAnalyze[9] Phase 1: begin scan
2013-09-17T22:01:37	FSAnalyze[9] Phase 1: end scan
2013-09-17T22:01:38	FSAnalyze[9] Phase 2: begin merge
2013-09-17T22:02:24	FSAnalyze[9] Phase 2: end merge
2013-09-17T22:02:26	FSAnalyze[9] Succeeded

Total: 14

The following command lists all the job events that happened on a specific day.

```
isi job events list --begin "2013-09-17T00:00" --end "2013-09-17T23:59"
```

The system displays output similar to the following example:

Time	Message
2013-09-17T22:00:04	SmartPools[8] Waiting
2013-09-17T22:00:05	FSAnalyze[9] Waiting
2013-09-17T22:00:06	SmartPools[8] Running
2013-09-17T22:00:07	SmartPools[8] Phase 1: begin lin policy update
2013-09-17T22:00:08	FSAnalyze[9] Running
2013-09-17T22:00:11	FSAnalyze[9] Phase 1: begin scan
2013-09-17T22:01:01	SmartPools[8] Phase 1: end lin policy update
2013-09-17T22:01:03	SmartPools[8] Phase 2: begin sin policy update
2013-09-17T22:01:06	SmartPools[8] Phase 2: end sin policy update
2013-09-17T22:01:09	SmartPools[8] Succeeded
2013-09-17T22:01:37	FSAnalyze[9] Phase 1: end scan
2013-09-17T22:01:38	FSAnalyze[9] Phase 2: begin merge
2013-09-17T22:02:24	FSAnalyze[9] Phase 2: end merge
2013-09-17T22:02:26	FSAnalyze[9] Succeeded

Total: 14

isi job jobs cancel

Cancels an active job.

Syntax

```
isi job jobs cancel <job>
```

Options

<job>

Specifies the job to cancel. You can specify the job by job ID or job type. Specify a job type only if one instance of that job type is active.

Examples

The following command cancels an active MultiScan job.

```
isi job jobs cancel multiscan
```

The following command cancels an active job with an instance ID of 14.

```
isi job jobs cancel 14
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job cancel 14
```

isi job jobs list

Displays information about active jobs.

Syntax

```
isi job jobs list
  [--state {running | paused_user | paused_priority | paused_policy | paused_system}]
  [--limit <integer>]
  [--sort {id | type | state | impact | policy | priority | start_time | running_time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--state {running | paused_user | paused_priority | paused_policy | paused_system}

Controls which jobs are listed according to status.

{--limit | -1} <integer>

Displays no more than the specified number of items. If no other parameters are specified, displays the most recently activated jobs up to the specified number.

--sort {id | type | state | impact | policy | priority | start_time | running_time}

Sorts the output by the specified attribute.

--descending

Sorts the output in descending order of activation time.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header}

Displays table and CSV output without headers.

{--no-footer}

Displays table output without footers.

{--verbose}

Displays more detailed information about active jobs.

Examples

The following example lists jobs that have been manually paused.

```
isi job jobs list --state paused_user
```

The system displays output similar to the following example.

ID	Type	State	Impact	Pri	Phase	Running Time
12	Collect	Paused by user	Low	4	1/2	11s
23	SmartPools	Paused by user	Low	6	1/8	40s

Total: 2

The following example outputs a CSV-formatted list of jobs to a file in the `/ifs/data` path.

```
isi job jobs list --format csv > /ifs/data/joblist.csv
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job list --format csv > /ifs/data/joblist.csv
```

isi job jobs modify

Changes the priority level or impact policy of a queued, running, or paused job.

Syntax

```
isi job jobs modify <job>  
  [--priority <integer>]  
  [--policy <string>]
```

Options

<job>

Specifies the job ID or job type to modify. If you specify job type (for example, FlexProtect), only one instance of that type can be active.

{--priority | -p} <integer>

Sets the priority level for the specified job.

```
{--policy | -o} <string>
```

Sets the impact policy for the specified job.

Examples

The following command changes the impact policy of an active MultiScan job. This command example, which specifies the job type, works only when a single instance of MultiScan is active.

```
isi job jobs modify multiscan --policy high
```

If more than one instance of a job type is active, you can specify the job ID number instead of job type. The following command changes the priority of an active job with an ID of 7.

```
isi job jobs modify 7 --priority 2
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job modify 7 --priority 2
```

isi job jobs pause

Pauses an active job.

Syntax

```
isi job jobs pause <job>
```

Options

<job>

Specifies the job to pause. You can specify the job by job type or job ID. If you use job type, only one instance of the job type can be active.

Examples

The following command pauses an active AutoBalance job.

```
isi job jobs pause autobalance
```

The following command pauses an active job with an ID of 18.

```
isi job jobs pause 18
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job pause 18
```

To resume a paused job, use the `isi job resume` command.

isi job jobs resume

Resumes a paused job.

You can confirm that a job has resumed by using the `isi job jobs list` command. Actual resumption of the job can take a while, depending on other activity in the Job Engine queue.

Syntax

```
isi job jobs resume <job>
```

Options

<job>

Specifies the job to resume. You can specify the job by job type or job ID. If you use the job type parameter, only one instance of this job type can be in the Job Engine queue.

Examples

The following command resumes a paused AutoBalance job.

```
isi job jobs resume autobalance
```

The following command resumes a paused job with an ID of 16.

```
isi job jobs resume 16
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job resume 16
```

isi job jobs start

Starts a new job.

The `isi job jobs start` command does not control jobs that are already in progress. If an active job is paused, you can use the `isi job jobs resume` command to start it from the point it was paused.

Syntax

```
isi job jobs start <type>
  [--policy <string>]
  [--priority <integer>]
  [--no-dup]
  [--paths <path>]
  [--delete]
  [--delete-quotas]
  [--root <path>]
  [--dm-type {snaprevert | synciq}]
  [--mapping-type {clone | sid | unix | native}]
  [--mode {clone | inherit | convert}]
  [--template <path>]
  [--zone <string>]
  [--snapid <integer>]
  [--verbose]
```

Options

<type>

Specifies the type of job to add to the job queue (for example, MediaScan).

{--priority} <integer>

Sets the priority level for the specified job, with 1 being the highest priority and 10 being the lowest.

{--policy} <string>

Sets the impact policy for the specified job.

{--no-dup}

Disallows duplicate jobs. If an instance of the specified job is already in the queue, the new job does not start.

--paths <path>

Specifies the path of the job, which must be within `/ifs`. This option is valid only for the TreeDelete and PermissionRepair jobs.

--delete

Valid for the DomainMark job only. Deletes the domain mark.

--delete-quotas

Valid for the TreeDelete job only. Automatically deletes quotas on the removed dataset.

This option may generate the following log message while attempting to remove a directory with a quota on it.

```
Attempt to remove directory <lin> with quota domains still defined on it,
operation not permitted.
```

Ignore this log message. The `--delete-quotas` option deletes the quota first and then it removes the directory.

--root <path>

Valid for the DomainMark job only. Specifies the root path location for the DomainMark job.

--dm-type {snaprevert | synciq}

Valid for the DomainMark job only. Specifies the domain type for the DomainMark job.

--mapping-type {global | sid | unix | native}

Valid for the PermissionRepair job only, and is only used with the `--mode convert` option. Specifies the type for PermissionRepair.

--mode {clone | inherit | convert}

Valid for the PermissionRepair job only. Specifies the mode for PermissionRepair.

--template <path>

Valid for the PermissionRepair job only. Specifies the pathname of a template file to use as a model for the PermissionRepair job. Must be within the `/ifs` path.

--zone <string>

Valid for the PermissionRepair job only. Specifies the access zone for PermissionRepair.

--snapid <integer>

Valid for the SnapRevert job only. Specifies a snapshot ID for the SnapRevert job.

{--verbose | -v}

Displays more detailed information.

Examples

The following command starts an AutoBalance job.

```
isi job jobs start autobalance
```

The following command starts a MultiScan job with a priority of 8 and a high impact policy.

```
isi job jobs start multiscan --priority 8 --policy high
```

The following command starts a TreeDelete job with a priority of 10 and a low impact policy that deletes the /ifs/data/old directory.

```
isi job jobs start treedelete --path /ifs/data/old --priority 10 --policy low
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job start autobalance
```

isi job jobs view

Displays information about a running or queued job, including the state, impact policy, priority, and schedule.

Syntax

```
isi job jobs view <job>
```

Options

<job>

Specifies the job to view. You can specify the job by job type or job ID. If you specify a job type, only one instance of this job can be active.

Examples

The following command displays information about an AutoBalance job with a job ID of 15.

```
isi job jobs view 15
```

The system displays information similar to the following example.

```
          ID: 15
          Type: AutoBalance
          State: Paused by user
          Impact: Low
          Policy: LOW
          Pri: 4
          Phase: 1/5
          Start Time: 2013-09-19T09:08:28
          Running Time: 24s
          Participants: 1, 2, 3
          Progress: Drives: 6 done, 0 in progress; last updated 3:0; Processed 4624 LINS and
918669 KB; 0 ECCs and 0 errors
          Waiting on job ID: -
          Description:
```

In all instructions that include the `isi job jobs` command, you can omit the `jobs` entry.

```
isi job view 15
```

isi job policies create

Creates a custom job impact policy.

By default, the new impact policy is assigned a low impact level. You can specify multiple time periods (intervals) during which the job can run at higher impact levels or be paused.

Syntax

```
isi job policies create <name>
  [--description <string>]
  [--impact {Low | Medium | High | Paused }]
  [--begin <interval_time>]
  [--end <interval_time>]
```

Options

<name>

Specifies a name for the new impact policy. The following names are reserved and cannot be used: LOW, MEDIUM, HIGH, and OFF_HOURS.

--description <string>

Describes the job policy.

--impact {Low | Medium High | Paused}

Specifies an impact level for the policy: Low, Medium, High, or Paused. You can specify an `--impact` parameter for each impact interval that you define.

--begin <interval_time>

Specifies the beginning time, on a 24-hour clock, of the period during which a job can run. For example: `--begin "Friday 20:00"`.

--end <interval_time>

Specifies the ending time, on a 24-hour clock, of the period during which a job can run. For example: `--end "Sunday 11:59"`.

Examples

The following command creates a new impact policy named HIGH-WKEND.

```
isi job policies create HIGH-WKEND --impact high --begin "Saturday 00:01" --end "Sunday 23:59"
```

The following command creates a more complex impact policy named HI-MED-WKEND. This policy includes multiple impact levels and time intervals. At the end of the specified intervals, a job running with this policy would automatically return to LOW impact.

```
isi job policies create HI-MED-WKEND --description "High to medium impact when run on the weekend" --impact high --begin "Friday 20:00" --end "Monday 03:00" --impact medium --begin "Monday 03:01" --end "Monday 08:00"
```

isi job policies delete

Deletes a job impact policy.

The following policies are reserved and cannot be deleted: LOW, MEDIUM, HIGH, and OFF_HOURS.

Syntax

```
isi job policies delete <id>
[--force]
```

Options

<id>

Specifies the name of the impact policy to delete. If you are unsure of the name, you can use the `isi job policies list` command.

--force

Forces deletion of the impact policy without the system asking for confirmation.

Examples

The following command deletes a custom impact policy named HIGH-MED.

```
isi job policies delete HIGH-MED
```

When you press ENTER, OneFS displays a confirmation message: Are you sure you want to delete the policy HIGH-MED? (yes/[no]):

Type **yes**, and then press ENTER.

The following command deletes a custom impact policy named HIGH-WKEND without the confirmation message being displayed.

```
isi job policies delete HIGH-WKEND --force
```

isi job policies list

Displays the names and descriptions of job impact policies.

Syntax

```
isi job policies list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

Examples

The following command displays a list of available impact policies.

```
isi job policies list
```

The system displays output similar to the following example:

```
ID          Description
-----
HIGH       Isilon template: high impact at all times
LOW        Isilon template: low impact at all times
MEDIUM     Isilon template: medium impact at all times
OFF_HOURS  Isilon template: paused M-F 9-5, low-impact
           at other times
HI-MED     High to medium to low
HI-WKEND   High impact when run on weekends
MED-WKEND  Medium impact when run on weekends
-----
Total: 7
```

The following command displays more information about available policies.

```
isi job policies list --verbose
```

The system displays verbose output in a list format as shown in the following partial example:

```
          ID: HIGH
Description: Isilon template: high impact at all times
System: True
Impact Intervals
          Impact : High
          Begin  : Sunday 00:00
          End    : Sunday 00:00
-----
          ID: LOW
Description: Isilon template: low impact at all times
System: True
Impact Intervals
          Impact : Low
          Begin  : Sunday 00:00
          End    : Sunday 00:00
-----
```

isi job policies modify

Change the description, impact levels and time intervals of a custom impact policy.

To confirm that the custom policy reflects your changes, you can use the `isi job policy view` command.

Syntax

```
isi job policy modify <ID>
  [--description<string>]
  [--impact {Low | Medium | High | Paused}]
  [--begin <interval_time>]
  [--end <interval_time>]
  [--reset_intervals]
```

Options

<ID>

Specifies the name of the policy to modify.

--description <string>

Specifies a description for the policy. Replaces an older description if one was in place.

--impact {**Low** | **Medium** **High** | **Paused**}

Specifies an impact level for the policy: **Low**, **Medium**, **High**, or **Paused**. Specify an **--impact** parameter for each additional impact interval that you define.

--begin <interval_time>

Specifies the beginning time, on a 24-hour clock, of the period during which a job can run. For example: **--begin "Friday 20:00"**.

--end <interval_time>

Specifies the ending time, on a 24-hour clock, of the period during which a job can run. For example: **--end "Sunday 11:59"**.

--reset-intervals

Clears all job policy intervals and restores the defaults.

Examples

The following command clears the custom intervals from a custom policy named MY_POLICY as the first step to adding new intervals.

```
isi job policies modify MY_POLICY --reset-intervals
```

The following command adds new intervals to a custom policy.

```
isi job policies modify MY_POLICY --impact high --begin "Friday 20:00" --end "Sunday 11:59"
```

isi job policies view

Displays the details for a specific Job Engine job policy.

Syntax

```
isi job policies view  
[<id> <string>]
```

Options

<id> <string>

Specifies the job policy to display by policy ID.

Examples

The following command displays the details for the default job policy, HIGH.

```
isi job policies view HIGH
```

The system displays the following policy details:

```
ID: HIGH
```

```
Description: Isilon template: high impact at all times
System: True
Impact Intervals
    Impact : High
    Begin : Sunday 00:00
    End : Sunday 00:00
```

isi job reports list

Displays information about successful job operations, including date and time, job ID, job type, and job phases that fully completed.

Syntax

```
isi job reports list
  [--job-type <string>]
  [--job-id <integer>]
  [--begin] <timestamp>]
  [--end] <timestamp>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

--job-type <string>

Displays reports for all instances of the specified job type.

--job-id <integer>

Displays the report for a job with the specified job ID. If a job has multiple phases, Job Engines displays a report for each phase of the specified job ID.

{--begin | -b} <interval_time>

Specifies the beginning of the time period for the job reports list. For example: **--begin "2013-09-19"**.

{--end | -e} <interval_time>

Specifies the end of the time period for the job reports list. For example: **--end "2013-09-20"**.

{--limit | -l} <integer>

Displays no more than the specified number of reports.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

Examples

The following command displays reports for all MultiScan jobs within a specified time period.

```
isi job reports list --job-type multiscan --begin "2013-9-19" --end "2013-9-20"
```

The system displays output similar to the following example.

Time	Job ID	Job Type	Phase
2013-09-19T10:00:08	1	MultiScan	1
2013-09-19T10:00:20	1	MultiScan	2
2013-09-19T10:00:21	1	MultiScan	3
2013-09-19T10:00:34	1	MultiScan	4
2013-09-19T10:02:50	2	MultiScan	1
2013-09-19T10:03:06	2	MultiScan	2
2013-09-19T10:03:09	2	MultiScan	3
2013-09-19T10:03:12	2	MultiScan	4
2013-09-20T10:04:53	4	MultiScan	1
2013-09-20T10:05:11	4	MultiScan	2
2013-09-20T10:05:15	4	MultiScan	3
2013-09-20T10:05:20	4	MultiScan	4

Total: 12

isi job reports view

Displays a detailed report for a specific job. Reports can be displayed only for successful jobs or for successful phases of a job.

Syntax

```
isi job reports view <id>
```

Options

<id>

Specifies the job ID for the reports you want to view.

Examples

The following command requests reports for an FSAnalyze job with an ID of 7.

```
isi job reports view 7
```

The system displays output similar to the following example. Note that when a job has more than one phase, a report for each phase is provided.

```
FSAnalyze[7] phase 1 (2013-09-19T22:01:58)
```

```
-----
FSA JOB QUERY PHASE
Elapsed time:           83 seconds
LINS traversed:        433
Errors:                 0
CPU usage:              max 30% (dev 2), min 0% (dev 1), avg 10%
Virtual memory size:   max 111772K (dev 1), min 104444K (dev 2), avg 109423K
Resident memory size: max 14348K (dev 1), min 9804K (dev 3), avg 12706K
Read:                  9 ops, 73728 bytes (0.1M)
Write:                 3035 ops, 24517120 bytes (23.4M)
```

```
FSAnalyze[7] phase 2 (2013-09-19T22:02:47)
```

```
-----
FSA JOB MERGE PHASE
Elapsed time:           47 seconds
Errors:                 0
CPU usage:              max 33% (dev 1), min 0% (dev 1), avg 8%
Virtual memory size:   max 113052K (dev 1), min 110748K (dev 2), avg 111558K
Resident memory size: max 16412K (dev 1), min 13424K (dev 3), avg 14268K
Read:                  2 ops, 16384 bytes (0.0M)
Write:                 2157 ops, 16871424 bytes (16.1M)
```

isi job statistics view

Displays statistics for an active job or jobs on an entire cluster or a specific node.

Syntax

```
isi job statistics view
[--job-id <integer>]
[--devid <integer>]
[--verbose]
[--format {table | json | csv | list}]
```

Options

--job-id <integer>

Displays statistics for a specific job ID.

--devid <integer>

Displays statistics for a specific node (device) in the cluster.

{--verbose | -v}

Displays more detailed statistics for an active job or jobs.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

Examples

The following command requests statistics for an AutoBalance job with an ID of 6.

```
isi job statistics view --job-id 6
```

The system displays output similar to the following example. In the example, PID is the process ID, and CPU indicates CPU utilization by the job. Also indicated are how many worker threads exist for the job on each node and what the sleep-to-work (STW) ratio is for each thread. The statistics represent how the system throttles the job based on impact policies.

```
Job ID: 6
Phase: 2
Nodes
  Node : 1
    PID : 17006
    CPU : 0.00% (0.00% min, 7.91% max, 4.50% avg)
  Memory
    Virtual : 104.62M (104.37M min, 104.62M max, 104.59M avg)
    Physical : 10.08M (10.01M min, 10.11M max, 10.09M avg)
  I/O
    Read : 4141 ops, 45.33M
    Write : 5035 ops, 35.28M
  Workers : 2 (0.60 STW avg.)
  Node : 2
    PID : 16352
    CPU : 13.96% (1.95% min, 13.96% max, 9.61% avg)
  Memory
    Virtual : 104.62M (104.37M min, 104.62M max, 104.59M avg)
    Physical : 10.01M (9.90M min, 10.01M max, 10.00M avg)
  I/O
    Read : 3925 ops, 43.39M
    Write : 4890 ops, 34.13M
  Workers : 2 (0.60 STW avg.)
  Node : 3
    PID : 15929
```



```
CPU : 0.98% (0.98% min, 12.89% max, 6.82% avg)
Memory
Virtual : 104.62M (104.37M min, 104.62M max, 104.57M avg)
Physical : 9.86M (9.84M min, 9.94M max, 9.92M avg)
I/O
Read : 3354 ops, 36.77M
Write : 772 ops, 2.12M
Workers : 2 (0.60 STW avg.)
```

isi job status

Displays a summary of active, completed, and failed jobs.

Syntax

```
isi job status
[--verbose]
```

Options

{**--verbose** | **-v**}

Displays more detailed job status information, including information about the cluster and nodes.

Examples

The following command provides basic job status.

```
isi job status
```

The system displays output that is similar to the following example.

The job engine is running.

No running or queued jobs.

Recent finished jobs:

ID	Type	State	Time
1	MultiScan	System Cancelled	2013-09-24T08:23:44
3	MultiScan	Succeeded	2013-09-24T08:26:37
2	SetProtectPlus	Succeeded	2013-09-24T08:27:16
4	FlexProtect	Succeeded	2013-09-24T09:14:27

Total: 4

The following command provides more detailed job status information.

```
isi job status --verbose
```

The system displays additional output that includes cluster and node information.

The job engine is running.

```
Coordinator: 1
Connected: True
Disconnected Nodes: -
Down or Read-Only Nodes: False
Statistics Ready: True
Cluster Is Degraded: False
Run Jobs When Degraded: False
```

No running or queued jobs.

Recent finished jobs:

ID	Type	State	Time
1	MultiScan	System Cancelled	2013-09-24T08:23:44
3	MultiScan	Succeeded	2013-09-24T08:26:37
2	SetProtectPlus	Succeeded	2013-09-24T08:27:16
4	FlexProtect	Succeeded	2013-09-24T09:14:27

Total: 4

isi job types list

Displays a list of job types and default settings.

Syntax

```
isi job types list
  [--all]
  [--sort {id | policy | exclusion_set | priority}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--all

Displays all job types available in the Job Engine.

--sort {id | policy | exclusion_set | priority}

Sorts the output by the specified parameter.

--descending

In conjunction with **--sort** option, specifies that output be sorted descending order. By default, output is sorted in ascending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated values (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information about a specific job type or all job types.

Examples

The following command provides detailed information about job types.

```
isi job types list --sort id --verbose
```

The system displays output similar to the following example.

```
      ID: AVScan
Description: Perform an antivirus scan on all files.
      Enabled: Yes
```

```
Policy: LOW
Schedule:
Exclusion Set: None
Priority: 6
```

```
-----
ID: AutoBalance
Description: Balance free space in a cluster. AutoBalance is most
             efficient in clusters that contain only HDDs.
Enabled: Yes
Policy: LOW
Schedule:
Exclusion Set: Restripe
Priority: 4
```

```
-----
ID: AutoBalanceLin
Description: Balance free space in a cluster. AutoBalanceLin is
             most efficient if file system metadata is stored on
             SSDs.
Enabled: Yes
Policy: LOW
Schedule:
Exclusion Set: Restripe
Priority: 4
```

isi job types modify

Modifies the parameters of a specified job type.

You can view the current parameters of any job type by using the `isi job types view` command.

Syntax

```
isi job types modify <id>
  [--enabled <boolean>]
  [--policy <string>]
  [--schedule<string>]
  [--priority <integer>]
  [--clear-schedule]
```

Options

<id>

Specifies the job type to modify.

--enabled <boolean>

Specifies whether the job type is enabled or disabled.

--policy<string>

Sets the policy for the specified job type.

--schedule <string>

Sets a recurring date pattern to run the specified job type.

--priority<integer>

Sets the priority level for the specified job type. Job types have a priority value between 1 and 10, with 1 being the highest priority and 10 being the lowest.

--clear-schedule

Clears any schedule associated with the specified job type.

--force

Forces the modification without a confirmation message.

Examples

The following command adds a recurring schedule to the MultiScan command.

```
isi job types modify multiscan --schedule "Every Friday at 22:00"
```

When you run this command, the system prompts you to confirm the change. Type **yes** or **no**, and then press ENTER.

isi job types view

Displays the parameters of a specific job type, including the description, schedule, policy, priority, and whether the job type is a member of an exclusion set.

Syntax

```
isi job types view <id>
```

Options

<id>

Specifies the job type to view.

Examples

The following command displays the parameters of the job type MultiScan.

```
isi job types view multiscan
```

The system displays output similar to the following example.

```
      ID: MultiScan
Description: Perform the work of the AutoBalance and Collect jobs simultaneously.
  Enabled: Yes
   Policy: LOW
  Schedule:
Exclusion Set: Restripe, Mark
  Priority: 4
```

isi license add

Activates a licensable product using a new or updated ELMS License file.

Syntax

```
isi license add
  [--path <string>]
  [--evaluation <string>]
  [--verbose]
```

Options

--path <string>

The location of the license file on the cluster.

--evaluation <string>

The name of a license to activate for a limited evaluation period. Repeat this option for every license you want to activate for evaluation.

{--verbose | -v}

Displays more detailed information.

isi license generate

Generates a license activation file.

Syntax

```
isi license generate
  [--include <module>]
  [--exclude <module>]
  [--only <module>...]
  [--action (license_list_only | generate_activation)]
  [--file <path>]
  [--format (table | json | csv | list)]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--include <module>

Adds a software module license to the activation file. Specify `--include` for each license you want to include in the activation file.

--exclude <module>

Removes a software module license from the activation file. Specify `--exclude` for each license you want to remove from the activation file.

--only <module>

Adds a software module license to the activation file. Specify `--only` for each license you want to include in the activation file.

--action (license_list_only | generate_activation)

Specifies the action you want the command to take. You can generate an activation file, or you can return a list of activated licenses without generating an activation file.

--file <path>

Sets the location on the cluster where you want to save the new activation file.

--format {table | json | csv | list}

Display licenses in table, JSON, CSV, or list format.

{--no-header | -a}

Do not display headers in table or CSV formats.

{--no-footer | -z}

Do not display table summary footer information.

{verbose | -v}

Displays more detailed information.

isi license list

Retrieves license information for all licensable products.

Syntax

```
isi license list
  [--limit <integer>]
  [--sort {name | module | status | expiration}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

The number of licenses to display.

--sort {name | module | status | expiration}

Sort data by the specified field.

{--descending | -d}

Sort data in descending order.

--format {table | json | csv | list}

Display licenses in table, JSON, CSV, or list format.

{--no-header | -a}

Do not display headers in table or CSV formats.

{--no-footer | -z}

Do not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi license view

Retrieves license information for any licensable product.

Syntax

```
isi license view <name>
  [--format {list | json}]
```

Options

<name>

Product name for the license to view.

{--format | -f} {list | json}

Display licenses in list or JSON format.

OneFS displays the license name, status and expiration for the license.

OneFS isi commands N through R

This chapter contains documentation of the OneFS CLI commands `isi ndmp contexts delete` through `isi remotesupport connectemc view`.

Topics:

- `isi ndmp contexts delete`
- `isi ndmp contexts list`
- `isi ndmp contexts view`
- `isi ndmp dumpdates delete`
- `isi ndmp dumpdates list`
- `isi ndmp sessions delete`
- `isi ndmp sessions list`
- `isi ndmp sessions view`
- `isi ndmp settings diagnostics modify`
- `isi ndmp settings diagnostics view`
- `isi ndmp settings global modify`
- `isi ndmp settings global view`
- `isi ndmp settings preferred-ips create`
- `isi ndmp settings preferred-ips delete`
- `isi ndmp settings preferred-ips list`
- `isi ndmp settings preferred-ips modify`
- `isi ndmp settings preferred-ips view`
- `isi ndmp settings variables create`
- `isi ndmp settings variables delete`
- `isi ndmp settings variables list`
- `isi ndmp settings variables modify`
- `isi ndmp users create`
- `isi ndmp users delete`
- `isi ndmp users list`
- `isi ndmp users modify`
- `isi ndmp users view`
- `isi network dnscache flush`
- `isi network dnscache modify`
- `isi network dnscache view`
- `isi network external modify`
- `isi network external view`
- `isi network groupnets create`
- `isi network groupnets delete`
- `isi network groupnets list`
- `isi network groupnets modify`
- `isi network groupnets view`
- `isi network interfaces list`
- `isi network pools create`
- `isi network pools delete`
- `isi network pools list`
- `isi network pools modify`
- `isi network pools rebalance-ips`
- `isi network pools sc-resume-nodes`
- `isi network pools sc-suspend-nodes`
- `isi network pools view`

- isi network rules create
- isi network rules delete
- isi network rules list
- isi network rules modify
- isi network rules view
- isi network sc-rebalance-all
- isi network subnets create
- isi network subnets delete
- isi network subnets list
- isi network subnets modify
- isi network subnets view
- isi nfs aliases create
- isi nfs aliases delete
- isi nfs aliases list
- isi nfs aliases modify
- isi nfs aliases view
- isi nfs exports check
- isi nfs exports create
- isi nfs exports delete
- isi nfs exports list
- isi nfs exports modify
- isi nfs exports reload
- isi nfs exports view
- isi nfs log-level modify
- isi nfs log-level view
- isi nfs netgroup check
- isi nfs netgroup flush
- isi nfs netgroup modify
- isi nfs nlm locks list
- isi nfs nlm locks waiters
- isi nfs nlm sessions check
- isi nfs nlm sessions delete
- isi nfs nlm sessions list
- isi nfs nlm sessions refresh
- isi nfs nlm sessions view
- isi nfs settings export modify
- isi nfs settings export view
- isi nfs settings global modify
- isi nfs settings global view
- isi nfs settings zone modify
- isi nfs settings zone view
- isi ntp servers create
- isi ntp servers delete
- isi ntp servers list
- isi ntp servers modify
- isi ntp servers view
- isi ntp settings modify
- isi ntp settings view
- isi performance datasets create
- isi performance datasets delete
- isi performance datasets list
- isi performance datasets modify
- isi performance datasets view
- isi performance filters apply
- isi performance filters list
- isi performance filters modify
- isi performance filters remove

- [isi performance filters view](#)
- [isi performance metrics list](#)
- [isi performance metrics view](#)
- [isi performance settings modify](#)
- [isi performance settings view](#)
- [isi performance workloads list](#)
- [isi performance workloads modify](#)
- [isi performance workloads pin](#)
- [isi performance workloads unpin](#)
- [isi performance workloads view](#)
- [isi_phone_home](#)
- [isi quota quotas create](#)
- [isi quota quotas delete](#)
- [isi quota quotas list](#)
- [isi quota quotas modify](#)
- [isi quota quotas notifications clear](#)
- [isi quota quotas notifications create](#)
- [isi quota quotas notifications delete](#)
- [isi quota quotas notifications disable](#)
- [isi quota quotas notifications list](#)
- [isi quota quotas notifications modify](#)
- [isi quota quotas notifications view](#)
- [isi quota quotas view](#)
- [isi quota reports create](#)
- [isi quota reports delete](#)
- [isi quota reports list](#)
- [isi quota settings mappings create](#)
- [isi quota settings mappings delete](#)
- [isi quota settings mappings list](#)
- [isi quota settings mappings modify](#)
- [isi quota settings mappings view](#)
- [isi quota settings notifications clear](#)
- [isi quota settings notifications create](#)
- [isi quota settings notifications delete](#)
- [isi quota settings notifications list](#)
- [isi quota settings notifications modify](#)
- [isi quota settings notifications view](#)
- [isi quota settings reports modify](#)
- [isi quota settings reports view](#)
- [isi readonly list](#)
- [isi readonly modify](#)
- [isi readonly view](#)
- [isi remotesupport connectemc modify](#)
- [isi remotesupport connectemc view](#)

isi ndmp contexts delete

Deletes an NDMP context.

Syntax

```
isi ndmp contexts delete --id <id>
  [--force]
  [--verbose]
```

Options

- `--id <id>`
The context ID string.
- `{--force | -f}`
Skips the confirmation prompt.
- `{verbose | -v}`
Displays more detailed information.

isi ndmp contexts list

Lists NDMP contexts.

Syntax

```
isi ndmp contexts list  
  [--type {bre | backup | restore}]  
  [--format {table | json | csv | list}]
```

Options

- `{--type | -t} {bre | backup | restore}`
Displays entries for the specified level: backup restartable extension (BRE), backup, or restore.
- `--format {table | json | csv | list}`
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

isi ndmp contexts view

Displays detailed information of an NDMP context.

Syntax

```
isi ndmp contexts view --id <id>  
  [--format {list | json}]
```

Options

- `--id <id>`
The context ID string.
- `--format {list | json}`
Lists the NDMP context in the specified format.

isi ndmp dumpdates delete

Deletes a snapshot created for a snapshot-based NDMP incremental backup.

Syntax

```
isi ndmp dumpdates delete --path <path>
  [--level <integer>]
  [--force]
  [--verbose]
```

Options

--path <path>

The path of the NDMP dumpdate. Must be within the `/ifs` directory structure.

--level <integer>

Deletes a dumpdate entry for a backup of the specified level for the given directory. If this option is not specified, deletes all dumpdate entries for the given directory.

Examples

The following command deletes the dumpdate entry for a level 0 backup of `/ifs/data/media`:

```
isi ndmp dumpdates delete /ifs/data/media --level=0
```

isi ndmp dumpdates list

Displays snapshots created for snapshot-based NDMP incremental backups.

Syntax

```
isi ndmp dumpdates list
  [--path <path>]
  [--level <integer>]
  [--limit <integer>]
  [--sort {path | level}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--path <path>

The path of the NDMP dumpdate. Must be within the `/ifs` directory structure.

--level <integer>

Displays dumpdate entries for a backup of the specified level for the given directory path.

{--limit | -1}<integer>

The number of NDMP dumpdates to display.

--sort {path | level}

Sorts data by the specified field.

{--descending | -d}

Sorts data in descending order.

--format {table | json | csv | list}

Displays NDMP dumpdates in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To view NDMP dumpdate entries, run the following command:

```
isi ndmp dumpdates list
```

The system displays output similar to the following example:

Date	Level	SnapID	Path
Fri May 29 12:06:26 2015	0	18028	/ifs/tmp/backup
Fri May 29 12:20:56 2015	1	18030	/ifs/tmp/backup

If a snapshot was created for a non-snapshot-based incremental backup, the snapshot ID is **0**.

isi ndmp sessions delete

Stops an NDMP session.

Syntax

```
isi ndmp sessions delete --session <session>
  [--force]
  [--level]
  [--path]
  [--session]
  [--verbose]
```

Options

{--force | -f}

Skips the confirmation prompt.

--level

Stops an NDMP session for a specified level.

--path

Stops an NDMP session that is running at a specified path.

--session <session>

The NDMP session identifier. The session ID consists of the logical node number (LNN) followed by a decimal point and then the process ID (PID), such as <lnn>.<pid>.

{verbose | -v}

Displays more detailed information.

Example for the `--session` option

The following command ends an NDMP session with the session ID 4.36339:

```
isi ndmp sessions delete --session=4.36339
```

isi ndmp sessions list

Lists all or specified NDMP sessions.

Syntax

```
isi ndmp sessions list
  [--node <integer>]
  [--session <string>]
  [--consolidate]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--node | -n} <integer>

Displays only sessions for the specified node.

{--session | -s} <string>

The NDMP session identifier. The session ID consists of the logical node number (LNN) followed by a decimal point and then the process ID (PID), such as <lnn>.<pid>.

{--consolidate | -c}

Consolidates sessions of a multi-stream context.

{--limit | -l} <integer>

The number of NDMP sessions to display.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{verbose | -v}

Displays more detailed information.

isi ndmp sessions view

Displays detailed information about an NDMP session.

Syntax

```
isi ndmp sessions view --session <session>
  [--probe]
  [--format {list | json}]
```

Options

--session <session>

The NDMP session identifier. The session ID consists of the logical node number (LNN) followed by a decimal point and then the process ID (PID), such as <lnn>.<pid>.

--probe

Displays probe information about the NDMP session.

--format {list | json}

Displays NDMP session information in list or JSON format.

isi ndmp settings diagnostics modify

Modifies NDMP diagnostics settings.

Syntax

```
isi ndmp settings diagnostics modify
  [--diag-level <integer>]
  [--protocol-version <integer>]
  [--trace-level {none | standard | include-file-history | log-file-history}]
```

Options

--diag-level <integer>

The diagnostics level for NDMP.

--protocol-version <integer>

The NDMP protocol version (3 or 4).

--trace-level {none | standard | include-file-history | log-file-history}

The NDMP trace log level. Select none for no log, standard for NDMP protocol tracing, include-file-history to log file history information into the trace file, or log-file-history to log file history into the file history log.

isi ndmp settings diagnostics view

Displays NDMP diagnostic settings.

Syntax

```
isi ndmp settings diagnostics view
  [--format {list | json}]
```

Options

--format {list | json}

Displays NDMP diagnostic settings information in list or JSON format.

isi ndmp settings global modify

Modifies NDMP global settings.

Syntax

```
isi ndmp settings global modify
  [--service {true | false}]
  [--dma {generic | atempo | bakbone | commvault | emc | symantec | tivoli | symantec-
netbackup | symantec-backupexec}]
  [--port <integer>]
  [--bre-max-num-contexts <integer>]
  [--msb-context-retention-duration <integer>]
  [--msr-context-retention-duration <integer>]
```

Options

{--service | -s} {true | false}

Enables or disables the NDMP service.

{--dma | -d} {generic | atempo | bakbone | commvault | emc | symantec | tivoli | symantec-netbackup | symantec-backupexec}

The data management application (DMA) that controls NDMP sessions.

{--port | -p} <integer>

Sets the TCP/IP port number on which the NDMP daemon listens for incoming connections. The default port is 10000.

--bre-max-num-contexts <integer>

Sets the maximum number of restartable backup contexts. The system maximum limit is 1024, and the default is 64. Set this option to zero (0) to disable restartable backups.

--msb-context-retention-duration <integer>

Sets the duration of multi-stream backup context retention. Express durations in YMWDHms integer format. The default duration is 5m (five minutes).

--msr-context-retention-duration <integer>

Sets the duration of multi-stream restore context retention. Express durations in YMWDHms integer format. The default duration is 10m (ten minutes).

--stub-file-open-timeout

During a backup or a restore, a smartlinked file may not be opened due to active operations. NDMP will retry, but will fail after this timeout. The default value is 15 (15 seconds). The maximum timeout value is 120 seconds (2 minutes).

--enable-redirector

Enable or disable NDMP redirector which does load distribution of NDMP backup and restore operations.

--throttler-cpu-threshold

Enable or disable NDMP throttler which limits CPU usage for NDMP backup and restore operations on each node.

--enable-throttler

When NDMP throttler is enabled, NDMP makes sure that NDMP backup or restore operations will not go over this CPU threshold. The value is represented in percentage which ranges from 1 to 100. The default value is 50.

isi ndmp settings global view

Displays NDMP global settings.

Syntax

```
isi ndmp settings global view
  [--format {list | json}]
```

Options

--format {list | json}

Displays NDMP global settings in list or JSON format.

Example

The following is an example of the output generated by this command:

```
Service: True
Port: 10000
Dma: emc
Bre Max Num Contexts: 64
Msb Context Retention Duration: 300
Msr Context Retention Duration: 600
```

isi ndmp settings preferred-ips create

For an NDMP three-way operation performed using Avamar, creates an NDMP preferred IP setting.

Syntax

```
isi ndmp settings preferred-ips create --scope <scope> --data-subnets <subnets>
  [{verbose | -v}]
```

Options

--scope <scope>

Specifies the scope of the preferred IP setting. The scope determines the conditions under which the IPs listed under the `data-subnets` will be preferred during a three-way NDMP backup or restore operation. The scope can either be the subnet that is receiving the incoming NDMP request or it can be a cluster in case of a cluster-wide preference. There can be upto one preference setting for each subnet scope and one for the *cluster* scope.

--data-subnets <subnets>

Specifies a comma-separated list of flexnet subnet names where the IPs in the subnets are preferred for outgoing data (during a backup) or incoming data (during a restore). The list of IPs are rearranged according to the order of subnets listed under `data-subnets`. If an IP is in the listed `data-subnets`, that IP is placed at the top of the list. A subnet in the `data-subnets` has no effect if none of the IPs in the list belong to the subnet. The preferences will be applied only under the condition specified by the `scope` parameter. The `scope` and `data-subnets` values can be set to the same subnet. In this case, the same subnet is used for the NDMP outgoing data even as the incoming data comes in on that subnet. For example, if the `scope` is `groupnet0.subnet0`, the `data-subnets` value is `10gnet.subnet0,globalnet0.subnet0`, and the NDMP data for a backup operation comes in over `groupnet0.subnet0`, the IP of `10gnet.subnet0` is placed at the top of the list. However, if that IP is not available, then the IP of `globalnet0.subnet0` is placed at the top of the list. The subnet names must always be separated by commas.

--verbose | -v

Display additional details.

isi ndmp settings preferred-ips delete

For an NDMP three-way operation performed using Avamar, deletes an NDMP preferred IP setting.

Syntax

```
isi ndmp settings preferred-ips delete --scope <scope>
  [--verbose | -v]
  [--help | -h]
```

Option

--scope <scope>

Scope of the preferred IP setting. You can set the preferred IP to have a cluster-wide scope by specifying a value for the cluster or you can select a OneFS-configured subnet, for example, `groupnet0.mysubnet1`.

--verbose | -v

Display additional details.

--help | -h

Display help for this command.

isi ndmp settings preferred-ips list

For an NDMP three-way operation performed using Avamar, lists all the NDMP preferred IP settings.

Syntax

```
isi ndmp settings preferred-ips list
  [--limit | -l] <integer>]
  [--format (table | json | csv | list)]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose | -v]
  [--help | -h]
```

Options

--limit | -l <integer>

The number of NDMP preferred IP settings to display.

--format (table | json | csv | list)

Display the NDMP preferred IP settings in a tabular, JSON, CSV, or list format.

--no-header | -a

Do not display headers in CSV or table formats.

--no-footer | -z

Do not display table summary footer information.

--verbose | -v

Display additional details.

isi ndmp settings preferred-ips modify

For an NDMP three-way operation performed using Avamar, modifies an existing NDMP preferred IP setting.

Syntax

```
isi ndmp settings preferred-ips modify --scope <scope>
  [--data-subnets <subnet> | --add-data-subnets <subnet> |
  --remove-data-subnets <subnet>]
  [--verbose | -v]
  [--help | -h]
```

Options

--scope <scope>

Scope of the NDMP preferred IP setting. You can set the preferred IP to have a cluster-wide scope by specifying a value for the cluster or you can select a OneFS-configured subnet, for example, `groupnet0.mysubnet1`.

--data-subnets <subnet>

A network subnet value. Specify a value for `--data-subnets` for each additional network subnet that you want to specify. The subnet names must be separated by commas.

--add-data-subnets <subnet>

Add a network subnet. Specify `--add-data-subnets` for each network subnet that you want to add. The subnet names must be separated by commas.

--remove-data-subnets <subnet>

Remove a network subnet. Specify `--remove-data-subnets` for each network subnet that you want to remove. The subnet names must be separated by commas.

--verbose | -v

Display additional details.

--help | -h

Display help for this command.

isi ndmp settings preferred-ips view

For an NDMP three-way operation performed using Avamar, displays details of an NDMP preferred IP setting.

Syntax

```
isi ndmp settings preferred-ips view --scope <scope>
  [--format (list | json)]
  [--help | -h]
```

Options

--scope <scope>

Scope of the NDMP preferred IP setting. You can set the preferred IP to have a cluster-wide scope by specifying `cluster` as the value or you can select a OneFS-configured subnet, for example, `groupnet0.mysubnet1`.

--format (list | json)

Display the NDMP preferred IP settings in a list or JSON format.

isi ndmp settings variables create

Sets the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables create --path <path> --name <name> --value <value>
```

For a list of available environment variables, see the *NDMP environment variables* section in the version-appropriate *OneFS Backup and Recovery Guide*.

Options

--path <path>

Applies the default NDMP environment variable value to the specified path. The directory path must be within `/ifs`.

--name <name>

Specifies the NDMP environment variable to define.

--value <value>

Specifies the value to be applied to the NDMP environment variable.

Examples

The following command enables snapshot-based incremental backups to be performed for `/ifs/data/media` by default:

```
isi ndmp settings variables create --path=/ifs/data/media BACKUP_MODE SNAPSHOT
```

isi ndmp settings variables delete

Deletes the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables delete  
  [--path <path>  
  [--name <name>  
  [--force]  
  [--verbose]
```

Options

For a list of available environment variables, see the *NDMP environment variables* section in the version-appropriate *OneFS Backup and Recovery Guide*.

<path>

Applies the default NDMP-environment-variable value to the specified path. This must be a valid directory path within `/ifs`.

<name>

The name of the variable to be deleted. If you do not specify a variable name, all environment variables are deleted for the specified path.

If this option is not specified, deletes default values for all the NDMP environment variables for the given directory.

{--force | -f}

Skips the confirmation prompt.

{--verbose | -v}

Displays more detailed information.

Examples

The following command removes all default NDMP settings for `/ifs/data/media`:

```
isi ndmp settings variables delete --path=/ifs/data/media
```

The following command removes the default file-history setting for backing up `/ifs/data/media`:

```
isi ndmp settings variables delete --path=/ifs/data/media --name=HIST
```

isi ndmp settings variables list

Lists all preferred NDMP environment variables.

Syntax

```
isi ndmp settings variables list  
  [--path <path>]  
  [--format {table | json | csv | list}]  
  [--no-header]  
  [--no-footer]
```

Options

--path <path>

Applies the default NDMP-environment-variable value to the specified path.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

isi ndmp settings variables modify

Modifies the default value for an NDMP environment variable for a given path.

Syntax

```
isi ndmp settings variables modify --path <path> --name <name> --value <value>
```

Options

For a list of available environment variables, see the *NDMP environment variables* section in the version-appropriate *OneFS Backup and Recovery Guide*.

<path>

Applies the default NDMP-environment-variable value to the specified path. This must be a valid directory path within */ifs*.

<name>

Specifies the NDMP environment variable to be defined.

<value>

Specifies the value to be applied to the NDMP environment variable.

isi ndmp users create

Creates a new NDMP user.

Syntax

```
isi ndmp users create --name <name>
[--password <string>]
```

Options

--name <name>

The name of the user.

--password <string>

The password for the new NDMP user. If you do not specify a password, the new user will be prompted to enter a password, and will be prompted to confirm the password by entering it again. This command fails if the specified user already exists.

Examples

The following command creates an NDMP user account with username `ndmp_user` and password `1234`:

```
isi ndmp users create --name=ndmp_user --password=1234
```

isi ndmp users delete

Deletes a specified NDMP user.

Syntax

```
isi ndmp users delete --name <name>
[--force]
[--verbose]
```

Options

--name <name>

The name of the NDMP user to delete.

{--force | -f}

Skips the confirmation prompt.

{verbose | -v}

Displays more detailed information.

isi ndmp users list

Lists all NDMP users.

Syntax

```
isi ndmp users list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

Example

This is an example of the output created by this command:

```
Name
-----
ndmp_nick
ndmp_lisa
ndmp_jason
-----
Total: 3
```

isi ndmp users modify

Changes the password for the specified NDMP user.

Syntax

```
isi ndmp users delete --name <name>
```

Options

--name <name>

The name of NDMP user you are modifying.

isi ndmp users view

View a specific NDMP user.

Syntax

```
isi ndmp users view --name <name>
[--format {list | json}]
```

Output

--name <name>

The name of the NDMP user.

--format {list | json}

Lists the NDMP user in the specified format.

Example

The following is an example of the output created by this command, for an NDMP user named `ndmp_lisa`, and with JSON format specified:

```
[{"id": "ndmp_lisa", "name": "ndmp_lisa"}]
```

isi network dnscache flush

Simultaneously flushes the DNS cache of each groupnet that has enabled DNS caching.

Syntax

```
isi network dnscache flush
[--verbose]
```

Options

{ **--verbose** | **-v**}

Displays more detailed information.

isi network dnscache modify

Modifies global DNS cache settings for each DNS cache that is enabled per groupnet.

Syntax

```
isi network dnscache modify
[--cache-entry-limit <integer>]
[--revert-cache-entry-limit]
[--cluster-timeout <integer>]
[--revert-cluster-timeout]
[--dns-timeout <integer>]
[--revert-dns-timeout]
[--eager-refresh <integer>]
```

```

[--revert-eager-refresh]
[--testping-delta <integer>]
[--revert-testping-delta]
[--ttl-max-noerror <integer>]
[--revert-ttl-max-noerror]
[--ttl-min-noerror <integer>]
[--revert-ttl-min-noerror]
[--ttl-max-nxdomain <integer>]
[--revert-ttl-max-nxdomain]
[--ttl-min-nxdomain <integer>]
[--revert-ttl-min-nxdomain]
[--ttl-max-other <integer>]
[--revert-ttl-max-other]
[--ttl-min-other <integer>]
[--revert-ttl-min-other]
[--ttl-max-servfail <integer>]
[--revert-ttl-max-servfail]
[--ttl-min-servfail <integer>]
[--revert-ttl-min-servfail]
[--verbose]

```

Options

--cache-entry-limit <integer>

Specifies the maximum number of entries that the DNS cache can contain. The limit must be a value between 1024 and 1048576. The default value is 65536 entries.

--revert-cache-entry-limit

Sets the value of `--cache-entry-limit` to the default system value.

--cluster-timeout <integer>

Specifies the timeout limit, in seconds, for calls made to other nodes in the cluster. The limit must be a value between 0 and 60. The default value is 5.

--revert-cluster-timeout

Sets the value of `--cluster-timeout` to the default system value.

--dns-timeout <integer>

Specifies the timeout limit, in seconds, for calls made to the DNS resolver. The limit must be a value between 0 and 60. The default value is 5.

--revert-dns-timeout

Sets the value of `--dns-timeout` to the default system value.

--eager-refresh <integer>

Specifies the lead time, in seconds, to refresh cache entries that are nearing expiration. The lead time must be a value between 0 and 30. The default value is 0.

--revert-eager-refresh

Sets the value of `--eager-refresh` to the default system value.

--testping-delta <integer>

Specifies the delta, in seconds, for checking the cbind cluster health. The delta must be a value between 0 and 60. The default value is 30.

--revert-testping-delta

Sets the value of `--testping-delta` to the default system value.

--ttl-max-noerror <integer>

Specifies the upper time-to-live boundary, in seconds, on cache hits. The boundary must be a value between 0 and 3600. The default value is 3600.

--revert-ttl-max-noerror

Sets the value of `--ttl-max-noerror` to the default system value.

--ttl-min-noerror <integer>

Specifies the lower time-to-live boundary, in seconds, on cache hits. The boundary must be a value between 0 and 3600. The default value is 30.

--revert-ttl-min-noerror

Sets the value of `--ttl-min-noerror` to the default system value.

--ttl-max-nxdomain <integer>

Specifies the upper time-to-live boundary, in seconds, for `nxdomain` failures. The boundary must be a value between 0 and 3600. The default value is 3600

--revert-ttl-max-nxdomain

Sets the value of `--ttl-max-nxdomain` to the default system value.

--ttl-min-nxdomain <integer>

Specifies the lower time-to-live boundary, in seconds, for `nxdomain` failures. The boundary must be a value between 0 and 3600. The default value is 15.

--revert-ttl-min-nxdomain

Sets the value of `--ttl-min-nxdomain` to the default system value.

--ttl-max-other <integer>

Specifies the upper time-to-live boundary, in seconds, for non-`nxdomain` failures. The boundary must be a value between 0 and 3600. The default value is 60.

--revert-ttl-max-other

Sets the value of `--ttl-max-other` to the default system value.

--ttl-min-other <integer>

Specifies the lower time-to-live boundary, in seconds, for non-`nxdomain` failures. The boundary must be a value between 0 and 3600. The default value is 0.

--revert-ttl-min-other

Sets the value of `--ttl-min-other` to the default system value.

--ttl-max-servfail <integer>

Specifies the upper time-to-live boundary, in seconds, for DNS server failures. The boundary must be a value between 0 and 3600. The default value is 3600.

--revert-ttl-max-servfail

Sets the value of `--ttl-max-servfail` to the default system value.

--ttl-min-servfail <integer>

Specifies the lower time-to-live boundary, in seconds, for DNS server failures. The boundary must be a value between 0 and 3600. The default value is 300.

--revert-ttl-min-servfail

Sets the value of `--ttl-min-servfail` to the default system value.

{--verbose | -v}

Displays more detailed information.

isi network dnscache view

Displays DNS cache settings.

Syntax

```
isi network dnscache view
```

Options

There are no options for this command.

isi network external modify

Modifies global external network settings on the cluster.

Syntax

```
isi network external modify
[--sbr {true | false}]
[--revert-sbr]
[--sc-rebalance-delay <integer>]
[--revert-sc-rebalance-delay]
[--tcp-ports <integer>]
[--clear-tcp-ports]
[--add-tcp-ports <integer>]
[--remove-tcp-ports <integer>]
[--revert-tcp-ports]
[--verbose]
```

Options

--sbr {true | false}

Enables or disables source-based routing on the cluster. Source-based routing is disabled by default.

--revert-sbr

Sets the value of `--sbr` to the default system value.

--sc-rebalance-delay <integer>

Specifies a period of time (in seconds) that should pass after a qualifying event before an automatic rebalance is performed. The default value is 0 seconds.

--revert-sc-rebalance-delay

Sets the value of `--sc-rebalance-delay` to the default system value.

--tcp-ports <integer>

Sets a list of recognized client TCP ports. 65535 is the maximum supported port number. You can specify multiple TCP ports separated by commas, or specify this option for each additional TCP port.

--clear-tcp-ports

Removes all client TCP ports.

--add-tcp-ports <integer>

Adds one or more recognized client TCP ports, separated by commas, to the existing list. 65535 is the maximum supported port number.

--remove-tcp-ports <integer>

Removes one or more recognized client TCP ports, separated by commas.

--revert-tcp-ports

Sets the value of `--tcp-ports` to the default system value.

{ --verbose | -v }

Displays more detailed information.

isi network external view

Displays configuration settings for the external network.

Syntax

```
isi network external view
```

Options

There are no options for this command.

isi network groupnets create

Creates a groupnet which defines the client DNS settings applied to services that connect through the groupnet.

Syntax

```
isi network groupnets create <id>
  [--description <string>]
  [--dns-cache-enabled {true | false}]
  [--dns-search <domain name>]
  [--dns-servers <ip address>]
  [--dns-options <string>]
  [--server-side-dns-search {true | false}]
  [--verbose]
```

Options

<id>

Specifies a unique ID for the groupnet. The ID can be up to 32 alphanumeric characters long and can include underscores or hyphens, but cannot include spaces or other punctuation. The ID cannot exceed 32 characters.

--description <string>

Specifies an optional description of the groupnet. The description cannot exceed 128 bytes.

--dns-cache-enabled {true | false}

Specifies whether DNS caching for the groupnet is enabled. DNS caching is enabled by default.

--dns-search <domain name>

Sets the list of DNS search suffixes. Suffixes are appended to domain names that are not fully qualified. The list cannot contain more than six suffixes.

 **NOTE: Do not begin suffixes with a leading dot; leading dots are automatically added.**

--dns-servers <ip address>

Sets a list of DNS IP addresses. Nodes issue DNS requests to these IP addresses. The list cannot contain more than three IP addresses.

--dns-options <string>

Sets the DNS resolver option. The only valid value for this option is `rotate`.

--server-side-dns-search {true | false}

Specifies whether server-side DNS searching is enabled, which appends DNS search lists to client DNS inquiries handled by a SmartConnect service IP address. Server-side search is enabled by default.

{--verbose | -v}

Displays more detailed information.

isi network groupnets delete

Deletes a groupnet from the cluster. You cannot delete the default groupnet from the system.

If the groupnet is associated with an access zone, an authentication provider, removal of the groupnet from the system might affect several other areas of OneFS and should be performed with caution. When you delete a groupnet, client connections to each subnet-pool association in the groupnet are lost. Deleting a groupnet that is in use can prevent access to the cluster. Client connections to the cluster through any subnet-pool in the deleted groupnet will be terminated.

Syntax

```
isi network groupnets delete <id>
  [--force]
  [--verbose]
```

Options

< id>

Specifies the ID of the groupnet to be deleted.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network groupnets list

Retrieves a list of groupnets that exist on the cluster.

Syntax

```
isi network groupnets list
  [--limit | -l] <integer>
  [--sort {description | dns_cache_enabled | id | name | server_side_dns_search}]
  [--descending | -d]
  [--format {true | table | json | csv | list}]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose | -v]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number.

--sort {description | dns_cache_enabled | id | name | server_side_dns_search}

Sorts output displayed by the specified attribute.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ **--no-header** | **-a**}

Displays table and CSV output without headers.

{ **--no-footer** | **-z**}

Displays table output without footers.

{ **--verbose** | **-v**}

Displays more detailed information.

isi network groupnets modify

Modifies a groupnet which defines the DNS settings applied to services that connect through the groupnet.

Syntax

```
isi network groupnets modify <id>
  [--description <string>]
  [--clear description]
  [--dns-cache-enabled {true | false}]
  [--revert-dns-cache-enabled]
  [--dns-search <domain name>]
  [--clear-dns-search]
  [--add-dns-search <domain name>]
  [--remove-dns-search <domain name>]
  [--dns-servers <IP address>]
  [--clear-dns-servers]
  [--add-dns-servers <IP address>]
  [--remove-dns-servers <IP address>]
  [--dns-options <string>]
  [--clear-dns-options]
  [--add-dns-options <string>]
  [--remove-dns-options <string>]
  [--name <string>]
  [--server-side-dns-search {true | false}]
  [--revert server-side-dns-search]
  [--verbose]
```

Options

<id>

Specifies the ID of the groupnet to be modified.

--description <string>

Specifies an optional description of the groupnet. This option overwrites the existing description. The description cannot exceed 128 bytes.

--clear-description

Clears the current description.

--dns-cache-enabled {true | false}


Specifies whether DNS caching for the groupnet is enabled. DNS caching is enabled by default.

--revert-dns-cache-enabled

Sets the value of `--dns-cache-enabled` to the system default value.

--dns-search <domain name>

Sets the list of DNS search suffixes. Suffixes are appended to domain names that are not fully qualified. The list cannot contain more than six suffixes. This option overwrites the current list of DNS search suffixes.

 **NOTE: Do not begin suffixes with a leading dot; leading dots are automatically added.**

--clear-dns-search

Removes the current list of DNS search suffixes.

- add-dns-search <domain name>**
Adds one or more DNS search suffixes to the current list.
- remove-dns-search <domain name>**
Removes one or more DNS search suffixes from the current list.
- dns-servers <IP address>**
Sets a list of DNS IP addresses. Nodes issue DNS requests to these IP addresses. The list cannot contain more than three IP addresses. This option overwrites the current list of DNS IP addresses.
- clear-dns-servers**
Removes the current list of DNS servers.
- add-dns-servers <IP address>**
Adds one or more DNS servers to the current list.
- remove-dns-servers <IP address>**
Removes one or more DNS servers from the current list.
- dns-options <string>**
Sets the DNS resolver option. The only valid value for this option is `rotate`.
- clear-dns-options**
Removes the current list of DNS resolver options.
- add-dns-options <string>**
Adds one or more DNS resolver options to the current list.
- remove-dns-options <string>**
Removes one or more DNS resolver options from the current list.
- name <string>**
Specifies a new name for the groupnet. The ID can be up to 32 alphanumeric characters long and can include underscores or hyphens, but cannot include spaces or other punctuation. The name cannot exceed 32 characters.
- server-side-dns-search {true | false}**
Specifies whether server-side DNS searching is enabled, which appends DNS search lists to client DNS inquiries handled by a SmartConnect service IP address. Server-side search is enabled by default.
- revert-server-side-dns-search**
Sets the value of `--server-side-dns-search` to the system default value.
- {--verbose | -v}**
Displays more detailed information.

isi network groupnets view

Displays the configuration details of a specific groupnet on the cluster.

Syntax

```
isi network groupnets view <id>
```

Options

- <id>**
Specifies the ID of the groupnet to be viewed.

isi network interfaces list

Displays a list of network interfaces on the cluster.

Syntax

```
isi network interfaces list
  [--nodes <integer>]
  [--show-inactive]
  [--limit <integer>]
  [--sort {lnn |
          name} | status]
  [--descending]
  [--format {table |
            json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, the command displays a list of all network interfaces on the cluster.

--nodes <lnn>

Lists interfaces only from the specified nodes. Specify nodes by Logical Node Number. Separate multiple nodes by commas.

--show-inactive

Includes inactive interfaces in the output.

{ --limit | -l } <integer>

Displays no more than the specified number of interfaces.

--sort {lnn | name | status}

Sorts output displayed by the specified attribute.

{ --descending | -d }

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

Examples

The following command lists network interfaces on node 1:

```
isi network interfaces list --nodes=1
```

The system displays output similar to the following example:

LNN	Name	Status	Owners	IP Addresses
1	10gige-1	Up	-	

```

1  10gige-2      No Carrier      -          -
1  10gige-agg-1 Not Available    -          -
1  ext-1         Up              groupnet0.subnet0.pool0 198.51.100.0
1  ext-2         Up              -          -
1  ext-agg       Not Available   -          -
-----

```

Total: 6

isi network pools create

Creates a pool of IP addresses within a subnet. A SmartConnect Advanced license is required to create more than one pool within a subnet.

Syntax

```

isi network pools create <id>
  [--access-zone <zone-name>]
  [--aggregation-mode {roundrobin | failover | lACP | fec}]
  [--alloc-method {dynamic | static}]
  [--description <string>]
  [--ifaces <node-interface-range>]...
  [--ranges <ip-address-range>]...
  [--rebalance-policy{manual | auto}]
  [--sc-auto-unsuspend-delay <integer>]
  [--sc-connect-policy {roundrobin | conn_count | throughput | cpu_usage}]
  [--sc-dns-zone <domain-name>]
  [--sc-dns-zone-aliases <domain-name>]...
  [--sc-failover-policy {roundrobin | conn_count | throughput | cpu_usage}]
  [--sc-subnet <string>]
  [--sc-ttl <integer>]
  [--static-routes <route>]...
  [--force]
  [--verbose]

```

Options

<id>

Specifies the ID of the new pool that you want to create. The pool must be added to an existing groupnet and subnet. The ID can be up to 32 alphanumeric characters long and can include underscores or hyphens, but cannot include spaces or other punctuation. Specify the pool ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0. The pool name must be unique in the subnet.

--access-zone <zone-name>

Associates an access zone with the pool. Clients will be allowed to connect to the specified access zone only through IP addresses in the pool. The access zone must belong to the same groupnet as the IP address pool.

--aggregation-mode {roundrobin | failover | lACP | fec}

Specifies how outgoing traffic is distributed across aggregated network interfaces. The aggregation mode is applied only if at least one aggregated network interface is a member of the IP address pool.

The following values are valid:

- roundrobin** Rotates connections through the nodes in a first-in, first-out sequence, handling all processes without priority. Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.
- failover** Switches to the next active interface when the primary interface becomes unavailable. Manages traffic only through a primary interface. The second interface takes over the work of the first as soon as it detects an interruption in communication.


lACP	Supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Balances outgoing traffic across the interfaces based on hashed protocol header information that includes the source and destination address and the VLAN tag, if available. Also assembles interfaces of the same speed into groups called Link Aggregated Groups (LAGs) and balances traffic across the fastest LAGs. This option is the default mode for new pools.
fec	Provides static balancing on aggregated interfaces through the Cisco Fast EtherChannel (FEC) driver, which is found on older Cisco switches. Capable of load-balancing traffic across Fast Ethernet links. Enables multiple physical Fast Ethernet links to combine into one logical channel.

--alloc-method {dynamic | static}

Specifies the method by which IP addresses are allocated to the network interfaces that are members of the pool.

The following values are valid:

static	Assigns each network interface in the IP address pool a single, permanent IP address from the pool. Depending on the number of IP addresses available, some IP addresses might go unused. The static option is the default setting.
dynamic	Specifies that all pool IP addresses must be assigned to a network interface at all times. Enables multiple IP addresses to be assigned to an interface. If a network interface becomes unavailable, this option helps to ensure that the assigned IP address are redistributed to another interface.

 **NOTE: This option is only available if a SmartConnect Advanced license is active on the cluster.**

--description <string>

Specifies an optional description of the IP address pool. The description cannot exceed 128 bytes.


--ifaces <node-interface-range>...

Specifies which network interfaces should be members of the IP address pool. Specify network interfaces in the following format:

```
<node>:<interface>
```

To specify a range of nodes, separate the lower and upper node IDs with a dash (-). To specify multiple network interfaces, separate each interface with a comma. The following example adds the interfaces from nodes 1, 2 and 3:

```
--ifaces 1-2:ext-1,3:ext-2,1:10gige-agg-1,3:10gige-1
```

 **NOTE: If you attempt to add an interface that has already been added as part of an aggregated interface, you will receive an error message.**

--ranges <ip-address-range>...

Specifies one or more IP address ranges for the pool. IP addresses within these ranges are assigned to the network interfaces that are members of the pool.

Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

--rebalance-policy {manual | auto}

Specifies when to redistribute pool IP addresses if a network interface that was previously unavailable becomes available.

manual Requires that connection rebalancing be performed manually after network interface failback.

To manually rebalance all IP addresses in a specific pool, run the following command:

```
isi network pools rebalance-ips
```

To manually rebalance all IP addresses across the cluster, run the following command:

```
isi network sc-rebalance-all
```

auto Causes connections to be rebalanced automatically after network interface failback. This is the default value.

--sc-auto-unsuspend-delay *<integer>*

Specifies the time delay (in seconds) before a node that is automatically unsuspended resumes SmartConnect DNS query responses for the node. During certain cluster operations such as rolling upgrades, general node splits, or node reboots, a node is automatically suspended and then unsuspended by the system.

--sc-connect-policy {**roundrobin** | **conn_count** | **throughput** | **cpu_usage**}

Specifies how incoming DNS queries for client connections are balanced across IP addresses.

The following values are valid:

round-robin Rotates connections through nodes equally. This value is the default policy.

conn-count Assigns connections to the node that has the fewest number of connections.

throughput Assigns connections to the node with the least throughput.

cpu-usage Assigns connections to the node with the lowest CPU usage.

--sc-dns-zone *<domain-name>*

Specifies the SmartConnect DNS zone name for this pool. IP addresses are returned in response to DNS queries to this SmartConnect zone.

--sc-dns-zone-aliases *<domain-name>*

Specifies a list of alternate SmartConnect DNS zone names for the pool. Multiple aliases can be specified in a comma-separated list.

--sc-failover-policy {**roundrobin** | **conn_count** | **throughput** | **cpu_usage**}

Specifies how IP addresses that belong to an unavailable interface are rebalanced across the remaining network interfaces.

The following values are valid:

round-robin Assigns IP addresses across nodes equally. This is the default policy.

conn-count Assigns IP addresses to the node that has fewest number of connections.

throughput Assigns IP addresses to the node with least throughput.

cpu-usage Assigns IP addresses to the node with lowest CPU usage.

--sc-subnet *<string>*

Specifies the name of the service subnet that is responsible for handling DNS requests for the SmartConnect zone.

--sc-ttl *<integer>*

Specifies the time-to-live value for SmartConnect DNS query responses (in seconds). DNS responses are only valid for the time specified. The default value is 0 seconds.

--static-routes *<route>*

Designates an IP addresses as a static route and specifies the destination gateway. If a client connects through a static route IP address, outgoing client traffic is routed through the specified gateway. Multiple routes can be specified in a comma-separated list.

Specify the static route in the following classless inter-domain routing (CIDR) notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

{**--verbose** | **-v**}

Displays more detailed information.

{**--force** | **-f**}

Forces commands without warnings.

Examples

The following command creates a new IP address pool called pool1 under groupnet0.subnet0 and assigns IP addresses 198.51.100.10-198.51.100.20 to ext-1 network on nodes 1, 2, and 3. The SmartConnect zone name of this pool is storage.company.com, but it accepts the alias of storage.company:

```
isi network pools create groupnet0.subnet0.pool1 \  
  --ranges=192.168.8.10-192.168.8.15 --ifaces=1-3:ext-1 \  
  --sc-dns-zone=storage.company.com --sc-dns-zone-aliases=storage.company
```

The following command creates a new IP address pool called pool1 under groupnet0.subnet0 and assigns IP addresses 198.51.100.10-198.51.100.20 to the pool. The command also includes aggregated interfaces from nodes 1-3 and specifies FEC as the aggregation mode:

```
isi network pools create groupnet0.subnet0.pool1 \  
  --ranges=192.168.8.10-192.168.8.15 --ifaces=1-3:10gige-agg-1 \  
  --aggregation-mode=fec
```

The following command creates a new IP address pool called pool1 under groupnet0.subnet0, assigns IP addresses 198.51.100.10-198.51.100.20 to the pool, and specifies that connection rebalancing must be performed manually:

```
isi network pools create groupnet0.subnet0.pool1 \  
  --ranges=192.168.8.10-192.168.8.15 --alloc-method=dynamic \  
  --rebalance-policy=manual
```

isi network pools delete

Deletes IP address pools.

Deleting an IP address pool that is in use can prevent access to the cluster. Client connections to the cluster through any IP address in the deleted pool will be terminated.

Syntax

```
isi network pools delete <id>  
  [--force]  
  [--verbose]
```

Options

<id>...

Specifies the ID of the IP address pool to be deleted. Specify the pool ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network pools list

Retrieves a list of IP address pools that exist on the cluster.

Syntax

```
isi network pools list
  [--subnet-id <string>...]
  [--groupnet <string>...]
  [--subnet <string>...]
  [--limit | -l] <integer>]
  [--sort {aggregation_mode | alloc_method | description | id | name | rebalance_policy |
sc_auto_suspend_dealy | sc_connect_policy | sc_dns_zone | sc_failover_policy | sc_subnet |
sc_ttl}]
  [--descending | -d}]
  [--format {table | json | csv | list}]
  [--no-header | -a}]
  [--no-footer | -z}]
  [--verbose | -v}]
```

Options

If no options are specified, the command displays a list of all IP address pool on the cluster.

--subnet-id <string>...

Displays IP address pools only from the specified subnet ID. Specify the subnet ID in the following format:

```
<groupnet_name>.<subnet_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0.

--groupnet <string>...

Displays IP address pools only from the specified groupnet name.

--subnet <string>...

Displays IP address pools only subnets with the specified name.

{ --limit | -l] <integer>

Displays no more than the specified number.

--sort {aggregation_mode | alloc_method | description | id | name | rebalance_policy | sc_auto_suspend_dealy | sc_connect_policy | sc_dns_zone | sc_failover_policy | sc_subnet | sc_ttl}

Sorts output displayed by the specified attribute.

{ --descending | -d}

Displays output in reverse order.

--format {true | table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a}

Displays table and CSV output without headers.

{ --no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi network pools modify

Modifies IP address pool settings.

Syntax

```
isi network pools modify <id>
  [--access-zone <zone-name>]
  [--revert-access-zone]
  [--aggregation-mode {roundrobin | failover | lacp | fec}]
  [--revert-aggregation-mode]
  [--alloc-method {dynamic | static}]
  [--revert-alloc-method]
  [--description <string>]
  [--clear-description]
  [--ifaces <node-interface-range>]...
  [--clear-ifaces]
  [--add-ifaces <node-interface-range>]...
  [--remove-ifaces <node-interface-range>]...
  [--name <string>]
  [--ranges <ip-address-range>]...
  [--clear-ranges]
  [--add-ranges <ip-address-range>]...
  [--remove-ranges <ip-address-range>]...
  [--rebalance-policy {manual | auto}]
  [--revert-rebalance-policy]
  [--sc-auto-unsuspend-delay <integer>]
  [--revert-sc-auto-unsuspend-delay]
  [--sc-connect-policy {roundrobin | conn_count | throughput | cpu_usage}]
  [--revert-sc-connect-policy]
  [--sc-dns-zone <domain-name>]
  [--sc-dns-zone-aliases <domain-name>]...
  [--clear-sc-dns-zone-aliases]
  [--add-sc-dns-zone-aliases <domain-name>]...
  [--remove-sc-dns-zone-aliases <domain-name>]...
  [--sc-failover-policy {roundrobin | conn_count | throughput | cpu_usage}]
  [--revert-sc-failover-policy]
  [--sc-subnet <string>]
  [--sc-ttl <integer>]
  [--revert-sc-ttl]
  [--static-routes <route>]...
  [--clear-static-routes]
  [--add-static-routes <route>]...
  [--remove-static-routes <route>]...
  [--force]
  [--verbose]
```

Options

<id>

Specifies the ID of the IP address pool that you want to modify. Specify the ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0. The pool name must be unique in the subnet.

--access-zone <zone-name>

Associates an access zone with the pool. Clients will be allowed to connect to the specified access zone only through IP addresses in the pool. The access zone must belong to the same groupnet as the IP address pool.

--revert-access-zone

Sets the value of **--access-zone** to the system default value.

--aggregation-mode {roundrobin | failover | lacp | fec}

Specifies how outgoing traffic is distributed across aggregated network interfaces. The aggregation mode is applied only if at least one aggregated network interface is a member of the IP address pool.

The following values are valid:

roundrobin	Rotates connections through the nodes in a first-in, first-out sequence, handling all processes without priority. Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.
failover	Switches to the next active interface when the primary interface becomes unavailable. Manages traffic only through a primary interface. The second interface takes over the work of the first as soon as it detects an interruption in communication.
lacp	Supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Balances outgoing traffic across the interfaces based on hashed protocol header information that includes the source and destination address and the VLAN tag, if available. Also assembles interfaces of the same speed into groups called Link Aggregated Groups (LAGs) and balances traffic across the fastest LAGs. This option is the default mode for new pools.
fec	Provides static balancing on aggregated interfaces through the Cisco Fast EtherChannel (FEC) driver, which is found on older Cisco switches. Capable of load-balancing traffic across Fast Ethernet links. Enables multiple physical Fast Ethernet links to combine into one logical channel.

--revert-aggregation-mode


Sets the value of `--aggregation-mode` to the system default value.

--alloc-method {dynamic | static}

Specifies the method by which IP addresses are allocated to the network interfaces that are members of the pool.

The following values are valid:

static	Assigns each network interface in the IP address pool a single, permanent IP address from the pool. Depending on the number of IP addresses available, some IP addresses might go unused. The static option is the default setting.
dynamic	Specifies that all pool IP addresses must be assigned to a network interface at all times. Enables multiple IP addresses to be assigned to an interface. If a network interface becomes unavailable, this option helps to ensure that the assigned IP address are redistributed to another interface.

 **NOTE: This option is only available if a SmartConnect Advanced license is active on the cluster.**

--revert-alloc-method

Sets the value of `--alloc-method` to the system default value.

--description <string>

Specifies an optional description of the IP address pool. This option overwrites the existing description. The description cannot exceed 128 bytes.

--clear-description

Clears the description of the IP address pool.

--ifaces <node-interface-range>...

Adds network interfaces to the IP address pool. Specify network interfaces in the following format:

```
<node>:<interface>
```

To specify a range of nodes, separate the lower and upper node IDs with a dash (-). To specify multiple network interfaces, separate each interface with a comma. The following example adds the interfaces from nodes 1, 2 and 3:

```
--ifaces 1-2:ext-1,3:ext-2,1:10gige-agg-1,3:10gige-1
```

--clear-ifaces

Removes all network interfaces from the IP address pool.

--add-ifaces <node-interface-range>...

Adds one or more network interfaces to the IP address pool.

--remove-ifaces <node-interface-range>...

Removes one or more network interfaces from the IP address pool.

--name <string>

Specifies a new name for the IP address pool. The name can be up to 32 alphanumeric characters long and can include underscores or hyphens, but cannot include spaces or other punctuation. The new pool name must be unique in the subnet.

--ranges <ip-address-range>...

Specifies one or more IP address ranges for the pool. IP addresses within these ranges are assigned to the network interfaces that are members of the pool.

Specify the IP address range in the following format:

```
<low-ip-address>-<high-ip-address>
```

This option overwrites the existing list of IP address ranges. Use the `--add-ranges` and `--remove-ranges` options to modify the existing list.

--clear-ranges

Removes all IP address ranges from the pool.

--add-ranges

Adds one or more IP address ranges to the pool.

--remove-ranges

Removes one or more IP address ranges from the pool.

--rebalance-policy{manual | auto}

Specifies when to redistribute pool IP addresses if a network interface that was previously unavailable becomes available.

manual

Requires that connection rebalancing be performed manually after network interface failback.

To manually rebalance all IP addresses in a specific pool, run the following command:

```
isi network pools rebalance-ips
```

To manually rebalance all IP addresses across the cluster, run the following command:

```
isi network sc-rebalance-all
```

auto

Causes connections to be rebalanced automatically after network interface failback. This is the default value.

--revert-rebalance-policy

Sets the value of `--rebalance-policy` to the system default value.

--sc-auto-unsuspend-delay <integer>

Specifies the time delay (in seconds) before a node that is automatically unsuspended resumes SmartConnect DNS query responses for the node. During certain cluster operations such as rolling upgrades, general node splits, or node reboots, a node is automatically suspended and then unsuspended by the system.

--revert-sc-auto-unsuspend-delay

Sets the value of `--sc-auto-unsuspend-delay` to the system default value.

--sc-connect-policy {roundrobin | conn_count | throughput | cpu_usage}

Specifies how incoming DNS requests for client connections are balanced across IP addresses.

The following values are valid:

roundrobin

Rotates connections through nodes equally. This value is the default policy.

conn_count

Assigns connections to the node that has the fewest number of connections.

throughput Assigns connections to the node with the least throughput.
cpu_usage Assigns connections to the node with the lowest CPU usage.

--revert-sc-connect-policy

Sets the value of `--sc-connect-policy` to the system default value.

--sc-dns-zone <domain-name>

Specifies the SmartConnect DNS zone name for this pool. IP addresses are returned in response to DNS queries to this SmartConnect zone.

--sc-dns-zone-aliases <domain-name>...

Specifies a list of alternate SmartConnect DNS zone names for the pool. Multiple aliases can be specified in a comma-separated list. This option overwrites the existing list of SmartConnect DNS zone aliases. Use the `--add-sc-dns-zone-aliases` and `--remove-sc-dns-zone-aliases` options to modify the existing list.

--clear-sc-dns-zone-aliases

Removes all SmartConnect DNS zone aliases from the pool.

--add-sc-dns-zone-aliases <domain-name>...

Adds one or more SmartConnect DNS zone aliases to the pool.

--remove-sc-dns-zone-aliases <domain-name>...

Removes one or more SmartConnect DNS zone aliases from the pool.

--sc-failover-policy {roundrobin | conn_count | throughput | cpu_usage}

Specifies how IP addresses that belong to an unavailable interface are rebalanced across the remaining network interfaces.

The following values are valid:

roundrobin Assigns IP addresses across nodes equally. This is the default policy.
conn_count Assigns IP addresses to the node that has fewest number of connections.
throughput Assigns IP addresses to the node with least throughput.
cpu_usage Assigns IP addresses to the node with lowest CPU usage.

--revert-failover-policy

Sets the value of `--sc-failover-policy` to the system default value.

--sc-subnet <string>

Specifies the name of the service subnet that is responsible for handling DNS requests for the SmartConnect zone.

--sc-ttl <integer>

Specifies the time-to-live value for SmartConnect DNS query responses (in seconds). DNS responses are only valid for the time specified. The default value is 0 seconds.

--static-routes <route>...

Designates an IP addresses as a static route and specifies the destination gateway. If a client connects through a static route IP address, outgoing client traffic is routed through the specified gateway. Multiple routes can be specified in a comma-separated list.

Specify the static route in the following classless inter-domain routing (CIDR) notation format:

```
<network-address>/<subnet-mask>-<gateway-ip-address>
```

This option overwrites the existing list of static routes. Use the `--add-static-routes` and `--remove-static-routes` options to modify the existing list.

--clear-static-routes

Removes all static routes from the pool.

--add-static-routes <route>...

Adds one or more static routes to the pool.

--remove-static-routes <route>...

Removes one or more static routes from the pool.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Forces commands without warnings.

isi network pools rebalance-ips

Redistributes the IP addresses in a specified pool across network interface members. Run this command for pools that specify a manual rebalance policy.

Syntax

```
isi network pools rebalance-ips <id>...  
  [--force]  
  [--verbose]
```

Options

<id>...

Specifies the name of the IP address pool to be rebalanced. Specify the pool name in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network pools sc-resume-nodes

Resumes SmartConnect DNS query responses on a node.

Syntax

```
isi network pools sc-resume-nodes <id> <lnn>...  
  [--force]  
  [--verbose]
```

Options

<id>...

Specifies the name of the IP address pool for which SmartConnect DNS query responses should be resumed. Specify the pool name in the following format:

```
<groupnet_id>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0.

<lnn>...

Specifies the Logical Node Number of the node for which SmartConnect DNS query responses should be resumed.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network pools sc-suspend-nodes

Suspends SmartConnect DNS query responses on a node.

Syntax

```
isi network pools sc-suspend-nodes <id> <lnn>...  
  [--force]  
  [--verbose]
```

Options

<id>...

Specifies the name of the IP address pool for which SmartConnect DNS query responses should be suspended. Specify the pool name in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters—for example, groupnet0:subnet1:pool0.

<lnn>...

Specifies the Logical Node Number of the node for which SmartConnect DNS query responses should be suspended.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network pools view

Displays the configuration details of a specific IP address pool on the cluster.

Syntax

```
isi network pools view <id>
```

Options

<id>

Specifies the ID of the IP address pool to be viewed. Specify the pool ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0.

isi network rules create

Creates a provisioning rule to automatically configure new network interfaces that are added to the cluster.

Syntax

```
isi network rules create <id> <iface>
  [--desc <string>]
  [--node-type {any | storage | accelerator | backup-accelerator}]
  [--verbose]
```

Options

<id>

Specifies the ID and location of the new provisioning rule. New network interfaces that meet the rule criteria will be assigned to the IP address pool that contains the rule. Valid IDs include the groupnet, subnet, pool, and rule name. The rule name must be unique throughout the given IP address pool. Specify the rule ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>.<rule_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0:rule3. The rule name must be unique in the pool.

<iface>

Specifies the network interface name the rule applies to. To view a list of interfaces on your system, run the `isi network interfaces list` command.

--description <string>

Specifies an optional description of the provisioning rule. The description cannot exceed 128 bytes.

--node-type {any | storage | accelerator | backup-accelerator}

Sets the provisioning rule to apply to one or more of the specified type of node. The default setting is `any`.

{--verbose | -v}

Displays more detailed information.

isi network rules delete

Deletes provisioning rules.

Syntax

```
isi network rules delete <id>
```

Options

<id>...

Specifies the ID of the provisioning rule to be deleted. Specify the rule ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>.<rule_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0:rule3.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network rules list

Retrieves a list of provisioning rules on the cluster.

Syntax

```
isi network rules list
  [--pool-id <string>]
  [--groupnet <string>]
  [--subnet <string>]
  [--pool <string>]
  [--limit | -l] <integer>
  [--sort {id | description | iface | node_type | name}]
  [--descending | -d]
  [--format {table | json | csv | list}]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose]
```

Options

If no options are specified, the command displays a list of all provisioning rules on the cluster.

--pool-id <string>

Displays provisioning rules only from the specified pool ID. Specify the pool ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0:rule3.

--groupnet <string>

Displays provisioning rules only from the specified groupnet name.

--subnet <string>

Displays provisioning rules only from subnets with the specified name.

--pool <string>

Displays provisioning rules only from IP address pools with the specified name.

{ --limit | -l } <integer>

Displays no more than the specified number.

--sort {id | description | iface | node_type | name}

Sorts output displayed by the specified attribute.

{ --descending | -d }

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

Examples

The following example displays a list of provisioning rules on a node:

```
isi networks list rules
```

The system displays the list of rules in output similar to the following example:

ID	Node Type	Interface
groupnet0.subnet0.pool0.rule0	any	ext-1
groupnet3.subnet3.pool3.rule3	any	ext-4

Total: 2

isi network rules modify

Modifies network provisioning rule settings.

Syntax

```
isi network rules modify <id>
  [--description <string>]
  [--clear-description ]
  [--iface <node_interface>]
  [--name <string>]
  [--node-type {any | storage | accelerator | backup-accelerator}]
  [--revert-node-type ]
  [--verbose]
```

Options

<id>

Specifies the ID of the provisioning rule to be modified. Specify the rule ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>.<rule_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0:rule3.

--description <string>

Specifies an optional description of the provisioning rule. This option overwrites the existing description. The description cannot exceed 128 bytes.

--clear-description

Clears the description of the provisioning rule.

--iface <node_interface>

Specifies the network interface name the rule applies to. This option overwrites the existing interface name.

--name <string>

Specifies a new name for the rule. The new rule name must be unique in the pool.

--node-type {any | storage | accelerator | backup-accelerator}

Sets the provisioning rule to apply to one or more of the specified type of node. The default node type is *any*.

--revert-node-type

Sets the value of `--node-type` to the system default value.

{--verbose | -v}

Displays more detailed information.

isi network rules view

Displays the configuration details of a specific provisioning rule on the cluster.

Syntax

```
isi network rules view <id>
```

Options

<id>

Specifies the ID of the provisioning rule to be viewed. Specify the rule ID in the following format:

```
<groupnet_name>.<subnet_name>.<pool_name>.<rule_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1:pool0:rule3.

isi network sc-rebalance-all

Redistributes IP addresses in all pools on the cluster.

To redistribute IP addresses in a specific pool, run the `isi network pools rebalance-ips` command.

Syntax

```
isi network sc-rebalance-all  
  [--force]  
  [--verbose]
```

Options

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network subnets create

Creates network subnets.

Syntax

```
isi networks create subnet <id> <addr-family> {ipv4 | ipv6} <prefixlen>  
  [--description <string>]  
  [--dsr-addr <ip-address>]...  
  [--gateway <ip-address>]  
  [--gateway-priority <integer>]  
  [--mtu <integer>]  
  [--sc-service-addr <ip-address>]  
  [--vlan-enabled {true | false}]  
  [--vlan-id <integer>]  
  [--verbose]
```

Options

`<id>`

Specifies the ID of the new subnet that you want to create. The subnet must be added to an existing groupnet. The ID can be up to 32 alphanumeric characters long and can include underscores or hyphens, but cannot include spaces or other punctuation. Specify the subnet ID in the following format:

```
<groupnet_name>.<subnet_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1. The subnet name must be unique in the groupnet.

`<addr-family> {ipv4 | ipv6}`

Specifies IP address format to be applied to the subnet. All subnet settings and IP address pools added to the subnet must use the specified address format. You cannot modify the address family once the subnet has been created.

`<prefixlen>`

Sets the prefix length of the subnet. Specify a prefix length appropriate for the selected address family.

`--description <string>`

Specifies an optional description of the subnet. The description cannot exceed 128 bytes.

`--dsr-addr <ip_address>...`

Sets one or more Direct Server Return addresses for the subnet. If an external hardware load balancer that uses DSR addresses is used, this parameter is required.

`--gateway <ip_address>`

Specifies the gateway IP address used by the subnet.

i **NOTE: The IP address must belong to the appropriate gateway. If no gateway is assigned or an incorrect IP address is specified, communication with the cluster might be disabled.**

`--gateway-priority <integer>`

Specifies the gateway priority for the subnet. Valid values start at 1. A lower value has a higher priority—for example, a gateway with priority 3 is given priority over a gateway with priority 7. When a new gateway is configured on the system, it is given a default priority of the current lowest priority plus 10 to ensure it does not take priority over existing gateways until you modify the priority level.

`--mtu <integer>`

Sets the maximum transmission unit (MTU) of the subnet. Common values are 1500 and 9000.

i **NOTE: Using a larger frame size for network traffic permits more efficient communication on the external network between clients and cluster nodes. For example, if a subnet is connected through a 10 GbE interface, we recommend that you set the MTU to 9000. To benefit from using jumbo frames, all devices in the network path must be configured to use jumbo frames.**

`--sc-service-addr <ip_address>`

Specifies the IP address on which the SmartConnect module listens for domain name server (DNS) requests on this subnet.

`--vlan-enabled {true | false}`

Enables or disables VLAN tagging on the subnet.

`--vlan-id <integer>`

Specifies the VLAN ID for all interfaces in the subnet.

`{--verbose | -v}`

Displays more detailed information.

isi network subnets delete

Deletes a subnet. Clients connected to the cluster through a pool in the subnet might lose their connection when the subnet is deleted.

Deleting a subnet that is in use can prevent access to the cluster. Client connections to the cluster through any IP address pool in the deleted subnet will be terminated.

Syntax

```
isi network subnets delete <id>
```

Options

<id>...

Specifies the ID of the subnet to be deleted. Specify the subnet ID in the following format:

```
<groupnet_name>.<subnet_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi network subnets list

Displays available subnets.

Syntax

```
isi network subnets list
  [--groupnet-id <string>]
  [--groupnet <string>]
  [--limit | -l] <integer>]
  [--sort {id | name | addr_family | base_addr | description | gateway | gateway_priority |
mtu | prefixlen | sc_service_addr | vlan_enabled | vlan_id}]
  [--descending | -d]
  [--format {true | table | json | csv | list}]
  [--no-header | -a]
  [--no-footer | -z]
  [--verbose]
```

Options

If no options are specified, the command displays a list of all subnets on the cluster.

--groupnet-id <string>

Displays subnets only from the specified groupnet ID.

--groupnet <string>

Displays subnets only from the specified groupnet ID.

{ --limit | -l} <integer>

Displays no more than the specified number.

--sort {id | description | iface | node_type | name}

Sorts output displayed by the specified attribute.

{ --descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

- `{ --no-header | -a }`
Displays table and CSV output without headers.
- `{ --no-footer | -z }`
Displays table output without footers.
- `{ --verbose | -v }`
Displays more detailed information.

Examples

The following command displays a list of all subnets:

```
isi networks list subnets
```

The system displays output similar to the following example:

```
ID                Subnet                Gateway|Prio  Pools SC Service
-----
groupnet0.subnet0 10.7.135.0/24 10.7.135.1|10 pool0 0.0.0.0
groupnet3.subnet3 10.7.136.0/24 0.0.0.0|20   pool3 0.0.0.0
-----
Total: 2
```

isi network subnets modify

Modifies network subnet settings.

Syntax

```
isi network subnets modify <id>
  [--description <string>|--clear-description]
  [--dsr-addr <ip_address>|--clear-dsr-addr
  | --add-dsr-addr <ip_address>|--remove-dsr-addr
  <ip_address>]
  [--revert-dsr-address]
  [--gateway <ip-address>]
  [--gateway-priority <integer>]
  [--mtu <integer>]
  [--revert-mtu]
  [--prefixlen <prefixlen>]
  [--name <string>]
  [--sc-service-addr <ip_address_range>
  | --clear-sc-service-addr|--add-sc-service-addr
  <ip_address_range> | --remove-sc-service-addr <ip_address_range>]
  [--sc-service-name <domain_name>]
  [--vlan-enabled <boolean>]
  [--revert-vlan-enabled]
  [--vlan-id <integer>]
  [{--force | -f}]
  [{--verbose | -v}]
  [{--help | -h}]
```

Options

`<id>`

Specifies the ID of the subnet that you want to modify. A subnet ID consists of a `<groupnet_id>`, followed by a `!`, followed by a subnet name. A subnet name must be unique throughout the cluster, and can be up to 32

characters long. Supported characters are [a-zA-Z0-9-]. The <groupnet_id> and '!' may be left off if the subnet is in the default groupnet.

```
A subnet in a groupnet: "example_groupnet:example_subnet"  
A subnet in the default groupnet: "example_subnet"
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1.

--description <string>

Specifies an optional description of the subnet. The description cannot exceed 128 bytes.

--clear-description

Clears the description of the subnet.

--dsr-addr <ip_address>...

Specifies a list of IP addresses that use Direct Server Return (DSR). You need to specify `--dsr-addr` for each additional IP address.

--clear-dsr-addr

Clears list of IP addresses that use Direct Server Return.

--add-dsr-addr <ip_address>...

Adds items to the list of IP addresses that use Direct Server Return. You need to specify `--add-dsr-addr` for each additional IP address that you want to add.

--remove-dsr-addr <ip_address>...

Removes items from the list of IP addresses that use Direct Server Return. You need to specify `--remove-dsr-addr` for each additional IP address that you want to remove.

--revert-dsr-addr

Sets the value of `--dsr-addr` to the system default value.

--gateway <ip_address>

Specifies the gateway IP address used by the subnet.

--gateway-priority <integer>

Specifies priority of the subnet gateway, where the lowest number is of the highest priority.

--mtu <integer>

Sets the size of the maximum transmission unit (MTU) of the subnet that the cluster uses in network communication.

--revert-mtu

Sets the value of `--mtu` to the system default value.

--prefixlen <ip_prefixlen>

Sets the prefix size of the subnet.

--name <string>

Specifies the name for the subnet.

--sc-service-addr <ip_address_range>...

Specifies a range of IP addresses in the subnet which will receive incoming DNS requests. You need to specify `--sc-service-addr` for each additional IP address.

--clear-sc-service-addr

Clears the list of SmartConnect Service IP addresses.

--add-sc-service-addr <ip_address_range>...

Adds the items to the list of SmartConnect Service IP addresses. You need to specify `--add-sc-service-addr` for each additional IP address you want to add.

--remove-sc-service-addr <ip_address_range>...

Removes the items from the list of SmartConnect Service IP addresses. You need to specify `--remove-sc-service-addr` for each additional IP address you want to remove.

--sc-service-name <domain_name>

Specifies the domain name corresponding to the SmartConnect Service IP address.

--vlan-enabled <boolean>

Enables VLAN tagging on the subnet.

--revert-vlan-enabled

Sets the value of `--vlan-enabled` to the system default value.

--vlan-id <integer>

Specifies the VLAN ID for all interfaces on this subnet.

{--force | -f }

Suppresses any prompts or warnings messages that would otherwise appear before or during the subnet modification operation.

{--verbose | -v}

Displays more detailed information.

{--help | -h}

Displays help for this command.

isi network subnets view

Displays the configuration details of a specific subnet on the cluster.

Syntax

```
isi network subnets view <id>
```

Options

<id>

Specifies the ID of the subnet to be viewed. Specify the subnet ID in the following format:

```
<groupnet_name>.<subnet_name>
```

The groupnet name is optional if referring to the default groupnet0. Colons are also acceptable as delimiters between component names—for example, groupnet0:subnet1.

isi nfs aliases create

Creates an NFS alias.

Syntax

```
isi nfs aliases create <name> <path>  
  [--zone <string>]  
  [--force]  
  [--verbose]
```

Options

<name>

The name of the alias. Alias names must be formed as Unix root directory with a single forward slash followed by the name. For example, /home.

<path>

The OneFS directory pathname the alias links to. The pathname must be an absolute path below the access zone root. For example, /ifs/data/ugroup1/home.

--zone

The access zone in which the alias is active.

{--force | -f}

Forces creation of the alias without requiring confirmation.

{--verbose | -v}

Displays more detailed information.

Example

The following command creates an alias in a zone named ugroup1:

```
isi nfs aliases create /home /ifs/data/ugroup1/home
--zone ugroup1
```

isi nfs aliases delete

Deletes an NFS alias.

Syntax

```
isi nfs aliases delete <name>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<name>

The name of the alias to be deleted.

--zone <string>

The access zone in which the alias is active.

{--force | -f}

Forces the alias to be deleted without requiring confirmation.

{--verbose | -v}

Displays more detailed information.

Example

The following command deletes an alias from a zone named ugroup1.

```
isi nfs aliases delete /projects --zone ugroup1
```

isi nfs aliases list

Lists NFS aliases available in the current access zone.

Syntax

```
isi nfs aliases list
  [--check]
  [--zone <string>]
```

```
[--limit <integer>]
[--sort {zone | name | path | health}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
```

Options

--check

For the current zone, displays a list of aliases and their health status.

--zone <string>

The access zone in which the alias is active.

{--limit | -l} <integer>

Displays no more than the specified number of NFS aliases.

--sort {zone | name | path | health}

Specifies the field to sort by.

{--descending | -d}

Specifies to sort the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

Example

The following command displays a table of the aliases in a zone named `ugroup1` including their health status.

```
isi nfs aliases list --zone ugroup1 --check
```

Output from the command is similar to the following example:

Zone	Name	Path	Health
ugroup1	/home	/ifs/data/offices/newyork	good
ugroup1	/root_alias	/ifs/data/offices	good
ugroup1	/project	/ifs/data/offices/project	good

Total: 3

isi nfs aliases modify

Modifies the name, zone, or absolute path of an alias.

Syntax

```
isi nfs aliases modify <alias>
  [--zone <string>]
  [--new-zone <string>]
  [--name <string>]
  [--path <path>]
```

```
[--force]
[--verbose]
```

Options

<alias>

The current name of the alias, for example, /home.

--zone <string>

The access zone in which the alias is currently active.

--new-zone <string>

The new access zone in which the alias is to be active.

--name <string>

A new name for the alias.

--path <path>

The new OneFS directory pathname the alias should link to. The pathname must be an absolute path below the access zone root. For example, /ifs/data/ugroup2/home.

{--force | -f}

Forces modification of the alias without requiring confirmation.

{--verbose | -v}

Displays more detailed information.

Example

The following command modifies the zone, name, and path of an existing alias:

```
isi nfs aliases modify /home --name /users --zone ugroup1 --new-zone ugroup2
--path /ifs/data/ugroup2/users
```

isi nfs aliases view

Shows information about an alias in the current zone.

Syntax

```
isi nfs aliases view <name>
  [--zone <string>]
  [--check]
```

Options

<name>

The name of the alias.

--zone <string>

The access zone in which the alias is active.

--check

Include the health status of the alias.

Example

The following command displays a table of information, including the health status, of an alias named `/projects` in the current zone.

```
isi nfs aliases view /projects --check
```

isi nfs exports check

Checks NFS exports for configuration errors, including conflicting export rules, bad paths, unresolvable host names, and unresolvable net groups.

Syntax

```
isi nfs exports check
  [--limit <integer>]
  [--zone <string>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--ignore-unresolvable-hosts]
  [--ignore-bad-paths]
  [--ignore-bad-auth]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of NFS exports.

--zone <string>

Specifies the access zone in which the export was created.

{--format {table | json | csv | list}}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

--ignore-unresolvable-hosts

Does not present an error condition on unresolvable hosts when creating or modifying an export.

--ignore-bad-paths

Does not present an error condition on bad paths when creating or modifying an export.

--ignore-bad-auth

Ignores bad authentication for mapping options when creating or modifying an export.

{--verbose | -v}

Displays more detailed information.

Examples

The following command checks the exports in a zone named **Zone-1**:

```
isi nfs exports check --zone Zone-1
```

If the check finds no problems, it returns an empty table. If, however, the check finds a problem, it returns a display similar to the following:

```
ID    Message
-----
3     '/ifs/data/project' does not exist
-----
Total: 1
```

isi nfs exports create

Creates an NFS export.

NOTE: To view the default NFS export settings that will be applied when creating an export, run the `isi nfs settings export view` command.

Syntax

```
isi nfs exports create <paths>
  [--block-size <size>]
  [--can-set-time {yes | no}]
  [--case-insensitive {yes | no}]
  [--case-preserving {yes | no}]
  [--chown-restricted {yes | no}]
  [--directory-transfer-size <size>]
  [--link-max <integer>]
  [--max-file-size <size>]
  [--name-max-size <integer>]
  [--no-truncate {yes | no}]
  [--return-32bit-file-ids {yes | no}]
  [--symlinks {yes | no}]
  [--zone <string>]
  [--clients <client>]
  [--description <string>]
  [--root-clients <client>]
  [--read-write-clients <client>]
  [--read-only-clients <client>]
  [--all-dirs {yes | no}]
  [--encoding <string>]
  [--security-flavors {unix | krb5 | krb5i | krb5p}]
  [--snapshot <snapshot>]
  [--map-lookup-uid {yes | no}]
  [--map-retry {yes | no}]
  [--map-root-enabled {yes | no}]
  [--map-non-root-enabled {yes | no}]
  [--map-failure-enabled {yes | no}]
  [--map-all <identity>]
  [--map-root <identity>]
  [--map-non-root <identity>]
  [--map-failure <identity>]
  [--map-full {yes | no}]
  [--commit-asynchronous {yes | no}]
  [--read-only {yes | no}]
  [--readdirplus {yes | no}]
  [--read-transfer-max-size <size>]
  [--read-transfer-multiple <integer>]
  [--read-transfer-size <size>]
  [--setattr-asynchronous {yes | no}]
  [--time-delta <time delta>]
  [--write-datasync-action {datasync | filesync | unstable}]
  [--write-datasync-reply {datasync | filesync}]
  [--write-filesync-action {datasync | filesync | unstable}]
  [--write-filesync-reply filesync]
  [--write-unstable-action {datasync | filesync | unstable}]
  [--write-unstable-reply {datasync | filesync | unstable}]
  [--write-transfer-max-size <size>]
  [--write-transfer-multiple <integer>]
  [--write-transfer-size <size>]
  [--ignore-unresolvable-hosts]
```



```
[--ignore-bad-paths]
[--ignore-bad-auth]
[--ignore-conflicts]
[--force]
[--verbose]
```

Options

<paths> ...

Required. Specifies the path to be exported, starting at /ifs. This option can be repeated to specify multiple paths.

--block-size <size>

Specifies the block size, in bytes.

--can-set-time {yes | no}

If set to yes, enables the export to set time. The default setting is no.

--case-insensitive {yes | no}

If set to yes, the server will report that it ignores case for file names. The default setting is no.

--case-preserving {yes | no}

If set to yes, the server will report that it always preserves case for file names. The default setting is no.

--chown-restricted {yes | no}

If set to yes, the server will report that only the superuser can change file ownership. The default setting is no.

--directory-transfer-size <size>

Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: b for bytes; K for KB; M for MB; or G for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128K.

--link-max <integer>

The reported maximum number of links to a file.

--max-file-size <size>

Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

--name-max-size <integer>

The reported maximum length of characters in a filename.

--no-truncate {yes | no}

If set to yes, too-long file names will result in an error rather than be truncated.

--return-32bit-file-ids {yes | no}

Applies to NFSv3 and NFSv4. If set to yes, limits the size of file identifiers returned from readdir to 32-bit values. The default value is no.



NOTE: This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

--symlinks {yes | no}

If set to yes, advertises support for symlinks. The default setting is no.

--zone <string>

Access zone in which the export should apply. The default zone is system.

--clients <client>

Specifies a client to be allowed access through this export. Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.



NOTE: This option replaces the entire list of clients. To add or remove a client from the list, specify --add-clients or --remove-clients.

--description <string>

The description for this NFS export.

--root-clients <client>

Allows the root user of the specified client to execute operations as the root user of the cluster. This option overrides the --map-all and --map-root option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--read-write-clients <client>

Grants read/write privileges to the specified client for this export. This option overrides the --read-only option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--read-only-clients <client>

Makes the specified client read-only for this export. This option overrides the --read-only option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--all-dirs {yes | no}

If set to yes, this export will cover all directories. The default setting is no.

--encoding <string>

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's /etc/encodings.xml file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)
iso-8859-7	ISO-8859-7 (Greek)
iso-8859-8	ISO-8859-8 (Hebrew)
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859- 14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)

Value	Encoding
iso-8859-16	ISO-8859-16 (Latin-10)

--security-flavors {unix | krb5 | krb5i | krb5p}

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

unix	UNIX (system) authentication.
krb5	Kerberos V5 authentication.
krb5i	Kerberos V5 authentication with integrity.
krb5p	Kerberos V5 authentication with privacy.

--snapshot {<snapshot> | <snapshot-alias>}

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify <snapshot> or <snapshot-alias> as the ID or name of a snapshot or snapshot alias.

--map-lookup-uid {yes | no}

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

--map-retry {yes | no}

If set to *yes*, the system retries failed user-mapping lookups. The default setting is *no*.

--map-root-enabled {yes | no}

Enable/disable mapping incoming root users to a specific account.

--map-non-root-enabled {yes | no}

Enable/disable mapping incoming non-root users to a specific account.

--map-failure-enabled {yes | no}

Enable/disable mapping users to a specific account after failing an auth lookup.

--map-all <identity>

Specifies the default identity that operations by any user will execute as. If this option is not set to *root*, you can allow the root user of a specific client to execute operations as the root user of the cluster by including the client in the *--root-clients* list.

--map-root <identity>

Map incoming root users to a specific user and/or group ID.

--map-non-root <identity>

Map non-root users to a specific user and/or group ID.

--map-failure <identity>

Map users to a specific user and/or group ID after a failed auth attempt.

--map-full {yes | no}

Determines how user mapping is accomplished if a user is specified in an export option such as *--map-root* or *--map-all*. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.

The default setting is *yes*.

--commit-asynchronous {yes | no}

If set to *yes*, enables commit data operations to be performed asynchronously. The default setting is *no*.

--read-only {yes | no}

Determines the default privileges for all clients accessing the export.

If set to `yes`, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.

If set to `no`, you can make a specific client read-only by including the client in the `--read-only-clients` list. The default setting is `no`.

--readdirplus {yes | no}

Applies to NFSv3 only. If set to `yes`, enables processing of readdir-plus requests. The default setting is `yes`.

--read-transfer-max-size <size>

Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 1M.

--read-transfer-multiple <integer>

Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are 0–4294967295. The initial default value is **512**.

--read-transfer-size <size>

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b, or lower if the `--read-transfer-max-size` is set to a lesser value. The initial default value is 128K.

--setattr-asynchronous {yes | no}

If set to `yes`, performs set-attributes operations asynchronously. The default setting is `no`.

--time-delta <float>

Specifies server time granularity, in seconds.

--write-datasync-action {datasync | filesync | unstable}

Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `datasync`, which performs the request as specified.

--write-datasync-reply {datasync | filesync}

Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync reply method. The following values are valid:

- `datasync`
- `filesync`

The default value is `datasync` (does not respond differently).

--write-filesync-action {datasync | filesync | unstable}

Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `filesync`, which performs the request as specified.

--write-filesync-reply {filesync}

Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync reply method. The only valid value is `filesync` (does not respond differently).

--write-unstable-action {datasync | filesync | unstable}

Specifies an alternate unstable-write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable`, which performs the request as specified.

--write-unstable-reply {datasync | filesync | unstable}

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

--write-transfer-max-size <size>

Specifies the preferred maximum write transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `k` for KB; `m` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`. The initial default value is `1M`.

--write-transfer-multiple <integer>

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are 0–4294967295. The initial default value is `512`.

--write-transfer-size <size>

Specifies the preferred write transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `k` for KB; `m` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is `4294967295b`, or lower if the `--write-transfer-max-size` is set to a lesser value. The initial default value is `512k`.

--ignore-unresolvable-hosts

Does not present an error condition on unresolvable hosts when creating or modifying an export.

--ignore-bad-paths

Does not present an error condition on bad paths when creating or modifying an export.

--ignore-bad-auth

Ignores bad authentication for mapping options when creating or modifying an export.

--ignore-conflicts

Ignores conflicts between the new or modified exports and the existing configuration.

{--force | -f}

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command executes without prompting for confirmation.

{--verbose | -v}

Displays more detailed information.

Examples

The following command creates an NFS export for a particular zone and set of clients:

```
isi nfs exports create /ifs/data/ugroup1/home
--description 'Access to home dirs for user group 1'
--zone ugroup1 --clients 10.1.28.1 --clients 10.1.28.2
```

The following command creates an NFS export with multiple directory paths and a custom security type (Kerberos 5):

```
isi nfs exports create /ifs/data/projects /ifs/data/templates
--security-flavors krb5
```

isi nfs exports delete

Deletes an NFS export.

Syntax

```
isi nfs exports delete <id>
[--zone <string>]
```

```
[--force]
[--verbose]
```

Options

<id>

Specifies the ID of the NFS export to delete. You can use the `isi nfs exports list` command to view a list of exports and their IDs in the current zone.

--zone <string>

Specifies the access zone in which the export was created. The default is the current zone.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

isi nfs exports list

Displays a list of NFS exports.

Syntax

```
isi nfs exports list
  [--zone <string>]
  [--limit <integer>]
  [--sort <field>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the name of the access zone in which the export was created.

{--limit | -l} <integer>

Displays no more than the specified number of NFS exports.

--sort <field>

Specifies the field to sort by. Valid values are as follows:

- id
- zone
- paths
- description
- clients
- root_clients
- read_only_clients
- read_write_clients
- unresolved_clients
- all_dirs
- block_size
- can_set_time

- commit_asynchronous
- directory_transfer_size
- encoding
- map_lookup_uid
- map_retry
- map_all
- map_root
- map_full
- max_file_size
- read_only
- readdirplus
- return_32bit_file_ids
- read_transfer_max_size
- read_transfer_multiple
- read_transfer_size
- security_flavors
- setattr_asynchronous
- symlinks
- time_delta
- write_datasync_action
- write_datasync_reply
- write_filesync_action
- write_filesync_reply
- write_unstable_action
- write_unstable_reply
- write_transfer_max_size
- write_transfer_multiple
- write_transfer_size

--descending

Specifies to sort the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

The following command lists NFS exports, by default in the current zone:

```
isi nfs exports list
```

The following command lists NFS exports in a specific zone:

```
isi nfs exports list --zone hq-home
```

isi nfs exports modify

Modifies an NFS export.

 **NOTE:** You can run the `isi nfs settings export view` command to see the full list of default settings for exports.

Syntax

```
isi nfs exports modify <id>
[--block-size <size>]
[--revert-block-size]
[--can-set-time {yes | no}]
[--revert-can-set-time]
[--case-insensitive {yes | no}]
[--revert-case-insensitive]
[--case-preserving {yes | no}]
[--revert-case-preserving]
[--chown-restricted {yes | no}]
[--revert-chown-restricted]
[--directory-transfer-size <size>]
[--revert-directory-transfer-size]
[--link-max <integer>]
[--revert-link-max]
[--max-file-size <size>]
[--revert-max-file-size]
[--name-max-size <integer>]
[--revert-name-max-size]
[--no-truncate {yes | no}]
[--revert-no-truncate]
[--return-32bit-file-ids {yes | no}]
[--revert-return-32bit-file-ids]
[--symlinks {yes | no}]
[--revert-symlinks]
[--new-zone <string>]
[--description <string>]
[--paths <path>]
[--clear-paths]
[--add-paths <string>]
[--remove-paths <string>]
[--clients <string>]
[--clear-clients]
[--add-clients <string>]
[--remove-clients <string>]
[--root-clients <string>]
[--clear-root-clients]
[--add-root-clients <string>]
[--remove-root-clients <string>]
[--read-write-clients <string>]
[--clear-read-write-clients]
[--add-read-write-clients <string>]
[--remove-read-write-clients <string>]
[--read-only-clients <string>]
[--clear-read-only-clients]
[--add-read-only-clients <string>]
[--remove-read-only-clients <string>]
[--all-dirs {yes | no}]
[--revert-all-dirs]
[--encoding <string>]
[--revert-encoding]
[--security-flavors {unix | krb5 | krb5i | krb5p}]
[--revert-security-flavors]
[--clear-security-flavors]
[--add-security-flavors {unix | krb5 | krb5i | krb5p}]
[--remove-security-flavors <string>]
[--snapshot <snapshot>]
[--revert-snapshot]
[--map-lookup-uid {yes | no}]
[--revert-map-lookup-uid]
[--map-retry {yes | no}]
```



```

[--revert-map-retry]
[--map-root-enabled {yes | no}]
[--revert-map-root-enabled]
[--map-non-root-enabled {yes | no}]
[--revert-map-non-root-enabled]
[--map-failure-enabled {yes | no}]
[--revert-map-failure-enabled]
[--map-all <identity>]
[--revert-map-all]
[--map-root <identity>]
[--revert-map-root]
[--map-non-root <identity>]
[--revert-map-non-root]
[--map-failure <identity>]
[--revert-map-failure]
[--map-full {yes | no}]
[--revert-map-full]
[--commit-asynchronous {yes | no}]
[--revert-commit-asynchronous]
[--read-only {yes | no}]
[--revert-read-only]
[--readdirplus {yes | no}]
[--revert-readdirplus]
[--read-transfer-max-size <size>]
[--revert-read-transfer-max-size]
[--read-transfer-multiple <integer>]
[--revert-read-transfer-multiple]
[--read-transfer-size <size>]
[--revert-read-transfer-size]
[--setattr-asynchronous {yes | no}]
[--revert-setattr-asynchronous]
[--time-delta <time delta>]
[--revert-time-delta]
[--write-datasync-action {datasync | filesync | unstable}]
[--revert-write-datasync-action]
[--write-datasync-reply {datasync | filesync}]
[--revert-write-datasync-reply]
[--write-filesync-action {datasync | filesync | unstable}]
[--revert-write-filesync-action]
[--write-filesync-reply filesync]
[--write-unstable-action {datasync | filesync | unstable}]
[--revert-write-unstable-action]
[--write-unstable-reply {datasync | filesync | unstable}]
[--revert-write-unstable-reply]
[--write-transfer-max-size <size>]
[--revert-write-transfer-max-size]
[--write-transfer-multiple <integer>]
[--revert-write-transfer-multiple]
[--write-transfer-size <size>]
[--revert-write-transfer-size]
[--zone <string>]
[--ignore-unresolvable-hosts]
[--ignore-bad-paths]
[--ignore-bad-auth]
[--ignore-conflicts]
[--force]
[--verbose]

```

Options

<id>

The export ID number. You can use the `isi nfs exports list` command to view all the exports and their ID numbers in the current access zone.

--block-size <size>

Specifies the block size, in bytes.

--revert-block-size

Restores the setting to the system default.

--can-set-time {yes | no}

If set to `yes`, enables the export to set time. The default setting is `no`.

--revert-can-set-time

Restores the setting to the system default.

--case-insensitive {yes | no}

If set to `yes`, the server will report that it ignores case for file names. The default setting is `no`.

--revert-case-insensitive

Restores the setting to the system default.

--case-preserving {yes | no}

If set to `yes`, the server will report that it always preserves case for file names. The default setting is `no`.

--revert-case-preserving

Restores the setting to the system default.

--chown-restricted {yes | no}

If set to `yes`, the server will report that only the superuser can change file ownership. The default setting is `no`.

--revert-chown-restricted

Restores the setting to the system default.

--directory-transfer-size <size>

Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128K.

--revert-directory-transfer-size

Restores the setting to the system default.

--link-max <integer>

The reported maximum number of links to a file.

--revert-link-max

Restores the setting to the system default.

--max-file-size <size>

Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

--revert-max-file-size

Restores the setting to the system default.

--name-max-size <integer>

The reported maximum length of characters in a filename.

--revert-name-max-size

Restores the setting to the system default.

--no-truncate {yes | no}

If set to `yes`, too-long file names will result in an error rather than be truncated.

--revert-no-truncate

Restores the setting to the system default.

--return-32bit-file-ids {yes | no}

Applies to NFSv3 and later. If set to `yes`, limits the size of file identifiers returned from `readdir` to 32-bit values. The default value is `no`.



NOTE: This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

--revert-return-32bit-file-ids

Restores the setting to the system default.

--symlinks {yes | no}

If set to `yes`, advertises support for symlinks. The default setting is `no`.

--revert-symlinks

Restores the setting to the system default.

--new-zone <string>
 Specifies a new access zone in which the export should apply. The default zone is `system`.

--description <string>
 The description for this NFS export.

--paths <paths> ...
 Required. Specifies the path to be exported, starting at `/ifs`. This option can be repeated to specify multiple paths.

--clear-paths
 Clear any of the paths originally specified for the export. The path must be within the `/ifs` directory.

--add-paths <paths> ...
 Add to the paths originally specified for the export. The path must be within `/ifs`. This option can be repeated to specify multiple paths.

--remove-paths <paths> ...
 Remove a path from the paths originally specified for the export. The path must be within `/ifs`. This option can be repeated to specify multiple paths to be removed.

--clients <string>
 Specifies a client to be allowed access through this export. Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.

--clear-clients
 Clear the full list of clients originally allowed access through this export.

--add-clients <string>
 Specifies a client to be added to the list of clients with access through this export. Specify clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple clients by repeating this option.

--remove-clients <string>
 Specifies a client to be removed from the list of clients with access through this export. Specify clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple clients by repeating this option.

--root-clients <string>
 Allows the root user of the specified client to execute operations as the root user of the cluster. This option overrides the `--map-all` and `--map-root` option for the specified client.
 Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--clear-root-clients
 Clear the full list of root clients originally allowed access through this export.

--add-root-clients <string>
 Specifies a root client to be added to the list of root clients with access through this export. Specify root clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple root clients by repeating this option.

--remove-root-clients <string>
 Specifies a root client to be removed from the list of root clients with access through this export. Specify root clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple root clients by repeating this option.

--read-write-clients <string>
 Grants read/write privileges to the specified client for this export. This option overrides the `--read-only` option for the specified client.
 Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--clear-read-write-clients
 Clear the full list of read-write clients originally allowed access through this export.

--add-read-write-clients <string>

Specifies a read-write client to be added to the list of read-write clients with access through this export. Specify read-write clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple read-write clients by repeating this option.

--remove-read-write-clients <string>

Specifies a read-write client to be removed from the list of read-write clients with access through this export. Specify read-write clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple read-write clients by repeating this option.

--read-only-clients <string>

Makes the specified client read-only for this export. This option overrides the `--read-only` option for the specified client.

Specify clients as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can specify multiple clients in a comma-separated list.

--clear-read-only-clients

Clear the full list of read-only clients originally allowed access through this export.

--add-read-only-clients <string>

Specifies a read-only client to be added to the list of read-only clients with access through this export. Specify read-only clients to be added as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can add multiple read-only clients by repeating this option.

--remove-read-only-clients <string>

Specifies a read-only client to be removed from the list of read-only clients with access through this export. Specify read-only clients to be removed as an IPv4 or IPv6 address, hostname, netgroup, or CIDR range. You can remove multiple read-only clients by repeating this option.

--all-dirs {yes | no}

If set to `yes`, this export will cover all directories. The default setting is `no`.

--revert-all-dirs

Restores the setting to the system default.

--encoding <string>

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's `/etc/encodings.xml` file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)
iso-8859-7	ISO-8859-7 (Greek)

Value	Encoding
iso-8859-8	ISO-8859-8 (Hebrew)
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859-14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)
iso-8859-16	ISO-8859-16 (Latin-10)

--revert-encoding

Restores the setting to the system default.

--security-flavors {unix | krb5 | krb5i | krb5p}

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

unix	UNIX (system) authentication.
krb5	Kerberos V5 authentication.
krb5i	Kerberos V5 authentication with integrity.
krb5p	Kerberos V5 authentication with privacy.

--revert-security-flavors

Restores the setting to the system default.

--clear-security-flavors

Clears the value for supported security flavors.

--add-security-flavors {unix | krb5 | krb5i | krb5p}

Adds supported security flavors. Repeat for each additional supported security flavor to add.

--remove-security-flavors

Removes supported security flavors. Repeat for each additional supported security flavor to remove from the list.

--snapshot {<snapshot> | <snapshot-alias>}

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify <snapshot> or <snapshot-alias> as the ID or name of a snapshot or snapshot alias.

--revert-snapshot

Restores the setting to the system default.

--map-lookup-uid {yes | no}

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

--revert-map-lookup-uid

Restores the setting to the system default.

--map-retry {yes | no}

If set to *yes*, the system will retry failed user-mapping lookups. The default setting is *no*.

--revert-map-retry

Restores the setting to the system default.

--map-root-enabled {yes | no}

Enable/disable mapping incoming root users to a specific account.

--revert-map-root-enabled
Restores the setting to the system default.

--map-non-root-enabled {yes | no}
Enable/disable mapping incoming non-root users to a specific account.

--revert-map-non-root-enabled
Restores the setting to the system default.

--map-failure-enabled {yes | no}
Enable/disable mapping users to a specific account after failing an auth lookup.

--revert-map-failure-enabled
Restores the setting to the system default.

--map-all <identity>
Specifies the default identity that operations by any user will execute as. If this option is not set to `root`, you can allow the root user of a specific client to execute operations as the root user of the cluster by including the client in the `--root-clients` list.

--revert-map-all
Restores the setting to the system default.

--map-root <identity>
Map incoming root users to a specific user and/or group ID.

--revert-map-root
Restores the setting to the system default.

--map-non-root <identity>
Map non-root users to a specific user and/or group ID.

--revert-map-non-root
Restores the setting to the system default.

--map-failure <identity>
Map users to a specific user and/or group ID after a failed auth attempt.

--revert-map-failure
Restores the setting to the system default.

--map-full {yes | no}
Determines how user mapping is accomplished if a user is specified in an export option such as `--map-root` or `--map-all`. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.
The default setting is `yes`.

--revert-map-full
Restores the `--map-full` setting to the system default, `yes`.

--commit-asynchronous {yes | no}
If set to `yes`, enables commit data operations to be performed asynchronously. The default setting is `no`.

--revert-commit-asynchronous
Restores the setting to the system default.

--read-only {yes | no}
Determines the default privileges for all clients accessing the export.
If set to `yes`, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.
If set to `no`, you can make a specific client read-only by including the client in the `--read-only-clients` list.
The default setting is `no`.

--revert-read-only
Restores the setting to the system default.

--readdirplus {yes | no}

Applies to NFSv3 only. If set to *yes*, enables processing of readdir-plus requests. The default setting is *yes*.

--revert-readdirplus

Restores the setting to the system default.

--read-transfer-max-size <size>

Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: *b* for bytes; *k* for KB; *m* for MB; or *G* for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 1M.

--revert-read-transfer-max-size

Restores the setting to the system default.

--read-transfer-multiple <integer>

Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are 0-4294967295. The initial default value is **512**.

--revert-read-transfer-multiple

Restores the setting to the system default.

--read-transfer-size <size>

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: *b* for bytes; *k* for KB; *m* for MB; or *G* for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b, or lower if the `--read-transfer-max-size` is set to a lesser value. The initial default value is 128k.

--revert-read-transfer-size

Restores the setting to the system default.

--setattr-asynchronous {yes | no}

If set to *yes*, performs set-attributes operations asynchronously. The default setting is *no*.

--revert-setattr-asynchronous

Restores the setting to the system default.

--time-delta <float>

Specifies server time granularity, in seconds.

--revert-time-delta

Restores the setting to the system default.

--write-datasync-action {datasync | filesync | unstable}

Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync write method. The following values are valid:

- *datasync*
- *filesync*
- *unstable*

The default value is *datasync*, which performs the request as specified.

--revert-write-datasync-action

Restores the setting to the system default.

--write-datasync-reply {datasync | filesync}

Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync reply method. The following values are valid:

- *datasync*
- *filesync*

The default value is *datasync* (does not respond differently).

--revert-write-datasync-reply

Restores the setting to the system default.

--write-filesync-action {datasync | filesync | unstable}

Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync write method. The following values are valid:

- *datasync*
- *filesync*

- `unstable`

The default value is `filesync`, which performs the request as specified.

--revert-write-filesync-action

Restores the setting to the system default.

--write-filesync-reply {filesync}

Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync reply method. The only valid value is `filesync` (does not respond differently).

--write-unstable-action {datasync | filesync | unstable}

Specifies an alternate unstable-write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable`, which performs the request as specified.

--revert-write-unstable-action

Restores the setting to the system default.

--write-unstable-reply {datasync | filesync | unstable}

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

--revert-write-unstable-reply

Restores the setting to the system default.

--write-transfer-max-size <size>

Specifies the preferred maximum write transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 1M.

--revert-write-transfer-max-size

Restores the setting to the system default.

--write-transfer-multiple <integer>

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are 0-4294967295. The initial default value is 512.

--revert-write-transfer-multiple

Restores the setting to the system default.

--write-transfer-size <size>

Specifies the preferred write transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b, or lower if the `--write-transfer-max-size` is set to a lesser value. The initial default value is 512K.

--revert-write-transfer-size

Restores the setting to the system default.

--zone

Access zone in which the export was originally created.

--ignore-unresolvable-hosts

Does not present an error condition on unresolvable hosts when creating or modifying an export.

--ignore-bad-paths

Does not present an error condition on bad paths when creating or modifying an export.

--ignore-bad-auth

Ignores bad authentication for mapping options when creating or modifying an export.

--ignore-conflicts

Ignores conflicts between the new or modified exports and the existing configuration.

{--force | -f}

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command executes without prompting for confirmation.

{--verbose | -v}

Displays more detailed information.

isi nfs exports reload

Reloads NFS export configurations.

Syntax

```
isi nfs exports reload  
  [--zone <string>
```

Options

--zone

The access zone for the exports you are reloading.

isi nfs exports view

View an NFS export.

Syntax

```
isi nfs exports view <id>  
  [--zone <string>]
```

Options

<id>

Specifies the ID of the NFS export to display. If you do not know the ID, use the `isi nfs exports list` command to view a list of exports and their associated IDs.

--zone <string>

Specifies the name of the access zone in which the export was created.

isi nfs log-level modify

Sets the logging level for the NFS service.

Syntax

```
isi nfs log-level modify <level>  
  [--verbose]
```

Options

<level>

Valid logging levels are:

Log level	Description
always	Specifies that all NFS events are logged in NFS log files.
error	Specifies that only NFS error conditions are logged in NFS log files.
warning	Specifies that only NFS warning conditions are logged in NFS log files.
info	Specifies that only NFS information conditions are logged in NFS log files.
verbose	Specifies verbose logging.
debug	Adds information that we can use to troubleshoot issues
trace	Adds tracing information that we can use to pinpoint issues

{--verbose | -v}

Displays more detailed information.

isi nfs log-level view

Shows the logging level for the NFS service.

Syntax

```
isi nfs log-level view
```

Options

There are no options for this command.

isi nfs netgroup check

Updates the NFS netgroup cache.

Syntax

```
isi nfs netgroup check  
  [--host <string>]  
  [--verbose]
```

Options

--host <string>

The IPv4 or IPv6 address of the node to check. The default is the localhost IP address.

{--verbose | -v}

Displays more detailed information.

isi nfs netgroup flush

Flushes the NFS netgroup cache.

Syntax

```
isi nfs netgroup flush
  [--host <string>]
  [--verbose]
```

Options

--host <string>

The IPv4 or IPv6 address of the node to flush. If you do not specify a node, all nodes are flushed (default).

{verbose | -v}

Displays more detailed information.

isi nfs netgroup modify

Modifies the NFS netgroup cache settings.

Syntax

```
isi nfs netgroup modify
  [--bgwrite <duration>]
  [--expiration <duration>]
  [--lifetime <duration>]
  [--retry <duration>]
  [--verbose]
```

Options

 **NOTE:** In the following option definitions, express duration in integer format as [YMWDHms].

{bgwrite | -w} <duration>

Sets the to-disk backup interval.

{expiration | -e} <duration>

Sets the netgroup expiration time.

{lifetime | -i} <duration>

Sets the netgroup lifetime.

{retry | -r} <duration>

Sets the retry interval.

{verbose | -v}

Displays more detailed information.

isi nfs nlm locks list

Applies to NFSv3 only. Displays a list of NFS Network Lock Manager (NLM) advisory locks.

Syntax

```
isi nfs nlm locks list
  [--limit <integer>]
  [--sort {client | path | lock_type | range | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of NFS nlm locks.

--sort {client | path | lock_type | range | created}

Specifies the field to sort by.

{--descending | -d}

Specifies to sort the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To view a detailed list of all current NLM locks, run the following command:

```
isi nfs nlm locks list --verbose
```

In the following sample output, there are currently three locks: one on `/ifs/home/test1/file.txt` and two on `/ifs/home/test2/file.txt`.

Client	Path	Lock Type	Range
machineName/10.72.134.119	/ifs/home/test1/file.txt	exclusive	[0, 2]
machineName/10.59.166.125	/ifs/home/test2/file.txt	shared	[10, 20]
machineName/10.63.119.205	/ifs/home/test2/file.txt	shared	[10, 20]

isi nfs nlm locks waiters

Displays a list of clients that are waiting to place a Network Lock Manager (NLM) lock on a currently locked file. This command applies to NFSv3 only.

Syntax

```
isi nfs nlm locks waiters
  [--limit <integer>]
  [--sort {client | path | lock_type | range | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l}<integer>

Displays no more than the specified number of NLM locks.

--sort {client | path | lock_type | range | created}

Specifies the field to sort by.

--descending

Specifies to sort the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

The following command displays a detailed list of clients waiting to lock a currently-locked file:

```
isi nfs nlm locks waiters --verbose
```

The system displays output similar to the following example:

Client	Path	Lock Type	Range
-----	-----	-----	-----
machineName/1.2.34.5	/ifs/home/test1/file.txt	exclusive	[0, 2]

isi nfs nlm sessions check

Searches for lost locks.

Syntax

```
isi nfs nlm sessions check
  [--cluster-ip <string>]
  [--zone <string>]
```

Options

- cluster-ip <string>**
The cluster IP address to which the client is connected.
- zone <string>**
The access zone to which the client is connected.

isi nfs nlm sessions delete

Deletes all states associated with an NFS Network Lock Manager (NLM) connection.

Syntax

```
isi nfs nlm sessions delete <hostname> <cluster-ip>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

- <hostname>**
The name of the client that initiated the session.
- <cluster-ip>**
The cluster IP address to which the client is connected.
- zone <string>**
The access zone to which the client is connected.
- {force | -f}**
Skips the confirmation prompt.
- {verbose | -v}**
Displays more detailed information.

isi nfs nlm sessions list

Displays a list of clients holding NFS Network Lock Manager (NLM) locks. This command applies to NFSv3 only.

Syntax

```
isi nfs nlm sessions list
  [--limit <integer>]
  [--sort {ID | client}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

The number of NFS NLM sessions to display.

--sort {hostname | cluster_ip | is_active | notify_attempts_remaining}

Specifies the field to sort by.

{--descending | -d}

Specifies to sort the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

To view a list of active NLM sessions, run the following command:

```
isi nfs nlm sessions list
```

isi nfs nlm sessions refresh

Refreshes an NFS Network Lock Manager (NLM) client.

Syntax

```
isi nfs nlm sessions refresh <hostname> <cluster-ip>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<hostname>

The name of the client that initiated the session.

<cluster-ip>

The cluster IP address to which the client is connected.

--zone <string>

The access zone to which the client is connected.

{--force | -f}

Skips the confirmation prompt.

{--verbose | -v}

Displays more detailed information.

isi nfs nlm sessions view

Displays information about NFS Network Lock Manager (NLM) client connections.

Syntax

```
isi nfs nlm sessions view <hostname>
  [--cluster-ip <string>]
  [--zone <string>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<hostname>

The name of the client that initiated the session.

--cluster-ip <string>

The cluster IP address to which the client is connected.

--zone <string>

The access zone to which the client is connected.

{--limit | -l} <integer>

Displays no more than the specified number of NFS nlm locks.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi nfs settings export modify

Modifies the default settings that are applied when creating NFS exports.

 **NOTE:** You can view the currently configured default NFS export settings by running the `isi nfs settings export view` command.

Syntax

```
isi nfs exports modify <ID>
  [--block-size <size>]
  [--revert-block-size]
  [--can-set-time {yes|no}]
  [--revert-can-set-time]
```



```

[--case-insensitive {yes|no}]
[--revert-case-insensitive]
[--case-preserving {yes|no}]
[--revert-case-preserving]
[--chown-restricted {yes|no}]
[--revert-chown-restricted]
[--directory-transfer-size <size>]
[--revert-directory-transfer-size]
[--link-max <integer>]
[--revert-link-max]
[--max-file-size <size>]
[--revert-max-file-size]
[--name-max-size <integer>]
[--revert-name-max-size]
[--no-truncate {yes|no}]
[--revert-no-truncate]
[--return-32bit-file-ids {yes|no}]
[--revert-return-32bit-file-ids]
[--symlinks {yes|no}]
[--revert-symlinks]
[--all-dirs {yes|no}]
[--revert-all-dirs]
[--encoding <string>]
[--revert-encoding]
[--security-flavors {unix|krb5|krb5i|krb5p}]
[--revert-security-flavors]
[--clear-security-flavors]
[--add-security-flavors {unix|krb5|krb5i|krb5p}]
[--remove-security-flavors <string>]
[--snapshot <snapshot>]
[--revert-snapshot]
[--map-lookup-uid {yes|no}]
[--revert-map-lookup-uid]
[--map-retry {yes|no}]
[--revert-map-retry]
[--map-root-enabled {yes|no}]
[--revert-map-root-enabled]
[--map-non-root-enabled {yes|no}]
[--revert-map-non-root-enabled]
[--map-failure-enabled {yes|no}]
[--revert-map-failure-enabled]
[--map-all <identity>]
[--revert-map-all]
[--map-root <identity>]
[--revert-map-root]
[--map-non-root <identity>]
[--revert-map-non-root]
[--map-failure <identity>]
[--revert-map-failure]
[--map-full {yes|no}]
[--revert-map-full]
[--commit-asynchronous {yes|no}]
[--revert-commit-asynchronous]
[--read-only {yes|no}]
[--revert-read-only]
[--readdirplus {yes|no}]
[--revert-readdirplus]
[--read-transfer-max-size <size>]
[--revert-read-transfer-max-size]
[--read-transfer-multiple <integer>]
[--revert-read-transfer-multiple]
[--read-transfer-size <size>]
[--revert-read-transfer-size]
[--setattr-asynchronous {yes|no}]
[--revert-setattr-asynchronous]
[--time-delta <integer>]
[--revert-time-delta]
[--write-datasync-action {datasync|filesync|unstable}]
[--revert-write-datasync-action]
[--write-datasync-reply {datasync|filesync}]
[--revert-write-datasync-reply]
[--write-filesync-action {datasync|filesync|unstable}]
[--revert-write-filesync-action]

```

```

[--write-filesystem-reply filesystem]
[--write-unstable-action {datasync|filesystem|unstable}]
[--revert-write-unstable-action]
[--write-unstable-reply {datasync|filesystem|unstable}]
[--revert-write-unstable-reply]
[--write-transfer-max-size <size>]
[--revert-write-transfer-max-size]
[--write-transfer-multiple <integer>]
[--revert-write-transfer-multiple]
[--write-transfer-size <size>]
[--revert-write-transfer-size]
[--zone <string>]
[--force]
[--verbose]

```

Options

--block-size <size>

Specifies the block size, in bytes.

--revert-block-size

Restores the setting to the system default.

--can-set-time {yes|no}

If set to *yes*, enables the export to set time. The default setting is *no*.

--revert-can-set-time

Restores the setting to the system default.

--case-insensitive {yes|no}

If set to *yes*, the server will report that it ignores case for file names. The default setting is *no*.

--revert-case-insensitive

Restores the setting to the system default.

--case-preserving {yes|no}

If set to *yes*, the server will report that it always preserves case for file names. The default setting is *no*.

--revert-case-preserving

Restores the setting to the system default.

--chown-restricted {yes|no}

If set to *yes*, the server will report that only the superuser can change file ownership. The default setting is *no*.

--revert-chown-restricted

Restores the setting to the system default.

--directory-transfer-size <size>

Specifies the preferred directory transfer size. Valid values are a number followed by a case-sensitive unit of measure: *b* for bytes; *K* for KB; *M* for MB; or *G* for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295*b*. The initial default value is 128*K*.

--revert-directory-transfer-size

Restores the setting to the system default.

--link-max <integer>

The reported maximum number of links to a file.

--revert-link-max

Restores the setting to the system default.

--max-file-size <size>

Specifies the maximum allowed file size on the server (in bytes). If a file is larger than the specified value, an error is returned.

--revert-max-file-size

Restores the setting to the system default.

--name-max-size <integer>

The reported maximum length of characters in a filename.

--revert-name-max-size

Restores the setting to the system default.

--no-truncate {yes|no}

If set to `yes`, too-long file names will result in an error rather than be truncated.

--revert-no-truncate

Restores the setting to the system default.

--return-32bit-file-ids {yes|no}

Applies to NFSv3 and later. If set to `yes`, limits the size of file identifiers returned from `readdir` to 32-bit values. The default value is `no`.



NOTE: This setting is provided for backward compatibility with older NFS clients, and should not be enabled unless necessary.

--revert-return-32bit-file-ids

Restores the setting to the system default.

--symlinks {yes|no}

If set to `yes`, advertises support for symlinks. The default setting is `no`.

--revert-symlinks

Restores the setting to the system default.

--new-zone <string>

Specifies a new access zone in which the export should apply. The default zone is `system`.

--all-dirs {yes|yesno}

If set to `yes`, this export will cover all directories. The default setting is `no`.

--revert-all-dirs

Restores the setting to the system default.

--encoding <string>

Specifies the character encoding of clients connecting through this NFS export.

Valid values and their corresponding character encodings are provided in the following table. These values are taken from the node's `/etc/encodings.xml` file, and are not case sensitive.

Value	Encoding
cp932	Windows-SJIS
cp949	Windows-949
cp1252	Windows-1252
euc-kr	EUC-KR
euc-jp	EUC-JP
euc-jp-ms	EUC-JP-MS
utf-8-mac	UTF-8-MAC
utf-8	UTF-8
iso-8859-1	ISO-8859-1 (Latin-1)
iso-8859-2	ISO-8859-2 (Latin-2)
iso-8859-3	ISO-8859-3 (Latin-3)
iso-8859-4	ISO-8859-4 (Latin-4)
iso-8859-5	ISO-8859-5 (Cyrillic)
iso-8859-6	ISO-8859-6 (Arabic)
iso-8859-7	ISO-8859-7 (Greek)

Value	Encoding
iso-8859-8	ISO-8859-8 (Hebrew)
iso-8859-9	ISO-8859-9 (Latin-5)
iso-8859-10	ISO-8859-10 (Latin-6)
iso-8859-13	ISO-8859-13 (Latin-7)
iso-8859-14	ISO-8859-14 (Latin-8)
iso-8859-15	ISO-8859-15 (Latin-9)
iso-8859-16	ISO-8859-16 (Latin-10)

--revert-encoding

Restores the setting to the system default.

--security-flavors {unix|krb5|krb5i|krb5p} ...

Specifies a security flavor to support. To support multiple security flavors, repeat this option for each additional entry. The following values are valid:

- sys** Sys or UNIX authentication.
- krb5** Kerberos V5 authentication.
- krb5i** Kerberos V5 authentication with integrity.
- krb5p** Kerberos V5 authentication with privacy.

--revert-security-flavors

Restores the setting to the system default.

--snapshot {<snapshot>|<snapshot-alias>}

Specifies the ID of a snapshot or snapshot alias to export. If you specify this option, directories will be exported in the state captured in either the specified snapshot or the snapshot referenced by the specified snapshot alias. If the snapshot does not capture the exported path, the export will be inaccessible to users.

If you specify a snapshot alias, and the alias is later modified to reference a new snapshot, the new snapshot will be automatically applied to the export.

Because snapshots are read-only, clients will not be able to modify data through the export unless you specify the ID of a snapshot alias that references the live version of the file system.

Specify *<snapshot>* or *<snapshot-alias>* as the ID or name of a snapshot or snapshot alias.

--revert-snapshot

Restores the setting to the system default.

--map-lookup-uid {yes|no}

If set to *yes*, incoming UNIX user identifiers (UIDs) will be looked up locally. The default setting is *no*.

--revert-map-lookup-uid

Restores the setting to the system default.

--map-retry {yes|no}

If set to *yes*, the system will retry failed user-mapping lookups. The default setting is *no*.

--revert-map-retry

Restores the setting to the system default.

--map-root-enabled {yes|no}

Enable/disable mapping incoming root users to a specific account.

--revert-map-root-enabled

Restores the setting to the system default.

--map-non-root-enabled {yes|no}

Enable/disable mapping incoming non-root users to a specific account.

--revert-map-non-root-enabled

Restores the setting to the system default.

--map-failure-enabled {yes|no}

Enable/disable mapping users to a specific account after failing an auth lookup.

--revert-map-failure-enabled

Restores the setting to the system default.

--map-all <identity>

Specifies the default identity that operations by any user will run as. If this option is not set to `root`, you can allow the root user of a specific client to run operations as the root user of the cluster by including the client in the `--root-clients` list.

--revert-map-all

Restores the setting to the system default.

--map-root <identity>

Map incoming root users to a specific user and/or group ID.

--revert-map-root

Restores the setting to the system default.

--map-non-root <identity>

Map non-root users to a specific user and/or group ID.

--revert-map-non-root

Restores the setting to the system default.

--map-failure <identity>

Map users to a specific user and/or group ID after a failed auth attempt.

--revert-map-failure

Restores the setting to the system default.

--map-full {yes|no}

Determines how user mapping is accomplished if a user is specified in an export option such as `--map-root` or `--map-all`. When enabled, a user mapping queries the OneFS user database and retrieves users from the applicable authentication subsystem, such as local authentication or Active Directory. When disabled, only local authentication is queried.

The default setting is `yes`.

--revert-map-full

Restores the `--map-full` setting to the system default, `yes`.

--commit-asynchronous {yes|no}

If set to `yes`, enables commit data operations to be performed asynchronously. The default setting is `no`.

--revert-commit-asynchronous

Restores the setting to the system default.

--read-only {yes|no}

Determines the default privileges for all clients accessing the export.

If set to `yes`, you can grant read/write privileges to a specific client by including the client in the `--read-write-clients` list.

If set to `no`, you can make a specific client read-only by including the client in the `--read-only-clients` list. The default setting is `no`.

--revert-read-only

Restores the setting to the system default.

--readdirplus {yes|no}

Applies to NFSv3 only. If set to `yes`, enables processing of readdir-plus requests. The default setting is `no`.

--revert-readdirplus

Restores the setting to the system default.

--read-transfer-max-size <size>

Specifies the maximum read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `K` for KB; `M` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 512K.

--revert-read-transfer-max-size
Restores the setting to the system default.

--read-transfer-multiple <integer>
Specifies the suggested multiple read size to report to NFSv3 and NFSv4 clients. Valid values are 0-4294967295. The initial default value is **512**.

--revert-read-transfer-multiple
Restores the setting to the system default.

--read-transfer-size <size>
Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: **b** for bytes; **k** for KB; **m** for MB; or **G** for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 128k.

--revert-read-transfer-size
Restores the setting to the system default.

--setattr-asynchronous {yes|no}
If set to *yes*, performs set-attributes operations asynchronously. The default setting is *no*.

--revert-setattr-asynchronous
Restores the setting to the system default.

--time-delta <integer>
Specifies server time granularity, in seconds.

--revert-time-delta
Restores the setting to the system default.

--write-datasync-action {datasync|filesync|unstable}
Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync write method. The following values are valid:

- *datasync*
- *filesync*
- *unstable*

The default value is *datasync*, which performs the request as specified.

--revert-write-datasync-action
Restores the setting to the system default.

--write-datasync-reply {datasync|filesync}
Applies to NFSv3 and NFSv4 only. Specifies an alternate datasync reply method. The following values are valid:

- *datasync*
- *filesync*

The default value is *datasync* (does not respond differently).

--revert-write-datasync-reply
Restores the setting to the system default.

--write-filesync-action {datasync|filesync|unstable}
Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync write method. The following values are valid:

- *datasync*
- *filesync*
- *unstable*

The default value is *filesync*, which performs the request as specified.

--revert-write-filesync-action
Restores the setting to the system default.

--write-filesync-reply {filesync}
Applies to NFSv3 and NFSv4 only. Specifies an alternate filesync reply method. The only valid value is *filesync* (does not respond differently).

--write-unstable-action {datasync|filesync|unstable}
Specifies an alternate unstable-write method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable`, which performs the request as specified.

--revert-write-unstable-action

Restores the setting to the system default.

--write-unstable-reply {`datasync`|`filesync`|`unstable`}

Specifies an alternate unstable-reply method. The following values are valid:

- `datasync`
- `filesync`
- `unstable`

The default value is `unstable` (does not respond differently).

--revert-write-unstable-reply

Restores the setting to the system default.

--write-transfer-max-size <size>

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `k` for KB; `m` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 512k.

--revert-write-transfer-max-size

Restores the setting to the system default.

--write-transfer-multiple <integer>

Specifies the suggested write transfer multiplier to report to NFSv3 and NFSv4 clients. Valid values are 0–4294967295. The initial default value is 512.

--revert-write-transfer-multiple

Restores the setting to the system default.

--write-transfer-size <size>

Specifies the preferred read transfer size to report to NFSv3 and NFSv4 clients. Valid values are a number followed by a case-sensitive unit of measure: `b` for bytes; `k` for KB; `m` for MB; or `G` for GB. If no unit is specified, bytes are used by default. The maximum value is 4294967295b. The initial default value is 512k.

--revert-write-transfer-size

Restores the setting to the system default.

--zone

Access zone in which the export was originally created.

--force

If set to `no` (default), a confirmation prompt displays when the command runs. If set to `yes`, the command runs without prompting for confirmation.

--verbose

Displays more detailed information.

isi nfs settings export view

Displays default NFS export settings.

Syntax

```
isi nfs settings export view
[--zone <string>]
```

Options

`--zone <string>`

Specifies the access zone in which the default settings apply.

Example

To view the currently-configured default export settings, run the following command:

```
isi nfs settings export view
```

The system displays output similar to the following example:

```
Read Write Clients: -
Unresolved Clients: -
  All Dirs: No
  Block Size: 8.0K
  Can Set Time: Yes
  Case Insensitive: No
  Case Preserving: Yes
  Chown Restricted: No
  Commit Asynchronous: No
Directory Transfer Size: 128.0K
  Encoding: DEFAULT
  Link Max: 32767
Map Lookup UID: No
  Map Retry: Yes
  Map Root
    Enabled: True
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Non Root
    Enabled: False
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Failure
    Enabled: False
    User: nobody
  Primary Group: -
  Secondary Groups: -
  Map Full: Yes
Max File Size: 8192.00000P
Name Max Size: 255
  No Truncate: No
  Read Only: No
  Readdirplus: Yes
Return 32Bit File Ids: No
Read Transfer Max Size: 1.00M
Read Transfer Multiple: 512
  Read Transfer Size: 128.0K
  Security Type: unix
Setattr Asynchronous: No
  Snapshot: -
  Symlinks: Yes
  Time Delta: 1.0 ns
Write Datasync Action: datasync
  Write Datasync Reply: datasync
Write Filesync Action: filesync
  Write Filesync Reply: filesync
Write Unstable Action: unstable
  Write Unstable Reply: unstable
Write Transfer Max Size: 1.00M
Write Transfer Multiple: 512
  Write Transfer Size: 512.0K
```


isi nfs settings global modify

Modifies the default NFS global settings.

Syntax

```
isi nfs settings global modify
  [--lock-protection <integer>]
  [--nfsv3-enabled {yes | no}]
  [--nfsv4-enabled {yes | no}]
  [--force]
```

Options

--lock-protection <integer>

Specifies the number of nodes failures that can happen before a lock might be lost.

--nfsv3-enabled {yes | no}

Specifies that NFSv3 is enabled.

--nfsv4-enabled {yes | no}

Specifies that NFSv4 is enabled.

{--force}

Causes the command to be executed without your confirmation.

isi nfs settings global view

Displays the global options for NFS settings.

Syntax

```
isi nfs settings global view
```

Options

There are no options for this command.

Example

The following is an example of the report generated by this command.

```
    NFSv3 Enabled: Yes
    NFSv4 Enabled: No
NFS Service Enabled: Yes
```

isi nfs settings zone modify

Modifies the default NFS zone settings for the NFSv4 ID mapper.

Syntax

```
isi nfs settings zone modify
  [--nfsv4-domain <string>]
```

```

[--revert-nfsv4-domain]
[--nfsv4-replace-domain {yes | no}]
[--revert-nfsv4-replace-domain]
[--nfsv4-no-domain {yes | no}]
[--revert-nfsv4-no-domain]
[--nfsv4-no-domain-uids {yes | no}]
[--revert-nfsv4-no-domain-uids]
[--nfsv4-no-names {yes | no}]
[--revert-nfsv4-no-names]
[--nfsv4-allow-numeric-ids {yes | no}]
[--revert-nfsv4-allow-numeric-ids]
[--zone <string>]
[--verbose]

```

Options

--nfsv4-domain <string>

Specifies the NFSv4 domain name.

--revert-nfsv4-domain

Returns the `--nfsv4-domain` setting to the system default (`localhost`).

--nfsv4-replace-domain {yes | no}

Replaces the owner/group domain with the NFSv4 domain name.

--revert-nfsv4-replace-domain

Returns setting to the system default. Default is `yes`.

--nfsv4-no-domain {yes | no}

Sends owners/groups without the NFSv4 domain name.

--revert-nfsv4-no-domain

Returns setting to the system default. Default is `no`.

--nfsv4-no-domain-uids {yes | no}

Sends UIDs/GIDs without the NFSv4 domain name.

--revert-nfsv4-no-domain-uids

Returns setting to the system default. Default is `yes`.

--nfsv4-no-names {yes | no}

Always sends owners/groups as UIDs/GIDs.

--revert-nfsv4-no-names

Returns setting to the system default. Default is `no`.

--nfsv4-allow-numeric-ids {yes | no}

Sends owners/groups as UIDs/GIDs when look-ups fail or if `--nfsv4-no-names` is enabled.

--revert-nfsv4-allow-numeric-ids

Returns setting to the system default. Default is `yes`.

--zone <string>

Specifies the access zone.

{--verbose | -v}

Displays more detailed information.

Example

The following command specifies that the NFS server would accept UIDs/GIDs in place of user names:

```
isi nfs settings zone modify --nfsv4-no-names yes
```

isi nfs settings zone view

Displays the default NFSv4-related access zone settings.

Syntax

```
isi nfs settings zone view  
[--zone <string>]
```

Options

--zone <string>

Specifies the access zone for which you want to view NFSv4-related settings.

Example

The following command specifies that you want to examine NFSv4-related settings for an access zone named Zone1:

```
isi nfs settings zone view --zone=Zone1
```

isi ntp servers create

Add a network time protocol (NTP) server to the cluster.

Syntax

```
isi ntp servers create <name>  
[--key <string>]  
[--verbose]
```

Options

<name> The host name of the NTP server you are adding to the cluster.

{--key | -k} <string> Value that maps the NTP server to a key in the key file.

{--verbose | -v} Display more detailed information.

isi ntp servers delete

Delete one or more network time protocol (NTP) servers from the cluster.

Syntax

```
isi ntp servers delete (<name> | --all)  
[--verbose]  
[--force]
```

Options

- <name>** The host name of an NTP server connected to the cluster.
- all** Delete all connected NTP servers.
- {--verbose | -v}** Display more detailed information.
- {--force | -f}** Do not prompt for confirmation of the delete.

isi ntp servers list

List network time protocol (NTP) servers.

Syntax

```
isi ntp servers list
  [--limit <integer>]
  [--sort (name | key)]
  [--descending]
  [--format (table | json | csv | list)]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- {--limit | -l}** The number of NTP servers to display.
<integer>
- sort (name | key)** Sort the list output by server name or key path.
- {--descending | -d}** Sort data in descending order.
- format (table | json | csv | list)** Displays NTP servers in table, JSON, CSV, or list format.
- {--no-header | -a}** Do not display headers in table or CSV formats.
- {--no-footer | -z}** Do not display table summary footer information.
- {--verbose | -v}** Display more detailed information.

OneFS displays a list of servers and key paths.

isi ntp servers modify

Modify network time protocol (NTP) server information.

Syntax

```
isi ntp servers modify <name>
  [--key (<string> | --clear-key)]
  [--verbose]
```

Options

- <name>** The host name of an NTP server connected to the cluster.
- {--key | -k} (<string> | --clear-key** **--clear-key** Clear the key value.
- {--verbose | -v}** Display more detailed information.

isi ntp servers view

View network time protocol (NTP) server properties.

Syntax

```
isi ntp servers view <name>
```

Options

- <name>** The host name of the NTP server.

OneFS displays information about the specified NTP server.

isi ntp settings modify

Modify network time protocol (NTP) server configuration settings.

Syntax

```
isi ntp settings modify
[--chimera <integer>]
[--excluded (<lnn> | --clear-excluded | --add-excluded <lnn> | --remove-excluded <lnn>)]
[--key-file (<path> | --clear-key-file]
[--verbose]
```

Options

- chimera** **<integer>** The number of chimer nodes that contact NTP servers.
- {--excluded | -x} (<lnn> | --clear-excluded | --add-excluded <lnn> | --remove-excluded <lnn>)** Specify logical node name (LNN) numbers for to exclude nodes from chimer duty. Specify this option again for each additional desired LNN.
- clear-excluded** **--clear-excluded** Clear the list of LNNs excluded from chimer duty.
- add-excluded <lnn>** **--add-excluded <lnn>** Add a server to the list of LNNs excluded from chimer duty. Specify this option again for each additional desired LNN.
- remove-excluded <lnn>** **--remove-excluded <lnn>** Remove a server from the list of LNNs excluded from chimer duty. Specify this option again for each additional desired LNN.
- {--key-file | -k} (<path> | --clear-key-file)** **<path>** Maintain NTP key file information. The path to the NTP key file within the OneFS file system.

--clear-key-file Clear the NTP key file path.

{--verbose | -v} Displays more detailed information.

isi ntp settings view

View cluster network time protocol (NTP) configuration settings.

Syntax

```
isi ntp settings view
```

Options

None

OneFS displays cluster NTP configuration settings.

isi performance datasets create

Create a new OneFS performance data set.

Syntax

```
isi performance datasets create <metrics> {groupname | local_address | path | protocol |  
remote_address | share_name | username | zone_name}  
[--filters]  
[--name]
```

Options

<metrics> The statistics metric(s) to include in the new data set. You can specify multiple metric options.

**{groupname |
local_address |
path | protocol |
remote_address
| share_name |
username |
zone_name}**

{--filter | -f} A statistics metric to filter the new dataset, from the available values in the *<metrics>* parameter.

{--name | -n} A custom name for the new data set.

isi performance datasets delete

Delete a performance data set.

Syntax

```
isi performance datasets delete <dataset>  
[--remove-filters]
```

```
[--unpin-workloads]
[--force]
```

Options

- <dataset>** The name or numeric ID of the data set you are deleting.
- {--remove-filters | -r}** Remove all filters from a performance data set before deleting the data set.
- {--unpin-workloads | -u}** Un-pin all workloads from a performance data set before deleting the data set.
- {--force | -f}** Do not prompt for confirmation of the data set deletion.

isi performance datasets list

List configured performance data sets.

Syntax

```
isi performance datasets list
  [--sort {id | name | statkey | creation_time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- sort {id | name | statkey | creation_time}** Sort data by the specified field.
- {--descending | -d}** Specifies to sort the data in descending order.
- format {table | json | csv | list}** Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- {--no-header | -a}** Displays table and CSV output without headers.
- {--no-footer | -z}** Displays table output without footers.
- {--verbose | -v}** Displays more detailed information.

isi performance datasets modify

Modify a configured performance data set.

Syntax

```
isi performance datasets modify <dataset>
[--name <string>]
```

Options

- <dataset>** The name or numeric ID of the data set you are modifying.
- {--name | -n}** A new custom name for the performance data set.

isi performance datasets view

View the properties of a configured OneFS performance data set.

Syntax

```
isi performance datasets view <dataset>
```

Options

- <dataset>** The name or numeric ID of the OneFS performance data set for which you are viewing properties.

OneFS displays the properties for the specified OneFS performance data set.

isi performance filters apply

Apply a new filter to a previously configured OneFS performance data set.

Syntax

```
isi performance filters apply <dataset> <metric-value>
[--name]
```

Options

- <dataset>** The name or numeric ID for the data set to which you are applying a filter.
- <metric-value>** A metric value for defining the new feature. Specify `--metric-value` for each additional value required. A metric value is a performance metric and a value for the metric, joined with a colon symbol (:). For example, `protocol:smb2`.
- {--name | -n}** A custom name for the new filter.

isi performance filters list

List the filters applied on a configured OneFS performance data set.

Syntax

```
isi performance filters list <dataset>
  [--sort {id | name | creation_time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- <dataset>** The name or numeric ID of the data set for which you are listing filters.
- sort {id | name | creation_time}** Sort data by the specified field.
- {--descending | -d}** Specifies to sort the data in descending order.
- format {table | json | csv | list}** Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- {--no-header | -a}** Displays table and CSV output without headers.
- {--no-footer | -z}** Displays table output without footers.
- {--verbose | -v}** Displays more detailed information.

isi performance filters modify

Modify a filter applied to a OneFS performance data set.

Syntax

```
isi performance filters modify <dataset> <filter>
  [--name <string>]
```

Options

- <dataset>** The name or numeric ID of the data set for which you are modifying a filter.
- <filter>** The name or numeric ID of the filter you are modifying.
- {--name | -n}** A new custom name for the filter.

isi performance filters remove

Remove a filter that is applied to a OneFS performance data set.

Syntax

```
isi performance filters remove <dataset> <filter>
[--force]
```

Options

- <dataset>** The name or numeric ID of the data set from which you are removing a filter.
- <filter>** The name or numeric ID of the filter to remove from the data set.
- {--force | -f}** Do not prompt for confirmation of the filter deletion.

isi performance filters view

View the properties of a filter that is applied to a OneFS performance data set.

Syntax

```
isi performance filters view <dataset> <filter>
```

Options

- <dataset>** The name or numeric ID of the data set to which you have applied the filter.
- <filter>** The name or numeric ID of the filter to view.

OneFS displays the properties for the specified filter applied to the OneFS performance data set.

isi performance metrics list

List statistics metrics that can be used to define OneFS performance data sets.

Syntax

```
isi performance metrics list
[--sort {id | datatype | system_only}]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

- sort {id | datatype | system_only}** Sort data by the specified field. Metrics classified as `system_only` are reserved for use by the system data set.

{--descending | -d} Specifies to sort the data in descending order.

--format {table | json | csv | list} Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a} Displays table and CSV output without headers.

{--no-footer | -z} Displays table output without footers.

{--verbose | -v} Displays more detailed information.

isi performance metrics view

View a statistics metric used to define data sets.

Syntax

```
isi performance metrics view <id>
```

Options

<id> The name or numeric ID of the statistics metric to view.

OneFS displays the properties for the specified statistics metric.

isi performance settings modify

Modify OneFS performance monitoring settings.

Syntax

```
isi performance settings modify <top-n-collection-count>
```

Options

<top-n-collection-count> The number of the highest resource consuming workloads, tracked and collected by OneFS for each configured performance data set. The default amount is 1,024.

isi performance settings view

View performance settings.

Syntax

```
isi performance settings view
```

OneFS displays performance settings.

isi performance workloads list

List the workloads pinned to a configured OneFS performance data set.

Syntax

```
isi performance workloads list <dataset>
  [--sort {id | name | creation_time}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

- <dataset>** The name or numeric ID of the data set for which you are listing pinned workloads.
- sort {id | name | creation_time}** Sort data by the specified field.
- {--descending | -d}** Specifies to sort the data in descending order.
- format {table | json | csv | list}** Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- {--no-header | -a}** Displays table and CSV output without headers.
- {--no-footer | -z}** Displays table output without footers.
- {--verbose | -v}** Displays more detailed information.

isi performance workloads modify

Modify a workload that is pinned to a configured OneFS performance data set.

Syntax

```
isi performance workloads modify <dataset> <workload>
  [--name <string>]
```

Options

- <dataset>** The name or numeric ID of the data set for which you are modifying a pinned workload.
- <workload>** The name or numeric ID of the workload you are modifying.
- {--name | -n}** A new custom name for the workload.

isi performance workloads pin

Pin a new workload to a configured OneFS performance data set.

Syntax

```
isi performance workloads pin <dataset> <metric-value>
[--name <string>]
```

Options

- <dataset>** The name or numeric ID of the data set for which you are pinning a workload.
- <metric-value>** A metric value for defining the new feature. Specify `--metric-value` for each additional value required. A metric value is a performance metric and a value for the metric, joined with a colon symbol (:). For example, `protocol:smb2`.
- {--name | -n}** A custom name for the workload.

isi performance workloads unpin

Remove a pinned workload from a configured OneFS performance data set.

Syntax

```
isi performance workloads unpin <dataset> <workload>
[--force]
```

Options

- <dataset>** The name or numeric ID of the data set from which you are un-pinning a workload.
- <workload>** The name or numeric ID of the workload to un-pin from the data set.
- {--force | -f}** Do not prompt for confirmation of the un-pinning of the workload from the data set.

isi performance workloads view

View the properties of a workload that is pinned to a configured OneFS performance data set.

Syntax

```
isi performance workloads view <dataset> <workload>
```

Options

- <dataset>** The name or numeric ID of the data set to view.
- <workload>** The name or numeric ID of the workload to view.

OneFS displays the properties for the specified workload pinned to the OneFS performance data set.

isi_phone_home

Modify the settings for the `isi_phone_home` feature, which gathers specific node- and cluster-related information to send to Isilon Technical Support on a weekly basis. This feature is enabled by default if you have EMC Secure Remote Services (ESRS) enabled.

Syntax

```
isi_phone_home
  [--enable]
  [--disable]
  [--logging-level {debug | info | warning | error | critical}]
  [--list-file <string>]
  [--script-file <string>]
  [--create-package]
  [--send-data]
  [--delete-data]
```

Options

(i) NOTE: We recommend that you run only the `--enable` or `--disable` options from the OneFS command-line interface. All others are run automatically when the tool is enabled, and are provided here for reference in case Isilon Technical Support asks you for a specific type of information.

`--enable | -e`

Enables `isi_phone_home`, providing that ESRS is configured and enabled. If you enabled ESRS when configuring the Isilon cluster, this feature is automatically enabled.

`--disable | -d`

Disables `isi_phone_home`.

`--logging-level | -o` {`debug` | `info` | `warning` | `error` | `critical`}

Emits logs specific to a log state, as well as all logs of higher priority. The default is `error`, which means all logs of condition `error` and `critical` are emitted. If you select a lower level log such as `warning`, all logs of level `warning`, `error`, and `critical` are emitted. We recommend that you do not change the default setting.

`--list-file | -l` <string>

Receives the name of a list file that contains `isi` commands to be run against the cluster. These list files are located in `/usr/local/isi_phone_home/list`.

`--script-file | -r` <string>

Receives the name of a Python script file to be run against the cluster. These script files are located in `/usr/local/isi_phone_home/script`.

`--create-package | -c`

Groups all the files in the `/ifs/data/Isilon_Support/phone_home/data` directory into a gzip file that is copied to `/ifs/data/Isilon_Support/phone_home/pkg`.

`--send-data | -s`

Scans `/ifs/data/Isilon_Support/phone_home/pkg` and uploads any unsent gzip files to Isilon Technical Support through ESRS.

`--delete-data | -t`

Deletes all gzip files older than 30 days from the `/ifs/data/Isilon_Support/phone_home/pkg` directory.

isi quota quotas create

Creates new file system quotas.

Syntax

```
isi quota quotas create <path> <type>
  [--user <name>]
  [--group <name>]
  [--gid <id>]
  [--uid <id>]
  [--sid <sid>]
  [--wellknown <name>]
  [--hard-threshold <size>]
  [--advisory-threshold <size>]
  [--soft-threshold <size>]
  [--soft-grace <duration>]
  [--container {yes | no}]
  [--ignore-limit-checks]
  [--include-snapshots {yes | no}]
  [--percent-advisory-threshold=<value>]
  [--percent-soft-threshold=<value>]
  [--thresholds-include-overhead {yes | no}]
  [--thresholds-on {fslogicalsize | physicalsize | applogicalsize}]
  [--zone <string>]
  [--enforced {yes | no}] [--zone <zone>]
  [--verbose]
```

Options

<path>

Specifies an absolute path within the `/ifs` file system.



CAUTION:

You should not create quotas of any type on the `/ifs` directory. A root-level quota may result in significant performance degradation.

<type> {directory | user | group | default-directory | default-user | default-group}

Specifies a quota type. The following values are valid:

directory	Creates a quota for all data in the directory, regardless of owner.
user	Creates a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Creates a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Sets a security identifier (SID). For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--hard-threshold <size>

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit. Size is a capacity value formatted as *<integer>[{b | K | M | G | T | P}]*.

--advisory-threshold <size>

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests. Size is a capacity value formatted as *<integer>[{b | K | M | G | T | P}]*.

--soft-threshold <size>

Specifies the soft threshold, which allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter. Size is a capacity value formatted as *<integer>[{b | K | M | G | T | P}]*.

--soft-grace <duration>

Specifies the soft threshold grace period, which is the amount of time to wait before disk write requests are denied.

Specify *<duration>* in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--container {yes | no}

Specifies whether to consider a share (container) or the entire Isilon filesystem when reporting total available space and amount of free space.

- **no**—Available space is relative to the entire Isilon cluster filesystem. This is the default setting.
- **yes**—Available space is relative to quotas set for each share. This feature is supported for protocols such as SMB, NFS, rquotas, and others. To determine the total available space in a share, OneFS considers both hard and soft quotas of all quota types (directory, groups, users). To report free space on a share, it finds the quota with the least amount of free space. This is not necessarily the smallest threshold set; it is the smallest currently available free space. For example, if a department has a 100T limit, and each user within the department has a 1T limit, it would be possible for the 100T share to run out of space before any one user runs out of their 1T.

The **yes** setting requires that either a hard or soft quota is set to define a share and that the **--enforced** setting is specified.

--ignore-limit-checks

Creates a child quota that has a higher quota limit than the parent quota.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

--percent-advisory-threshold=<value>

Specifies an advisory threshold as a percentage of the quota hard limit.

--percent-soft-threshold=<value>

Specifies a soft limit as a percentage of the quota hard limit.

--thresholds-include-overhead {yes | no}

Includes OneFS storage overhead in the quota threshold when set to *yes*.

--thresholds-on {fslogicalsize | physicalsize | applogicalsize}

Enforces the limits for this quota based on the following parameters.

fslogicalsize Base quota enforcement on file system logical size; storage usage which does not include metadata and data protection.

physicalsize Base quota enforcement on physical size; storage usage which includes metadata and data protection.

applogicalsize Base quota enforcement on application logical size; storage usage which includes capacity consumption on the cluster as well as data tiered to the cloud. This storage usage is usually equal to or less than the file system logical size.

--enforced {yes | no}

Enforces this quota when set to *yes*. Specifying any threshold automatically sets this value to *yes* on create.

--zone <zone>

Specifies an access zone.

{--verbose | -v}

Displays more detailed information.

isi quota quotas delete

Deletes a file system quota or multiple quotas.

Syntax

```
isi quota quotas delete <path> <type>
  [--uid <id>]
  [--user <name>]
  [--gid <id>]
  [--group <name>]
  [--sid <sid>]
  [--wellknown <name>]
  [--recurse-path-parents]
  [--recurse-path-children]
  [--include-snapshots {yes | no}]
  [--zone <zone>]
  [--verbose]
```

Options

<path>

Specifies an absolute path within the */ifs* file system.

<type> {directory | user | group | default-directory | default-user | default-group | --all}

Deletes quotas of the specified type. Argument must be specified with the *<path>* variable. The following values are valid:

directory	Specifies a quota for all data in the directory, regardless of owner.
user	Specifies a quota for one specific user. Requires specification of <code>--user</code> , <code>--uid</code> , or <code>--sid</code> .
group	Specifies a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , or <code>--sid</code> option.
default-directory	Specifies a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Specifies a master quota that creates a linked quota for every user who has data in the directory.
default-group	Specifies a master quota that creates a linked quota for every group that owns data in the directory.
--all	Deletes all quotas. Flag may not be specified with <i><type></i> or <i><path></i> .

--uid <id>

Deletes a quota by the specified numeric user identifier (UID).

--user <name>

Deletes a quota associated with the user identified by name.

--gid <id>

Deletes a quota by the specified numeric group identifier (GID).

--group <name>

Deletes a quota associated with the group identified by name.

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Deletes a quota associated with the wellknown persona.

--recurse-path-parents

Searches parent paths for quotas.

--recurse-path-children

Searches child paths for quotas.

--include-snapshots {yes | no}

Deletes quotas that include snapshot data usage.

--zone <zone>

Specifies an access zone.

{--verbose | -v}

Displays more detailed information.

isi quota quotas list

Displays a list of quotas.

Syntax

```
isi quota quotas list
  [--user <name> | --group <name> | --gid <id> | --uid <id> | --sid <sid> | --wellknown <name>]
  [--type (directory | user | group | default-directory | default-user | default-group)]
  [--path <path>]
  [--recurse-path-parents]
  [--recurse-path-children]
  [--include-snapshots <boolean>]
```

```

[--exceeded]
[--enforced <boolean>]
[--zone <string>]
[{--limit | -l} <integer>]
[--format (table | json | csv | list)]
[{--no-header | -a}]
[{--no-footer | -z}]
[{--verbose | -v}]
[{--help | -h}]

```

Options

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--type

Specifies a quota type. The following values are valid:

directory	Creates a quota for all data in the directory, regardless of owner.
user	Creates a quota for one specific user. Requires specification of the --user , --uid , --sid , or --wellknown option.
group	Creates a quota for one specific group. Requires specification of the --group , --gid , --sid , or --wellknown option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--path <path>

Specifies quotas on the specified path.

--recurse-path-parents

Specifies parent paths for quotas.

--recurse-path-children

Specifies child paths for quotas.

--include-snapshots <boolean>

Specifies quotas that include snapshot data usage.

--exceeded

Specifies only quotas that have an exceeded threshold.

--enforced <boolean>

Specifies quotas that have an enforced threshold.

--zone <string>

Specifies quotas in the specified zone.

{--limit | -l} <integer>

Specifies the number of quotas to display.

--format

Displays quotas in the specified format. The following values are valid:

- table
- json
- csv
- list

{--no-header | -a}

Suppresses headers in CSV or table formats.

{--no-footer | -z}

Suppresses table summary footer information.

{--verbose | -v}

Displays more detailed information.

{--help | -h}

Displays help for this command.

isi quota quotas modify

Modifies a file system quota.

Syntax

```
isi quota quotas modify <path> <type>
  [--user <name> | --group <name> | --gid <id> | --uid <id>
  | --sid <sid> | --wellknown <name>]
  [--hard-threshold <size>]
  [--clear-hard-threshold]
  [--advisory-threshold <size>]
  [--clear-advisory-threshold]
  [--soft-threshold <size>]
  [--clear-soft-threshold]
  [--soft-grace <duration>]
  [--container {yes | no}]
  [--ignore-limit-checks]
  [--include-snapshots {yes | no}]
  [--percent-advisory-threshold=<value>]
  [--thresholds-include-overhead {yes | no}]
  [--thresholds-on {fslogicalsize | physicalsize | applogicalsize}]
  [--enforced {yes | no}]
  [--linked {yes | no}]
  [--zone <string>]
  [--verbose]
```

Options

--path <path>

Specifies an absolute path within the /ifs file system.

--type

Specifies a quota type. The following values are valid:

- | | |
|------------------|--|
| directory | Creates a quota for all data in the directory, regardless of owner. |
| user | Creates a quota for one specific user. Requires specification of the --user, --uid, or --sid option. |

group	Creates a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , or <code>--sid</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota that you want to modify. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--hard-threshold <size>

Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

--clear-hard-threshold

Clears an absolute limit for disk usage.

--advisory-threshold <size>

Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

--clear-advisory-threshold

Clears the advisory threshold.

--soft-threshold <size>

Specifies the soft threshold, which allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter. Size is a capacity value formatted as `<integer>[b | K | M | G | T | P]`.

--clear-soft-threshold

Clears the soft threshold.

--soft-grace <duration>

Specifies the soft threshold grace period, which is the amount of time to wait before disk write requests are denied.

Specify `<duration>` in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

m	Specifies minutes
s	Specifies seconds

--container {yes | no}

Specifies whether to consider a share (container) or the entire Isilon filesystem when reporting total available space and amount of free space.

- **no**—Available space is relative to the entire Isilon cluster filesystem. This is the default setting.
- **yes**—Available space is relative to quotas set for each share. This feature is supported for protocols such as SMB, NFS, rquotas, and others. To determine the total available space in a share, OneFS considers both hard and soft quotas of all quota types (directory, groups, users). To report free space on a share, it finds the quota with the least amount of free space. This is not necessarily the smallest threshold set; it is the smallest currently available free space. For example, if a department has a 100T limit, and each user within the department has a 1T limit, it would be possible for the 100T share to run out of space before any one user runs out of their 1T.

The **yes** setting requires that either a hard or soft quota is set to define a share and that the **--enforced** setting is specified.

--ignore-limit-checks

Creates a child quota that has a higher quota limit than the parent quota.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

--percent-advisory-threshold=<value>

Specifies a soft limit or advisory threshold as a percentage of the quota hard limit.

--thresholds-include-overhead {yes | no}

Includes OneFS storage overhead in the quota threshold when set to **yes**.

--thresholds-on {fslogicalsize | physicalsize | applogicalsize}

Enforces the limits for this quota based on the following parameters.

fslogicalsize	Base quota enforcement on file system logical size; storage usage which does not include metadata and data protection.
physicalsize	Base quota enforcement on physical size; storage usage which includes metadata and data protection.
applogicalsize	Base quota enforcement on application logical size; storage usage which includes capacity consumption on the cluster as well as data tiered to the cloud. This storage usage is usually equal to or less than the file system logical size.

--enforced {yes | no}

Enforces this quota when set to **yes**. Specifying any threshold automatically sets this value to **yes** on create.

--linked {yes | no}

Unlinks a linked quota created automatically by a **default-directory**, **default-user**, or **default-group** quota. Unlinking allows the quota to be modified separately. To modify a linked quota, you must modify the original **default-directory**, **default-user**, or **default-group** quota it originated from, instead of the linked quota itself.

--zone <string>


The zone used by the quota. Use this parameter only to resolve personas used by the quota.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications clear

Clears rules for a quota and uses system notification settings.

 **NOTE:** Use the `isi quota quotas notifications disable` command to disable all notifications for a quota.

Syntax

```
isi quota quotas notifications clear <path> <type>
  [--user <name>]
  [--group <name>]
  [--gid <id>]
  [--uid <id>]
  [--sid <sid>]
  [--wellknown <name>]
  [--include-snapshots {yes | no}]
  [--force]
```

Options

<path>

Specifies an absolute path within the */ifs* file system.

<type>

Specifies a quota type. The following values are valid:

- directory** Creates a quota for all data in the directory, regardless of owner.
- user** Creates a quota for one specific user. Requires specification of the `--user`, `--uid`, `--sid`, or `--wellknown` option.
- group** Creates a quota for one specific group. Requires specification of the `--group`, `--gid`, `--sid`, or `--wellknown` option.
- default-directory** Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
- default-user** Creates a master quota that creates a linked quota for every user who has data in the directory.
- default-group** Creates a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

{--force | -f}

Skips the confirmation prompt.

isi quota quotas notifications create

Creates a notification rule for a quota.

Syntax

```
isi quota quotas notifications create
--path <path>
--type {directory | user | group | default-directory | default-user | default-group}
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--schedule <name>]
[--holdoff <duration>]
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address <address>]
[--verbose]
```

Options

--path <path>

Specifies an absolute path within the `/ifs` file system.

--type

Specifies a quota type. The following values are valid:

directory	Creates a quota for all data in the directory, regardless of owner.
user	Creates a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Creates a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold.

violated Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.

expired Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Sets a security identifier (SID). For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Specifies quotas that include snapshot data usage.

--schedule <name>

Specifies the date pattern at which recurring notifications are made.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--holdoff <duration>

Specifies the length of time to wait before generating a notification.

Specify <duration> in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
s	Specifies seconds

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

--action-email-address <address>

Specifies the email address of user to be notified. Specify `--action-email-address` for each additional email address of user to notify.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications delete

Deletes a quota notification rule.

Syntax

```
isi quota quotas notifications delete
--path <path>
--type {directory | user | group | default-directory | default-user | default-group}
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--verbose]
```

Options

--path <path>

Deletes quota notifications set on an absolute path within the `/ifs` file system.

--type

Deletes a quota notification by specified type. The following values are valid:

directory	Specifies a quota for all data in the directory, regardless of owner.
user	Specifies a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.

group	Specifies a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Specifies a master quota that creates a linked quota for every user who has data in the directory.
default-group	Specifies a master quota that creates a linked quota for every group that owns data in the directory.

--threshold

Deletes a quota notification by specified threshold. The following values are valid:

hard	Specifies an absolute limit for disk usage.
soft	Specifies the soft threshold.
advisory	Specifies the advisory threshold..

--condition

Deletes a quote notification by the specified condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Deletes a quota notification by the specified user name.

--group <name>

Deletes a quota notification by the specified group name.

--gid <id>

Deletes a quota notification by the specified numeric group identifier (GID).

--uid <id>

Deletes a quota notification by the specified numeric user identifier (UID).

--sid <sid>

Deletes a quota notification by the specified security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Deletes a quota notification by the specified well-known user, group, machine, or account name.

--include-snapshots {yes | no}


Deletes a quota notification by the specified settings for Included snapshots in the quota size.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications disable

Disables all quota notifications.

 **CAUTION:** When you disable all quota notifications, system notification behavior is disabled also. Use the `--clear` options to remove specific quota notification rules and fall back to the system default.

Syntax

```
isi quota quotas notifications disable
--path <path>
--type {directory | user | group | default-directory | default-user | default-group}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
```

Options

--path <path>

Specifies an absolute path within the `/ifs` file system.

--type

Disables quotas of the specified type. Argument must be specified with the `--path` option. The following values are valid:

directory	Specifies a quota for all data in the directory, regardless of owner.
user	Specifies a quota for one specific user. Requires specification of <code>-user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Specifies a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Specifies a master quota that creates a linked quota for every user who has data in the directory.
default-group	Specifies a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Disables a quota associated with the user identified by name.

--gid <id>

Disables a quota by the specified numeric group identifier (GID).

--uid <id>

Disables a quota by the specified numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting a quota. For example, `S-1-5-21-13`.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Disables quotas that include snapshot data usage.

isi quota quotas notifications list

Displays a list of quota notification rules.

Syntax

```
isi quota quotas notifications list
--path <path>
--type {directory | user | group | default-directory | default-user | default-group}
[--user <name> | --group <name> | --gid <id> | --uid <id>
 | --sid <sid> | --wellknown <name>]
[--include-snapshots {yes | no}]
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--path <path>

Specifies an absolute path within the `/ifs` file system.

--type

Specifies a quota type. The following values are valid:

directory	Creates a quota for all data in the directory, regardless of owner.
user	Creates a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Creates a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

{--limit | -l} <integer>

Specifies the number of quota notification rules to display.

--format

Displays quota notification rules in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Suppresses headers in CSV or table formats.

{--no-footer | -z}

Suppresses table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications modify

Modifies a notification rule for a quota.

Syntax

```
isi quota quotas notifications modify
  --path <path>
  --type {directory | user | group | default-directory | default-user | default-group}
  --threshold {hard | soft | advisory}
  --condition {exceeded | denied | violated | expired}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
  [--schedule <string>]
  [--holdoff <duration>]
  [--clear-holdoff]
  [--action-alert {yes | no}]
  [--action-email-owner {yes | no}]
  [--action-email-address <address>]
  [--clear-action-email-address]
  [--add-action-email-address <address>]
  [--remove-action-email-address <address>]
  [--email-template <path>]
  [--clear-email-template]
  [--verbose]
```

Options

--path <path>

Specifies an absolute path within the `/ifs` file system.

--type

Specifies a quota type. The following values are valid:

directory

Creates a quota for all data in the directory, regardless of owner.

user	Creates a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Creates a quota for one specific group. Requires specification of <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Sets a security identifier (SID). For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

--schedule <name>

Specifies the date pattern at which recurring notifications are made.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--holdoff <duration>

Specifies the length of time to wait before generating a notification.

Specify <duration> in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
s	Specifies seconds

--clear-holdoff

Clears the value for the --holdoff duration.

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires --action-email-address.

--action-email-address <address>

Specifies the email address of the user to be notified. Specify --action-email-address for each additional email address of user to notify.

--clear-action-email-address

Clears the value for the email address of the user to notify.

--add-action-email-address<address>

Adds the email address of the user to be notified. Specify `--add-action-email-address` for each additional email address of user to notify.

--remove-action-email-address<address>

Removes the email address of the user to notify. Specify `--remove-action-email-address` for each email address of user to notify.

--email-template <path>

Specifies the path in `/ifs` to the email template.

--clear-email-template

Clears the setting for the path to the email template.

{--verbose | -v}

Displays more detailed information.

isi quota quotas notifications view

Displays the properties of a quota notification rule.

Syntax

```
isi quota quotas notifications view
  --path <path>
  --type {directory | user | group | default-directory | default-user | default-group}
  --threshold {hard | soft | advisory}
  --condition {exceeded | denied | violated | expired}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
```

Options

--path <path>

Specifies an absolute path within the `/ifs` file system.

--type

Specifies a quota type. The following values are valid:

directory	Creates a quota for all data in the directory, regardless of owner.
user	Creates a quota for one specific user. Requires specification of the <code>--user</code> , <code>--uid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
group	Creates a quota for one specific group. Requires specification of the <code>--group</code> , <code>--gid</code> , <code>--sid</code> , or <code>--wellknown</code> option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Creates a master quota that creates a linked quota for every user who has data in the directory.
default-group	Creates a master quota that creates a linked quota for every group that owns data in the directory.

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--user <name>

Specifies a user name.

--group <name>

Specifies a group name.

--gid <id>

Specifies the numeric group identifier (GID).

--uid <id>

Specifies a numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Includes snapshots in the quota size.

isi quota quotas view

Displays detailed properties of a single file system quota.

Syntax

```
isi quota quotas view
  --path <path>
  --type {directory | user | group | default-directory | default-user | default-group}
  [--user <name> | --group <name> | --gid <id> | --uid <id>
   | --sid <sid> | --wellknown <name>]
  [--include-snapshots {yes | no}]
  [--zone <string>]
```

Options

--path <path>

Specifies an absolute path within the /ifs file system.

--type

Specifies quotas of the specified type. Argument must be specified with the **--path** option. The following values are valid:

directory	Specifies a quota for all data in the directory, regardless of owner.
user	Specifies a quota for one specific user. Requires specification of -user , --uid , --sid , or --wellknown option.
group	Specifies a quota for one specific group. Requires specification of the --group , --gid , --sid , or --wellknown option.
default-directory	Creates a master quota that creates a linked quota for every immediate subdirectory created in the directory.
default-user	Specifies a master quota that creates a linked quota for every user who has data in the directory.
default-group	Specifies a master quota that creates a linked quota for every group that owns data in the directory.

--user <name>

Specifies a quota associated with the user identified by name.

--group <name>

Specifies a quota associated with the group identified by name.

--gid <id>

Specifies a quota by the numeric group identifier (GID).

--uid <id>

Specifies a quota by the specified numeric user identifier (UID).

--sid <sid>

Specifies a security identifier (SID) for selecting the quota. For example, S-1-5-21-13.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

--include-snapshots {yes | no}

Specifies quotas that include snapshot data usage.

--zone <zone>

Specifies an access zone.

isi quota reports create

Generates a quota report.

Syntax

```
isi quota reports create  
[--verbose]
```

Options

{--verbose | -v}

Displays more detailed information.

isi quota reports delete

Deletes a specified report.

Syntax

```
isi quota reports delete
  --time <string>
  --generated {live | scheduled | manual}
  --type {summary | detail}
  [--verbose]
```

Options

--time <string>

Specifies the timestamp of the report.

Specify *<time-and-date>* in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify *<time>* as one of the following values.

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
h	Specifies hours
s	Specifies seconds

--generated

Specifies the method used to generate the report. The following values are valid:

```
live
scheduled
manual
```

--type

Specifies a report type. The following values are valid:

```
summary
detail
```

{--verbose | -v}

Displays more detailed information.

isi quota reports list

Displays a list of quota reports.

Syntax

```
isi quota reports list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--limit <integer>

Specifies the number of quotas to display.

--format

Displays quotas in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Suppresses headers in CSV or table formats.

{--no-footer | -z}

Suppresses table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi quota settings mappings create

Creates a SmartQuotas email mapping rule.

Syntax

```
isi quota settings mappings create <type> <domain> <mapping>
  [--verbose]
```

Options

<type> {ad | local | nis | ldap}

The authentication provider type for the source domain.

<domain>

The fully-qualified domain name for the source domain you are mapping.

<mapping>

The fully-qualified domain name for the destination domain you are mapping to.

{--verbose | -v}

Displays more detailed information.

isi quota settings mappings delete

Deletes SmartQuotas email mapping rules.

Syntax

```
isi quota settings mappings delete <type> <domain>
  [--all]
  [--verbose]
  [--force]
```

Options

<type> {**ad** | **local** | **nis** | **ldap**}

The authentication provider type for the source domain.

<domain>

The fully-qualified domain name for the source domain you are mapping.

--all

Deletes all mapping rules.

{**--verbose** | **-v**}

Displays more detailed information.

{**force** | **-f**}

Forces the deletion without displaying a confirmation prompt.

isi quota settings mappings list

Lists SmartQuotas email mapping rules.

Syntax

```
isi quota settings mappings list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{**--limit** | **-l**} **<integer>**

The number of quota mapping settings to display.

--format {**table** | **json** | **csv** | **list**}

Display quota mappings settings in table, JSON, CSV, or list format.

{**--no-header** | **-a**}

Do not display headers in table or CSV formats.

{**--no-footer** | **-z**}

Do not display table summary footer information.

{**--verbose** | **-v**}

Displays more detailed information.

isi quota settings mappings modify

Modifies an existing SmartQuotas email mapping rule.

Syntax

```
isi quota settings mappings modify <type> <domain> <mapping>
[--verbose]
```

Options

<type> {ad | local | nis | ldap}

The authentication provider type for the source domain.

<domain>

The fully-qualified domain name for the source domain you are mapping.

<mapping>

The fully-qualified domain name for the destination domain you are mapping to.

{--verbose | -v}

Displays more detailed information.

isi quota settings mappings view

View a SmartQuotas email mapping rule.

Syntax

```
isi quota settings mappings view <type> <domain>
```

Options

<type> {ad | local | nis | ldap}

The authentication provider type for the source domain.

<domain>

The fully-qualified domain name for the source domain you are mapping.

isi quota settings notifications clear

Clears all default quota notification rules.

When you clear all default notification rules, the system reverts to system notification behavior. Use the `--disable` option to disable notification settings for a specific quota notification rule.

Syntax

```
isi quota settings notifications clear
```

isi quota settings notifications create

Creates a default notification rule.

Syntax

```
isi quota settings notifications create
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
--schedule <string>
--holdoff <duration>
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address<address>]
[--email-template <path>]
[--verbose]
```

Options

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--schedule <string>

Specifies the date pattern at which recurring notifications are made.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month

- The `<integer>` {weekday | `<day>`} of every [{other | `<integer>`}] month
- Yearly on `<month>` `<integer>`
- Yearly on the {last | `<integer>`} [weekday | `<day>`] of `<month>`

Specify `<frequency>` in one of the following formats:

- at `<hh>[:<mm>]` [{AM | PM}]
- every [`<integer>`] {hours | minutes} [between `<hh>[:<mm>]` [{AM | PM}] and `<hh>[:<mm>]` [{AM | PM}]]
- every [`<integer>`] {hours | minutes} [from `<hh>[:<mm>]` [{AM | PM}] to `<hh>[:<mm>]` [{AM | PM}]]

You can optionally append "st", "th", or "rd" to `<integer>`. For example, you can specify "Every 1st month"

Specify `<day>` as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--holdoff <duration>

Specifies the length of time to wait before generating a notification.

Specify `<duration>` in the following format:

```
<integer> <units>
```

The following `<units>` are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
s	Specifies seconds

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

--action-email-address <address>

Specifies the email address of user to be notified. Specify `--action-email-address` for each additional email address of user to notify.

--email-template <path>

Specifies the path in `/ifs` to the email template.

{--verbose | -v}

Displays more detailed information.

isi quota settings notifications delete

Delete a default quota notification rule.

Syntax

```
isi quota settings notifications delete
  --threshold {hard | soft | advisory}
  --condition {exceeded | denied | violated | expired}
  [--verbose]
```

Options

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

{--verbose | -v}

Displays more detailed information.

isi quota settings notifications list

Displays a list of global quota notification rules.

Syntax

```
isi quota settings notifications list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Specifies the number of quota notification rules to display.

--format

Displays quotas in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Suppresses headers in CSV or table formats.

{--no-footer | -z}

Suppresses table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi quota settings notifications modify

Modifies a quota notification rule.

Syntax

```
isi quota settings notifications modify
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
[--schedule <string>]
[--holdoff <duration>]
[--clear-holdoff]
[--action-alert {yes | no}]
[--action-email-owner {yes | no}]
[--action-email-address <address>]
[--clear-action-email-address]
[--add-action-email-address <address>]
[--remove-action-email-address <address>]
[--email-template <path>]
[--clear-email-template]
[--verbose
```

Options

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

--schedule <string>

Specifies the date pattern at which recurring notifications are made.

--holdoff <duration>

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

Specifies the length of time to wait before generating a notification.

Specify <duration> in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks

D	Specifies days
H	Specifies hours
s	Specifies seconds

--clear-holdoff

Clears the value for the `--holdoff` duration.

--action-alert {yes | no}

Generates an alert when the notification condition is met.

--action-email-owner {yes | no}

Specifies that an email be sent to a user when the threshold is crossed. Requires `--action-email-address`.

--action-email-address <address>

Specifies the email address of user to be notified. Specify `--action-email-address` for each additional email address of user to notify.

--clear-action-email-address

Clears the value for the email address of the user to notify.

--add-action-email-address<address>

Adds the email address of the user to be notified. Specify `--add-action-email-address` for each additional email address of user to notify.

--remove-action-email-address<address>

Removes the email address of the user to notify. Specify `--remove-action-email-address` for each email address of user to notify.

--email-template <path>

Specifies the path in `/ifs` to the email template.

--clear-email-template

Clears the setting for the path to the email template.

{--verbose | -v}

Displays more detailed information.

isi quota settings notifications view

Displays properties of a system default notification rule.

Syntax

```
isi quota settings notifications view
--threshold {hard | soft | advisory}
--condition {exceeded | denied | violated | expired}
```

Options

--threshold

Specifies the threshold type. The following values are valid:

hard	Sets an absolute limit for disk usage. Attempts to write to disk are generally denied if the request violates the quota limit.
soft	Specifies the soft threshold. Allows writes to disk above the threshold until the soft grace period expires. Attempts to write to disk are denied thereafter.
advisory	Sets the advisory threshold. For notification purposes only. Does not enforce limitations on disk write requests.

--condition

Specifies the quota condition on which to send a notification. The following values are valid:

denied	Specifies a notification when a hard threshold or soft threshold outside of its soft grace period causes a disk write operation to be denied.
exceeded	Specifies a notification when disk usage exceeds the threshold. Applies to only soft thresholds within the soft-grace period.
violated	Specifies a notification when disk usage exceeds a quota threshold but none of the other conditions apply.
expired	Specifies a notification when disk usage exceeds the soft threshold and the soft-grace period has expired.

isi quota settings reports modify

Modifies cluster-wide quota report settings.

Syntax

```
isi quota settings reports modify
  [--schedule <schedule>]
  [--revert-schedule]
  [--scheduled-dir <dir>]
  [--revert-scheduled-dir]
  [--scheduled-retain <integer>]
  [--revert-scheduled-retain]
  [--live-dir <dir> | --revert-live-dir]
  [--live-retain <integer> | --revert-live-retain]
  [--verbose]
```

Options

--schedule <schedule>

Specifies the date pattern at which recurring notifications are made.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify *<frequency>* in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]

- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--revert-schedule

Sets the `--schedule` value to system default.

--scheduled-dir <dir>

Specifies the location where scheduled quota reports are stored.

--revert-scheduled-dir

Sets the `--scheduled-dir` value to system default.

--scheduled-retain <integer>

Specifies the maximum number of scheduled reports to keep.

--revert-scheduled-retain

Sets the `--scheduled-retain` value to system default.

--live-dir <dir>

Specifies the location where live quota reports are stored.

--revert-live-dir

Sets the `--live-dir` value to system default.

--live-retain <integer>

Specifies the maximum number of live quota reports to keep.

--revert-live-retain

Sets the `--live-retain` value to system default.

{--verbose | -v}

Displays more detailed information.

isi quota settings reports view

Displays cluster-wide quota report settings.

Syntax

```
isi quota settings reports view
```

Options

There are no options for this command.

isi readonly list

List read-only information for nodes.

Syntax

```
isi readonly list
  [--format (table | json | csv | list)]
```

```
[--no-header]
[--no-footer]
[--verbose]
```

Options

--format (table | json | csv | list) Displays ESRS data items in table, JSON, CSV, or list format.

{--no-header | -a} Do not display headers in table or CSV formats.

{--no-footer | -z} Do not display table summary footer information.

{--verbose | -v} Display more detailed information.

OneFS displays a list of current read-only modes and statuses for nodes, sorted by LNN.

isi readonly modify

Modify the current read-only status for a node.

Syntax

```
isi readonly modify
  [--allowed (yes | no)]
  [--enabled (yes | no)]
  [--node-lnn <string>]
  [--verbose]
```

Options

--allowed (yes | no) Read-only status allowed for the node. Default is yes.

--enabled (yes | no) Read-only status enabled for the node. Default is no.

--node-lnn <string> The logical node number (LNN) of the node to modify read-only status. If you do not specify an LNN, the local node is selected.

{--verbose | -v} Displays more detailed information.

isi readonly view

View the current read-only status for a specific node.

Syntax

```
isi readonly view
  [--node-lnn <integer>]
```


Options

--node-lnn *<integer>* The logical node number (LNN) for the specific. If you do specify an LNN, the local node is displayed.

OneFS displays read-only status information for a specific node.

isi remotesupport connectemc modify

Enables or disables support for EMC Secure Remote Services (ESRS) on an Isilon node.

Syntax

```
isi remotesupport connectemc modify
  [--enabled {yes|no}]
  [--primary-esrs-gateway <string>]
  [--secondary-esrs-gateway <string>]
  [--use-smtp-failover {yes|no}]
  [--email-customer-on-failure {yes|no}]
  [--gateway-access-pools <string>]...
  [--clear-gateway-access-pools]
  [--add-gateway-access-pools <string>]...
  [--remove-gateway-access-pools <string>]...
```

Options

--enabled {yes|no}

Specifies whether support ESRS is enabled on the Isilon cluster.

--primary-esrs-gateway <string>

Specifies the primary ESRS gateway server. The gateway server acts as the single point of entry and exit for IP-based remote support activities and monitoring notifications. You can specify the gateway as an IPv4 address or the gateway name.

--secondary-esrs-gateway <string>

Specifies an optional secondary ESRS gateway server that acts as a failover server. You can specify the gateway as an IPv4 address or the gateway name.

--use-smtp-failover {yes|no}

Specifies whether to send event notifications to a failover SMTP address upon ESRS transmission failure. The SMTP email address is specified through the `isi email settings modify` command.

--email-customer-on-failure {yes|no}

Specifies whether to send an alert to a customer email address upon failure of other notification methods. The customer email address is specified through the `isi_promptesrs -e` command.

--gateway-access-pools <string>...

Specifies the IP address pools on the Isilon cluster that will handle remote support connections through the ESRS gateway.

The IP address pools must belong to a subnet under `groupnet0`, which is the default system `groupnet`.



NOTE: We recommend that you designate pools with static IP addresses that are dedicated to remote connections through ESRS.

--clear-gateway-access-pools

Deletes the list of IP address pools that handle remote support connections.

--add-gateway-access-pools <string>...

Adds one or more IP address pools to the list of pools that will handle remote support connections through the ESRS gateway.

--remove-gateway-access-pools <string>...

Deletes one or more IP address pools from the list of pools that will handle remote support connections through the ESRS gateway.

Examples

The following command enables ESRS, specifies an IPv4 address as the primary gateway, directs OneFS to email the customer if all transmission methods fail, and removes an IP address pool from the list of pools that handle gateway connections:

```
isi remotesupport connectemc modify --enabled=yes \  
--primary-esrs-gateway=192.0.2.1 --email-customer-on-failure=yes \  
--remove-gateway-access-pools=subnet3.pool1
```

isi remotesupport connectemc view

Displays EMC Secure Remote Services (ESRS) settings on an Isilon node.

Syntax

```
isi remotesupport connectemc view
```

Options

This command has no options.

OneFS isi commands S through Z

This chapter contains documentation of the OneFS CLI commands `isi servicelight list` through `isi zone zones view`.

Topics:

- `isi servicelight list`
- `isi servicelight modify`
- `isi servicelight view`
- `isi services`
- `isi set`
- `isi smb log-level filters create`
- `isi smb log-level filters delete`
- `isi smb log-level filters list`
- `isi smb log-level filters view`
- `isi smb log-level modify`
- `isi smb log-level view`
- `isi smb openfiles close`
- `isi smb openfiles list`
- `isi smb sessions delete`
- `isi smb sessions delete-user`
- `isi smb sessions list`
- `isi smb settings global modify`
- `isi smb settings global view`
- `isi smb settings shares modify`
- `isi smb settings shares view`
- `isi smb settings zone modify`
- `isi smb settings zone view`
- `isi smb shares create`
- `isi smb shares delete`
- `isi smb shares list`
- `isi smb shares modify`
- `isi smb shares permission create`
- `isi smb shares permission delete`
- `isi smb shares permission list`
- `isi smb shares permission modify`
- `isi smb shares permission view`
- `isi smb shares view`
- `isi snapshot aliases create`
- `isi snapshot aliases delete`
- `isi snapshot aliases list`
- `isi snapshot aliases modify`
- `isi snapshot aliases view`
- `isi snapshot locks create`
- `isi snapshot locks delete`
- `isi snapshot locks list`
- `isi snapshot locks modify`
- `isi snapshot locks view`
- `isi snapshot schedules create`
- `isi snapshot schedules delete`
- `isi snapshot schedules list`
- `isi snapshot schedules modify`

- `isi snapshot schedules pending list`
- `isi snapshot schedules view`
- `isi snapshot settings modify`
- `isi snapshot settings view`
- `isi snapshot snapshots create`
- `isi snapshot snapshots delete`
- `isi snapshot snapshots list`
- `isi snapshot snapshots modify`
- `isi snapshot snapshots view`
- `isi snmp settings modify`
- `isi snmp settings view`
- `isi auth settings modify`
- `isi statistics client`
- `isi statistics data-reduction`
- `isi statistics data-reduction view`
- `isi statistics drive`
- `isi statistics heat`
- `isi statistics list keys`
- `isi statistics list operations`
- `isi statistics protocol`
- `isi statistics pstat`
- `isi statistics query current`
- `isi statistics query history`
- `isi status`
- `isi storagepool compatibilities active create`
- `isi storagepool compatibilities active delete`
- `isi storagepool compatibilities active list`
- `isi storagepool compatibilities active view`
- `isi storagepool compatibilities available list`
- `isi storagepool compatibilities class active create`
- `isi storagepool compatibilities class active delete`
- `isi storagepool compatibilities class active list`
- `isi storagepool compatibilities class active view`
- `isi storagepool compatibilities class available list`
- `isi storagepool compatibilities ssd active create`
- `isi storagepool compatibilities ssd active delete`
- `isi storagepool compatibilities ssd active list`
- `isi storagepool compatibilities ssd active view`
- `isi storagepool compatibilities ssd available list`
- `isi storagepool health`
- `isi storagepool list`
- `isi storagepool nodepools create`
- `isi storagepool nodepools delete`
- `isi storagepool nodepools list`
- `isi storagepool nodepools modify`
- `isi storagepool nodepools view`
- `isi storagepool settings modify`
- `isi storagepool settings modify`
- `isi storagepool settings view`
- `isi storagepool tiers create`
- `isi storagepool tiers delete`
- `isi storagepool tiers list`
- `isi storagepool tiers modify`
- `isi storagepool tiers view`
- `isi storagepool unprovisioned view`
- `isi swift accounts create`
- `isi swift accounts delete`

- `isi swift accounts list`
- `isi swift accounts modify`
- `isi swift accounts view`
- `isi sync certificates peer delete`
- `isi sync certificates peer import`
- `isi sync certificates peer list`
- `isi sync certificates peer modify`
- `isi sync certificates peer view`
- `isi sync certificates server delete`
- `isi sync certificates server import`
- `isi sync certificates server list`
- `isi sync certificates server modify`
- `isi sync certificates server view`
- `isi sync jobs cancel`
- `isi sync jobs list`
- `isi sync jobs pause`
- `isi sync jobs reports list`
- `isi sync jobs reports view`
- `isi sync jobs resume`
- `isi sync jobs start`
- `isi sync jobs view`
- `isi sync policies create`
- `isi sync policies delete`
- `isi sync policies disable`
- `isi sync policies enable`
- `isi sync policies list`
- `isi sync policies modify`
- `isi sync policies reset`
- `isi sync policies resolve`
- `isi sync policies view`
- `isi sync recovery allow-write`
- `isi sync recovery resync-prep`
- `isi sync reports list`
- `isi sync reports rotate`
- `isi sync reports subreports list`
- `isi sync reports subreports view`
- `isi sync reports view`
- `isi sync rules create`
- `isi sync rules delete`
- `isi sync rules list`
- `isi sync rules modify`
- `isi sync rules reports list`
- `isi sync rules reports view`
- `isi sync rules view`
- `isi sync service policies create`
- `isi sync service policies delete`
- `isi sync service policies disable`
- `isi sync service policies enable`
- `isi sync service policies list`
- `isi sync service policies modify`
- `isi sync service policies reset`
- `isi sync service policies resolve`
- `isi sync service policies view`
- `isi sync service recovery allow-write`
- `isi sync service recovery resync-prep`
- `isi sync service target break`
- `isi sync service target cancel`

- isi sync service target list
- isi sync service target view
- isi sync settings modify
- isi sync settings view
- isi sync target break
- isi sync target cancel
- isi sync target list
- isi sync target reports list
- isi sync target reports subreports list
- isi sync target reports subreports view
- isi sync target reports view
- isi sync target view
- isi tape delete
- isi tape list
- isi tape modify
- isi tape rename
- isi tape rescan
- isi tape view
- isi upgrade cluster add-nodes
- isi upgrade cluster add-remaining-nodes
- isi upgrade cluster archive
- isi upgrade cluster assess
- isi upgrade cluster commit
- isi upgrade cluster firmware
- isi upgrade cluster from-version
- isi upgrade cluster nodes firmware
- isi upgrade cluster nodes list
- isi upgrade cluster nodes view
- isi upgrade cluster retry-last-action
- isi upgrade cluster rollback
- isi upgrade cluster rolling-reboot
- isi upgrade cluster settings
- isi upgrade cluster start
- isi upgrade cluster to-version
- isi upgrade cluster view
- isi upgrade patches abort
- isi upgrade patches install
- isi upgrade patches list
- isi upgrade patches uninstall
- isi upgrade patches view
- isi version
- isi worm cdate set
- isi worm cdate view
- isi worm create
- isi worm domains create
- isi worm domains list
- isi worm domains modify
- isi worm domains view
- isi worm files delete
- isi worm files view
- isi zone restrictions create
- isi zone restrictions delete
- isi zone restrictions list
- isi zone zones create
- isi zone zones delete
- isi zone zones list
- isi zone zones modify

isi servicelight list

Displays a list of service LEDs in the cluster by node, along with the status of each service LED.

Syntax

```
isi servicelight list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{ --no-header | -a }

Displays table and CSV output without headers.

{ --no-footer | -z }

Displays table output without footers.

{ --verbose | -v }

Displays more detailed information.

isi servicelight modify

Turns a node's service LED on or off.

Syntax

```
isi servicelight modify
  [--enabled <boolean>]
  [--node-lnn <integer>]
  [--verbose]
```

Options

--enabled <boolean>

Enables or disables a node's service LED.

--node-lnn <integer>

Specifies the node on which you want to modify the service light status. If omitted, the local node will be used.

{ --verbose | -v }

Displays more detailed information.

isi servicelight view

Displays the status of a node's service LED.

Syntax

```
isi servicelight view
  [--node-lnn <integer>]
```

Options

--node-lnn <integer>

Specifies the node you want to view. If omitted, service LED status for the local node is displayed.

isi services

Displays a list of available services. The **-l** and **-a** options can be used separately or together.

Syntax

```
isi services
  [-l | -a]
  [<service> [{enable | disable}]]
```

Options

-l

Lists all available services and the current status of each. This is the default value for this command.

- a

Lists all services, including hidden services, and the current status of each.

<service> {enable | disable}

Enables or disables the specified service.

Examples

The following example shows the command to enable a specified hidden service.

```
isi services -a <hidden-service> enable
```

isi set

Works similar to `chmod`, providing a mechanism to adjust OneFS-specific file attributes, such as the requested protection, or to explicitly restripe files. Files can be specified by path or LIN.

Syntax

```
isi set
  [-f -F -L -n -v -r -R]
```



```

[-p <policy>]
[-w <width>]
[-c {on | off}]
[-g <restripe_goal>]
[-e <encoding>]
[-d <@r drives>]
[-a {<default> | <streaming> | <random> | <custom{1..5}> | <disabled>}]
[-l {<concurrency> | <streaming> | <random>}]
[--diskpool {<id> | <name>}]
[-A {on | off}]
[-P {on | off}]
[[--strategy | -s] {<avoid> | <metadata> | <metadata-write> |
<data>}]
[<file> {<path> | <lin>}]

```

Options

- f**
Suppresses warnings on failures to change a file.
- F**
Includes the `/ifs/.ifsvvar` directory content and any of its subdirectories. Without `-F`, the `/ifs/.ifsvvar` directory content and any of its subdirectories are skipped. This setting allows the specification of potentially dangerous, unsupported protection policies.
- L**
Specifies file arguments by LIN instead of path.
- n**
Displays the list of files that would be changed without taking any action.
- v**
Displays each file as it is reached.
- r**
Runs a restripe.
- R**
Sets protection recursively on files.
- p <policy>**
Specifies protection policies in the following forms:
 - +M** Where *M* is the number of node failures that can be tolerated without loss of data. *M* must be a number from, where numbers 1 through 4 are valid.
 - +D:M** Where *D* indicates the number of drive failures and *M* indicates number of node failures that can be tolerated without loss of data. *D* must be a number from 1 through 4 and *M* must be any value that divides into *D* evenly. For example, `+2:2` and `+4:2` are valid, but `+1:2` and `+3:2` are not.
 - Nx** Where *N* is the number of independent mirrored copies of the data that will be stored. *N* must be a number, with 1 through 8 being valid choices.
- w <width>**
Specifies the number of nodes across which a file is striped. Typically, $w = N + M$, but width can also mean the total of the number of nodes that are used.
You can set a maximum width policy of 32, but the actual protection is still subject to the limitations on *N* and *M*.
- c {on | off}**
Specifies whether write-coalescing is turned on.
- g <restripe goal>**

Specifies the restripe goal. The following values are valid:

repair
reprotect
rebalance
retune

-e <encoding>

Specifies the encoding of the filename. The following values are valid:

EUC-JP
EUC-JP-MS
EUC-KR
ISO-8859-1
ISO-8859-10
ISO-8859-13
ISO-8859-14
ISO-8859-15
ISO-8859-160
ISO-8859-2
ISO-8859-3
ISO-8859-4
ISO-8859-5
ISO-8859-6
ISO-8859-7
ISO-8859-8
ISO-8859-9
UTF-8
UTF-8-MAC
Windows-1252
Windows-949
Windows-SJIS

-d <@r drives>

Specifies the minimum number of drives that the file is spread across.

-a <value>

Specifies the file access pattern optimization setting. The following values are valid:

default
streaming
random
custom1
custom2
custom3
custom4
custom5
disabled

-l <value>

Specifies the file layout optimization setting. This is equivalent to setting both the `-a` and `-d` flags.

concurrency

streaming

random

--diskpool *<id | name>*

Sets the preferred diskpool for a file.

-A {on | off}

Specifies whether file access and protections settings should be managed manually.

-P {on | off}

Specifies whether the file inherits values from the applicable file pool policy.

{--strategy | -s} *<value>*

Sets the SSD strategy for a file. The following values are valid:

If the value is `metadata-write`, all copies of the file's metadata are laid out on SSD storage if possible, and user data still avoids SSDs. If the value is `data`, Both the file's meta- data and user data (one copy if using mirrored protection, all blocks if FEC) are laid out on SSD storage if possible.

avoid	Writes all associated file data and metadata to HDDs only. The data and metadata of the file are stored so that SSD storage is avoided, unless doing so would result in an out-of-space condition.
metadata	Writes both file data and metadata to HDDs. One mirror of the metadata for the file is on SSD storage if possible, but the strategy for data is to avoid SSD storage.
metadata-write	Writes file data to HDDs and metadata to SSDs, when available. All copies of metadata for the file are on SSD storage if possible, and the strategy for data is to avoid SSD storage.
data	Uses SSD node pools for both data and metadata. Both the metadata for the file and user data, one copy if using mirrored protection and all blocks if FEC, are on SSD storage if possible.
<i><file></i> { <i><path></i> <i><lin></i> }	Specifies a file by path or LIN.

isi smb log-level filters create

Creates a new SMB log filter.

Syntax

```
isi smb log-level filters create <level>  
  [--ops <string>]  
  [--ip-addr <string>]  
  [--verbose]
```

Options

<level>

The logging level for the new filter. Valid logging levels are:

- always
- error
- warning
- info

- verbose
- debug
- trace

{--ops | -o} <string>

List all SMB operations to filter against. Repeat for each operation.

{--ip-addr | -i} <string>

List IPv4 and IPv6 addresses to filter against. Repeat for each IP address.

isi smb log-level filters delete

Deletes SMB log filters.

Syntax

```
isi smb log-level filters delete <id> <level>
  [--all]
  [--force]
  [--verbose]
```

Options

<id>

Deletes a specific SMB log filter, by ID.

<level>

Deletes all SMB log filters at a specified level. The following levels are valid:

- always
- error
- warning
- info
- verbose
- debug
- trace

--all

Deletes all SMB log-level filters.

{--force | -f}

Skips the delete confirmation prompt.

{verbose | -v}

Displays more detailed information.

isi smb log-level filters list

Lists SMB log filters.

Syntax

```
isi smb log-level filters list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`{--limit | -l} <integer>`

Displays the specified number of SMB log-level filters.

`--format {table | json | csv | list}`

Displays SMB log-level filters in table, JSON, comma-separated, or list format.

`{--no-header | -a}`

Does not display headers in comma-separated or table format.

`{--no-footer | -z}`

Does not display table summary footer information.

`{--verbose | -v}`

Displays more detailed information.

isi smb log-level filters view

View an individual SMB log-level filter.

Syntax

```
isi smb log-level filters view <id>
  [--level <string>]
```

Options

`<id>`

The ID of the SMB log-level filter to view.

`{--level | -l} <string>`

Specifies a log-level to view. The following levels are valid:

- always
- error
- warning
- info
- verbose
- debug
- trace

isi smb log-level modify

Sets the log level for the SMB service.

Syntax

```
isi smb log-level modify <level>
  [--verbose]
```

Options

`<level>`

Specifies a log level to set for the SMB service. The following levels are valid:

- always
- error
- warning
- info
- verbose
- debug
- trace

{--verbose | -v}

Displays more detailed information.

isi smb log-level view

Shows the current log level for the SMB service.

Syntax

```
isi smb log-level view
```

Options

There are no options for this command.

isi smb openfiles close

Closes an open file.

 **NOTE:**

To view a list of open files, run the `isi smb openfiles list` command.

Syntax

```
isi smb openfiles close <id>  
[--force]
```

Options

<id>

Specifies the ID of the open file to close.

{--force | -f}

Suppresses command-line prompts and messages.

Examples

The following command closes a file with an ID of 32:

```
isi smb openfiles close 32
```

isi smb openfiles list

Displays a list of files that are open in SMB shares.

Syntax

```
isi smb openfiles list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of smb openfiles.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi smb sessions delete

Deletes SMB sessions, filtered first by computer and then optionally by user.

 **NOTE: Any open files are automatically closed before an SMB session is deleted.**

Syntax

```
isi smb sessions delete <computer-name>
  [{--user <name> | --uid <id> | --sid <sid>}]
  [--force]
  [--verbose]
```

Options

<computer-name>

Required. Specifies the computer name. If a **--user**, **--uid**, or **--sid** option is not specified, the system deletes all SMB sessions associated with this computer.

--user <string>

Specifies the name of the user. Deletes only those SMB sessions to the computer that are associated with the specified user.

--uid <id>

Specifies a numeric user identifier. Deletes only those SMB sessions to the computer that are associated with the specified user identifier.

--sid <sid>

Specifies a security identifier. Deletes only those SMB sessions to the computer that are associated with the security identifier.

{--force | -f}

Specifies that the command execute without prompting for confirmation.

Examples

The following command deletes all SMB sessions associated with a computer named computer1:

```
isi smb sessions delete computer1
```

The following command deletes all SMB sessions associated with a computer named computer1 and a user named user1:

```
isi smb sessions delete computer1 --user=user1
```

isi smb sessions delete-user

Deletes SMB sessions, filtered first by user then optionally by computer.

 **NOTE:**

Any open files are automatically closed before an SMB session is deleted.

Syntax

```
isi smb sessions delete-user {<user> | --uid <id> | --sid <sid> }  
  [--computer-name <string>]  
  [--force]  
  [--verbose]
```

Options

<user>

Required. Specifies the user name. If the `--computer-name` option is omitted, the system deletes all SMB sessions associated with this user.

{--computer-name | -C} <string>

Deletes only the user's SMB sessions that are associated with the specified computer.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

Examples

The following command deletes all SMB sessions associated with a user called user1:

```
isi smb sessions delete-user user1
```

The following command deletes all SMB sessions associated with a user called user1 and a computer called computer1:

```
isi smb sessions delete-user user1 \  
--computer-name=computer1
```


isi smb sessions list

Displays a list of open SMB sessions.

Syntax

```
isi smb sessions list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Specifies the maximum number of SMB sessions to list.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi smb settings global modify

Modifies global SMB settings.

Syntax

```
isi smb settings global modify
  [--access-based-share-enum {yes | no}]
  [--revert-access-based-share-enum]
  [--dot-snap-accessible-child {yes | no}]
  [--revert-dot-snap-accessible-child]
  [--dot-snap-accessible-root]
  [--revert-dot-snap-accessible-root]
  [--dot-snap-visible-child {yes | no}]
  [--revert-dot-snap-visible-child]
  [--dot-snap-visible-root {yes | no}]
  [--revert-dot-snap-visible-root]
  [--enable-security-signatures {yes | no}]
  [--revert-enable-security-signatures]
  [--guest-user <string>]
  [--revert-guest-user]
  [--ignore-eas {yes | no}]
  [--revert-ignore-eas]
  [--onefs-cpu-multiplier <integer>]
  [--revert-onefs-cpu-multiplier]
  [--onefs-num-workers <integer>]
  [--revert-onefs-num-workers]
  [--reject-unencrypted-access {yes | no}]
  [--revert-reject-unencrypted-access]
  [--require-security-signatures {yes | no}]
```

```

[--revert-require-security-signatures]
[--server-side-copy {yes | no}]
[--revert-server-side-copy]
[--server-string <string>]
[--revert-server-string]
[--support-multichannel {yes | no}]
[--revert-support-multichannel]
[--support-netbios {yes | no}]
[--revert-support-netbios]
[--support-smb2 {yes | no}]
[--revert-support-smb2]
[--support-smb3-encryption {yes | no}]
[--revert-support-smb3-encryption]
[--verbose]

```

Options

--access-based-share-enum {yes | no}

Enumerates only the files and folders that the requesting user has access to.

--revert-access-based-share-enum

Sets the value to the system default for `--access-based-share-enum`.

--dot-snap-accessible-child {yes | no}

Specifies whether to make the `/ifs/.snapshot` directory visible in subdirectories of the share root. The default setting is `no`.

--revert-dot-snap-accessible-child

Sets the value to the system default for `--dot-snap-accessible-child`.

--dot-snap-accessible-root {yes | no}

Specifies whether to make the `/ifs/.snapshot` directory accessible at the share root. The default setting is `yes`.

--revert-dot-snap-accessible-root

Sets the value to the system default for `--dot-snap-accessible-root`.

--dot-snap-visible-child {yes | no}

Specifies whether to make the `/ifs/.snapshot` directory visible in subdirectories of the share root. The default setting is `no`.

--revert-dot-snap-visible-child

Sets the value to the system default for `--dot-snap-visible-child`.

--dot-snap-visible-root {yes | no}

Specifies whether to make the `/ifs/.snapshot` directory visible at the root of the share. The default setting is `no`.

--revert-dot-snap-visible-root

Sets the value to the system default for `--dot-snap-visible-root`.

--enable-security-signatures {yes | no}

Indicates whether the server supports signed SMB packets.

--revert-enable-security-signatures

Sets the value to the system default for `--enable-security-signatures`.

--guest-user <integer>

Specifies the fully qualified user to use for guest access.

--revert-guest-user

Sets the value to the system default for `--guest-user`.

--ignore-eas {yes | no}

Specifies whether to ignore EAs on files.

--revert-ignore-eas

Sets the value to the system default for `--ignore-eas`.

--onefs-cpu-multiplier <integer>
 Specifies the number of OneFS worker threads to configure based on the number of CPUs. Valid numbers are 1-4.

--revert-onefs-cpu-multiplier
 Sets the value to the system default for **--onefs-cpu-multiplier**.

--onefs-num-workers <integer>
 Specifies the number of OneFS worker threads that are allowed to be configured. Valid numbers are 0-1024. If set to 0, the number of SRV workers will equal the value specified by **--onefs-cpu-multiplier** times the number of CPUs.

--revert-onefs-num-workers
 Sets the value to the system default for **--onefs-num-workers**.

--reject-unencrypted-access {yes | no}
 Rejects unencrypted client sessions.

--revert-reject-unencrypted-access
 Sets the value to the system default for **--reject-unencrypted-access**.

--require-security-signatures {yes | no}
 Specifies whether packet signing is required. If set to *yes*, signing is always required. If set to *no*, signing is not required but clients requesting signing will be allowed to connect if the **--enable-security-signatures** option is set to *yes*.

--revert-require-security-signatures
 Sets the value to the system default for **--require-security-signatures**.

--server-side-copy {yes | no}
 Enables or disables SMB server-side copy functionality. The default is *yes*.

--revert-server-side-copy
 Sets the value to the system default for **--server-side-copy**.

--server-string <string>
 Provides a description of the server.

--revert-server-string
 Sets the value to the system default for **--revert-server-string**.

--support-multichannel {yes | no}
 Specifies whether Multichannel for SMB 3.0 is enabled on the cluster. SMB Multichannel is enabled by default.

--revert-support-multichannel
 Set the value of **--support-multichannel** back to the default system value.

--support-netbios {yes | no}
 Specifies whether to support the NetBIOS protocol.

--revert-support-netbios
 Sets the value to the system default for **--support-netbios**.

--support-smb2 {yes | no}
 Specifies whether to support the SMB 2.0 protocol. The default setting is *yes*.

--revert-support-smb2
 Sets the value to the system default for **--support-smb2**.

--support-smb3-encryption {yes | no}
 Enables SMBv3 encryption on the cluster. The default setting is *no*.

--revert-support-smb3-encryption
 Sets the value to the system default for **--support-smb3-encryption**.

isi smb settings global view

Displays the default SMB configuration settings.

Syntax

```
isi smb settings global view
```

Options

There are no options for this command.

isi smb settings shares modify

Modifies default settings for SMB shares.

Syntax

```
isi smb settings shares modify
  [--access-based-enumeration {yes | no}]
  [--revert-access-based-enumeration]
  [--access-based-enumeration-root-only {yes | no}]
  [--revert-access-based-enumeration-root-only]
  [--allow-delete-readonly {yes | no}]
  [--revert-allow-delete-readonly]
  [--allow-execute-always {yes | no}]
  [--revert-allow-execute-always]
  [--ca-timeout <integer>]
  [--revert-ca-timeout]
  [--strict-ca-lockout {yes | no}]
  [--revert-strict-ca-lockout]
  [--ca-write-integrity {none | write-read-coherent | full}]
  [--revert-ca-write-integrity]
  [--change-notify {all | norecurse | none}]
  [--revert-change-notify]
  [--create-permissions {"default acl" | "inherit mode bits" | "use create mask and mode"}]
  [--revert-create-permissions]
  [--directory-create-mask <integer>]
  [--revert-directory-create-mask]
  [--directory-create-mode <integer>]
  [--revert-directory-create-mode]
  [--file-create-mask <integer>]
  [--revert-file-create-mask]
  [--file-create-mode <integer>]
  [--revert-file-create-mode]
  [--file-filtering-enabled {yes | no}]
  [--revert-file-filtering-enabled]
  [--file-filter-extensions <string>]
  [--clear-file-filter-extensions]
  [--add-file-filter-extensions <string>]
  [--remove-file-filter-extensions <string>]
  [--revert-file-filter-extensions]
  [--file-filter-type {deny | allow}]
  [--revert-file-filter-type]
  [--hide-dot-files {yes | no}]
  [--revert-hide-dot-files]
  [--host-acl <host-acl>]
  [--revert-host-acl]
  [--clear-host-acl]
  [--add-host-acl <string>]
  [--remove-host-acl <string>]
  [--impersonate-guest {always | "bad user" | never}]
  [--revert-impersonate-guest]
```

```

[--impersonate-user <string>]
[--revert-impersonate-user]
[--mangle-byte-start <integer>]
[--revert-mangle-byte-start]
[--mangle-map <mangle-map>]
[--revert-mangle-map]
[--clear-mangle-map]
[--add-mangle-map <string>]
[--remove-mangle-map <string>]
[--ntfs-acl-support {yes | no}]
[--revert-ntfs-acl-support]
[--oplocks {yes | no}]\
[--revert-oplocks]
[--support-smb3-encryption {yes | no}]
[--revert-support-smb3-encryption]
[--strict-flush {yes | no}]
[--revert-strict-flush]
[--strict-locking {yes | no}]
[--revert-strict-locking]
[--zone <string>]

```

Options

--access-based-enumeration {yes | no}

Specifies whether access-based enumeration is enabled.

--revert-access-based-enumeration

Sets the value to the system default for `--access-based-enumeration`.

--access-based-enumeration-root-only {yes | no}

Specifies whether access-based enumeration is only enabled on the root directory of the share.

--revert-access-based-enumeration-root-only

Sets the value to the system default for `--access-based-enumeration-root-only`.

--allow-delete-readonly {yes | no}

Specifies whether read-only files can be deleted.

--revert-allow-delete-readonly

Sets the value to the system default for `--allow-delete-readonly`.

--allow-execute-always {yes | no}

Specifies whether a user with read access to a file can also execute the file.

--revert-allow-execute-always

Sets the value to the system default for `--allow-execute-always`.

--ca-timeout <integer>

The amount of time, in seconds, a persistent handle is retained after a client is disconnected or a server fails. The default is 120 seconds.

--revert-ca-timeout

Sets the value to the system default for `--ca-timeout`.

--strict-ca-lockout {yes | no}

If set to `yes`, prevents another client from opening a file if a client has an open but disconnected persistent handle for that file. If set to `no`, OneFS issues persistent handles, but discards them if any client other than the original opener attempts to open the file. This option is only relevant if `--continuously-available` was set to `yes` when the share was created. The default is `yes`.

--revert-strict-ca-lockout

Sets the value to the system default for `--strict-ca-lockout`.

--ca-write-integrity {none | write-read-coherent | full}

Specifies the level of write integrity on continuously available shares:

none	Continuously available writes are not handled differently than other writes to the cluster. If you specify <code>none</code> and a node fails, you may experience data loss without notification. Therefore, we do not recommend this option.
write-read-coherent	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data. This is the default setting.
full	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data, and prevents OneFS from granting SMB clients write-caching and handle-caching leases.

--revert-ca-write-integrity

Sets the value to the system default for `--ca-write-integrity`.

--change-notify {norecurse | all | none}

Defines the change notify setting. The acceptable values are `norecurse`, `all`, and `none`.

--revert-change-notify

Sets the value to the system default for `--change-notify`.

--create-permissions {"default acl" | "inherit mode bits" | "use create mask and mode"}

Sets the default permissions to apply when a file or directory is created.

--revert-create-permissions

Sets the value to the system default for `--create-permissions`.

--directory-create-mask <integer>

Defines which mask bits are applied when a directory is created.

--revert-directory-create-mask

Sets the value to the system default for `--directory-create-mask`.

--directory-create-mode <integer>

Defines which mode bits are applied when a directory is created.

--revert-directory-create-mode

Sets the value to the system default for `--directory-create-mode`.

--file-create-mask <integer>

Defines which mask bits are applied when a file is created.

--revert-file-create-mask

Sets the value to the system default for `--file-create-mask`.

--file-create-mode <integer>

Defines which mode bits are applied when a file is created.

--revert-file-create-mode

Sets the value to the system default for `--file-create-mode`.

--file-filtering-enabled {yes | no}

If set to `yes`, enables file filtering at the share level. The default setting is `no`.

--revert-file-filtering-enabled

Sets the value for the system default of `--file-filtering-enabled`.

--file-filter-type {deny | allow}

If set to `allow`, allows the specified file types to be written to the share. The default setting is `deny`.

--revert-file-filter-type

Sets the value for the system default of `--file-filter-type`.

--file-filter-extensions <string>

Specifies the list of file types to deny or allow writes to the share, depending on the setting of `--file-filter-type`. File types may be specified in a list of comma separated values.

--clear-file-filter-extensions

Clears the list of file filtering extensions for the share.

--add-file-filter-extensions <string>

Adds entries to the list of file filter extensions. Repeat for each file extension to add.

--remove-file-filter-extensions <string>

Removes entries to the list of file filter extensions. Repeat for each file extension to remove.

--revert-file-filter-extensions

Sets the value for the system default of `--file-filter-extensions`.

--hide-dot-files {yes | no}

Specifies whether to hide files that begin with a period—for example, UNIX configuration files.

--revert-hide-dot-files

Sets the value to the system default for `--hide-dot-files`.

--host-acl <string>

Specifies which hosts are allowed access. Specify `--host-acl` for each additional host ACL clause. This will replace any existing ACL.

--revert-host-acl

Sets the value to the system default for `--host-acl`.

--clear-host-acl <string>

Clears the value for an ACL expressing which hosts are allowed access.

--add-host-acl <string>

Adds an ACE to the already-existing host ACL. Specify `--add-host-acl` for each additional host ACL clause to be added.

--remove-host-acl <string>

Removes an ACE from the already-existing host ACL. Specify `--remove-host-acl` for each additional host ACL clause to be removed.

--impersonate-guest {always | "bad user" | never}

Allows guest access to the share. The acceptable values are `always`, `"bad user"`, and `never`.

--revert-impersonate-guest

Sets the value to the system default for `--impersonate-guest`.

--impersonate-user <string>

Allows all file access to be performed as a specific user. This must be a fully qualified user name.

--revert-impersonate-user

Sets the value to the system default for `--impersonate-user`.

--mangle-byte-start <string>

Specifies the `wchar_t` starting point for automatic invalid byte mangling.

--revert-mangle-byte-start

Sets the value to the system default for `--mangle-byte-start`.

--mangle-map <string>

Maps characters that are valid in OneFS but are not valid in SMB names.

--revert-mangle-map

Sets the value to the system default for `--mangle-map`.

--clear-mangle-map <string>

Clears the values for character mangle map.

--add-mangle-map <string>

Adds a character mangle map. Specify `--add-mangle-map` for each additional Add character mangle map.

--remove-mangle-map <string>

Removes a character mangle map. Specify `--remove-mangle-map` for each additional Remove character mangle map.

--ntfs-acl-support {yes | no}

Specifies whether ACLs can be stored and edited from SMB clients.

--revert-ntfs-acl-support

Sets the value to the system default for `--ntfs-acl-support`.

--oplocks {yes | no}

Specifies whether to allow oplock requests.

--revert-oplocks

Sets the value to the system default for `--oplocks`.

--support-smb3-encryption {yes | no}

Enables SMBv3 encryption on the share. The default setting is no.

--revert-support-smb3-encryption

Sets the value to the system default for `--support-smb3-encryption`.

--strict-flush {yes | no}

Specifies whether to always honor flush requests.

--revert-strict-flush

Sets the value to the system default for `--strict-flush`.

--strict-locking {yes | no}

Specifies whether the server will check for and enforce file locks.

--revert-strict-locking

Sets the value to the system default for `--strict-locking`.

--zone <string>

Specifies the name of the access zone.

isi smb settings shares view

Displays default settings for all SMB shares or for SMB shares in a specified access zone.

Syntax

```
isi smb settings shares view
  [--zone <string>]
```

Options

--zone <string>

Specifies the name of the access zone. Displays only the settings for shares in the specified zone.

isi smb settings zone modify

Modify SMB settings for a specific access zone.

Syntax

```
isi smb settings zone modify
  [--access-based-share-enum {yes | no}]
  [--revert-access-based-share-enum]
  [--enable-security-signatures {yes | no}]
  [--revert-enable-security-signatures]
  [--reject-unencrypted-access {yes | no}]
  [--revert-reject-unencrypted-access]
  [--require-security-signatures {yes | no}]
  [--revert-require-security-signatures]
  [--server-side-copy {yes | no}]
  [--revert-server-side-copy]
  [--support-multichannel {yes | no}]
  [--revert-support-multichannel]
```



```
[--support-smb2 {yes | no}]
[--revert-support-smb2]
[--support-smb3-encryption {yes | no}]
[--revert-support-smb3-encryption]
[--zone <string>]
[--verbose]
```

Options

--access-based-share-enum {yes | no}

Enumerates only the files and folders that the requesting user has access to.

--revert-access-based-share-enum

Sets the value to the system default for `--access-based-share-enum`.

--enable-security-signatures {yes | no}

Indicates whether the server supports signed SMB packets.

--revert-enable-security-signatures

Sets the value to the system default for `--enable-security-signatures`.

--reject-unencrypted-access {yes | no}

Rejects unencrypted client sessions.

--revert-reject-unencrypted-access

Sets the value to the system default for `--reject-unencrypted-access`.

--require-security-signatures {yes | no}

Specifies whether packet signing is required. If set to `yes`, signing is always required. If set to `no`, signing is not required but clients requesting signing will be allowed to connect if the `--enable-security-signatures` option is set to `yes`.

--revert-require-security-signatures

Sets the value to the system default for `--require-security-signatures`.

--server-side-copy {yes | no}

Enables or disables SMB server-side copy functionality. The default is `yes`.

--revert-server-side-copy

Sets the value to the system default for `--server-side-copy`.

--support-multichannel {yes | no}

Specifies whether Multichannel for SMB 3.0 is enabled on the cluster. SMB Multichannel is enabled by default.

--revert-support-multichannel

Set the value of `--support-multichannel` back to the default system value.

--support-smb2 {yes | no}

Specifies whether to support the SMB 2.0 protocol. The default setting is `yes`.

--revert-support-smb2

Sets the value to the system default for `--support-smb2`.

--support-smb3-encryption {yes | no}

Supports SMBv3 encryption for the access zone. The default setting is `no`.

--revert-support-smb3-encryption

Sets the value to the system default for `--support-smb3-encryption`.

--zone <string>

Access zone.

isi smb settings zone view

View SMB settings for a specific access zone.

Syntax

```
isi smb settings zone view
[--zone <string>]
```

Options

--zone <string>

The name of the access zone for which you are viewing SMB settings.

isi smb shares create

Creates an SMB share.

Syntax

```
isi smb shares create <name> <path>
[--zone <string>]
[--inheritable-path-acl {yes | no}]
[--create-path]
[--host-acl <string>]
[--description <string>]
[--csc-policy {none | documents | manual | programs}]
[--allow-variable-expansion {yes | no}]
[--auto-create-directory {yes | no}]
[--browsable {yes | no}]
[--allow-execute-always {yes | no}]
[--directory-create-mask <integer>]
[--strict-locking {yes | no}]
[--hide-dot-files {yes | no}]
[--impersonate-guest {always | "bad user" | never}]
[--strict-flush {yes | no}]
[--access-based-enumeration {yes | no}]
[--access-based-enumeration-root-only {yes | no}]
[--continuously-available {yes | no}]
[--ca-timeout <integer>]
[--strict-ca-lockout {yes | no}]
[--ca-write-integrity {none | write-read-coherent | full}]
[--mangle-byte-start <string>]
[--file-create-mask <integer>]
[--create-permissions {"default acl" | "inherit mode bits"
| "use create mask and mode"}]
[--mangle-map <string>]
[--impersonate-user <string>]
[--change-notify <string>]
[--oplocks {yes | no}]
[--allow-delete-readonly {yes | no}]
[--directory-create-mode <integer>]
[--ntfs-acl-support {yes | no}]
[--file-create-mode <integer>]
[--file-filtering-enabled {yes | no}]
[--file-filter-type {deny | allow}]
[--file-filter-extensions <string>]
[--smb3-encryption-enabled {yes | no}]
```

Options

<name>

Required. Specifies the name for the new SMB share.

<path>

Required. Specifies the full path of the SMB share to create, beginning at `/ifs`.

--zone <string>

Specifies the access zone the new SMB share is assigned to. If no access zone is specified, the new SMB share is assigned to the default `System` zone.

{--inheritable-path-acl | -i} {yes | no}

If set to `yes`, if the parent directory has an inheritable access control list (ACL), its ACL will be inherited on the share path. The default setting is `no`.

--create-path

Creates the SMB-share path if one doesn't exist.

--host-acl <string>

Specifies the ACL that defines host access. Specify `--host-acl` for each additional host ACL clause.

--description <string>

Specifies a description for the SMB share.

--csc-policy {none | documents | manual | programs}

Sets the client-side caching policy for the share.

--allow-variable-expansion {yes | no}

Specifies automatic expansion of variables for home directories.

--directory-create-mask <integer>

Creates home directories automatically.

--browsable {yes | no}, -b {yes | no}

If set to `yes`, makes the share visible in net view and the browse list. The default setting is `yes`.

--allow-execute-always {yes | no}

If set to `yes`, allows a user with read access to a file to also execute the file. The default setting is `no`.

--directory-create-mask <integer>

Defines which mask bits are applied when a directory is created.

--strict-locking {yes | no}

If set to `yes`, directs the server to check for and enforce file locks. The default setting is `no`.

--hide-dot-files {yes | no}

If set to `yes`, hides files that begin with a decimal—for example, UNIX configuration files. The default setting is `no`.

--impersonate-guest {always | "bad user" | never}

Allows guest access to the share. The acceptable values are `always`, `"bad user"`, and `never`.

--strict-flush {yes | no}

If set to `yes`, flush requests are always honored. The default setting is `yes`.

--access-based-enumeration {yes | no}

If set to `yes`, enables access-based enumeration only on the files and folders that the requesting user can access. The default setting is `no`.

--access-based-enumeration-root-only {yes | no}

If set to `yes`, enables access-based enumeration only on the root directory of the SMB share. The default setting is `no`.

--continuously-available {yes | no}

If set to `yes`, the share allows certain Windows clients to open persistent handles that can be reclaimed after a network disconnect or server failure. The default is `no`.

--ca-timeout <integer>

The amount of time, in seconds, a persistent handle is retained after a client is disconnected or a server fails. The default is 120 seconds.

--strict-ca-lockout {yes | no}

If set to *yes*, prevents a client from opening a file if another client has an open but disconnected persistent handle for that file. If set to *no*, OneFS issues persistent handles, but discards them if any client other than the original opener attempts to open the file. The default is *yes*.

--ca-write-integrity {none | write-read-coherent | full}

Specifies the level of write integrity on continuously available shares:

none	Continuously available writes are not handled differently than other writes to the cluster. If you specify <i>none</i> and a node fails, you may experience data loss without notification. Therefore, we do not recommend this option.
write-read-coherent	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data. This is the default setting.
full	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data, and prevents OneFS from granting SMB clients write-caching and handle-caching leases.

--mangle-byte-start <string>

Specifies the `wchar_t` starting point for automatic invalid byte mangling.

--file-create-mask <integer>

Defines which mask bits are applied when a file is created.

--create-permissions {"default acl" | "inherit mode bits" | "use create mask and mode"}

Sets the default permissions to apply when a file or directory is created. Valid values are "default acl", "inherit mode bits", and "use create mask and mode"

--mangle-map <string>

Maps characters that are valid in OneFS but are not valid in SMB names.

--impersonate-user <string>

Allows all file access to be performed as a specific user. This value must be a fully qualified user name.

--change-notify {norecurse | all | none}

Defines the change notify setting. The acceptable values are *norecurse*, *all*, or *none*.

--oplocks {yes | no}

If set to *yes*, allows oplock requests. The default setting is *yes*.

--allow-delete-readonly {yes | no}

If set to *yes*, allows read-only files to be deleted. The default setting is *no*.

--directory-create-mode <integer>

Defines which mode bits are applied when a directory is created.

--ntfs-acl-support {yes | no}

If set to *yes*, allows ACLs to be stored and edited from SMB clients. The default setting is *yes*.

--file-create-mode <integer>

Defines which mode bits are applied when a file is created.

--file-filtering-enabled {yes | no}

If set to *yes*, enables file filtering at the share level. The default setting is *no*.

--file-filter-type {deny | allow}

If set to *allow*, allows the specified file types to be written to the share. The default setting is *deny*.

--file-filter-extensions <string>

Specifies the list of file extensions to deny or allow writes to the share, depending on the setting of `--file-filter-type`. File types may be specified in a list of comma separated values.

--smb3-encryption-enabled {yes | no}

Enables SMBv3 encryption on the share.

isi smb shares delete

Deletes an SMB share.

Syntax

```
isi smb shares delete <share>
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<share>

Specifies the name of the SMB share to delete.

--zone <string>

Specifies the access zone the SMB share is assigned to. If no access zone is specified, the system deletes the SMB share with the specified name assigned to the default `System` zone, if found.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays more detailed information.

Examples

The following command deletes a share named "test-smb" in the "example-zone" access zone without displaying a warning prompt:

```
isi smb shares delete test-smb --zone example-zone --force
```

isi smb shares list

Displays a list of SMB shares.

Syntax

```
isi smb shares list
  [--zone <string>]
  [--limit <integer>]
  [--sort {name | path | description}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the access zone. Displays all SMB shares in the specified zone. If no access zone is specified, the system displays all SMB shares in the default `System` zone.

{--limit | -l} <integer>

Specifies the maximum number of items to list.

--sort {name | path | description}

Specifies the field to sort items by.

{--descending | -d}

Sorts the data in descending order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

--verbose | -v

Displays more detailed information.

isi smb shares modify

Modifies an SMB share's settings.

Syntax

```
isi smb shares modify <share>
  [--name <string>]
  [--path <path>]
  [--zone <string>]
  [--new-zone <string>]
  [--host-acl <host-acl>]
  [--revert-host-acl]
  [--clear-host-acl]
  [--add-host-acl <string>]
  [--remove-host-acl <string>]
  [--description <string>]
  [--csc-policy {manual | documents | programs | none}]
  [--revert-csc-policy]
  [--allow-variable-expansion {yes | no}]
  [--revert-allow-variable-expansion]
  [--auto-create-directory {yes | no}]
  [--revert-auto-create-directory {yes | no}]
  [--browsable {yes | no}]
  [--revert-browsable]
  [--allow-execute-always {yes | no}]
  [--revert-allow-execute-always]
  [--directory-create-mask <integer>]
  [--revert-directory-create-mask]
  [--strict-locking {yes | no}]
  [--revert-strict-locking]
  [--hide-dot-files {yes | no}]
  [--revert-hide-dot-files]
  [--impersonate-guest {always | "bad user" | never}]
  [--revert-impersonate-guest]
  [--strict-flush {yes | no}]
  [--revert-strict-flush]
  [--access-based-enumeration {yes | no}]
  [--revert-access-based-enumeration]
  [--access-based-enumeration-root-only {yes | no}]
  [--revert-access-based-enumeration-root-only]
  [--ca-timeout <integer>]
  [--revert-ca-timeout]
  [--strict-ca-lockout {yes | no}]
  [--revert-strict-ca-lockout]
  [--ca-write-integrity {none | write-read-coherent | full}]
  [--revert-ca-write-integrity]
  [--mangle-byte-start <integer>]
  [--revert-mangle-byte-start]
```

```

[--file-create-mask <integer>]
[--revert-file-create-mask]
[--create-permissions {"default acl" | "inherit mode bits"
 | "use create mask and mode"}]
[--revert-create-permissions]
[--mangle-map <mangle-map>]
[--revert-mangle-map]
[--clear-mangle-map]
[--add-mangle-map <string>]
[--remove-mangle-map <string>]
[--impersonate-user <string>]
[--revert-impersonate-user]
[--change-notify {all | norecurse | none}]
[--revert-change-notify]
[--oplocks {yes | no}]
[--revert-oplocks]
[--allow-delete-readonly {yes | no}]
[--revert-allow-delete-readonly]
[--directory-create-mode <integer>]
[--revert-directory-create-mode]
[--ntfs-acl-support {yes | no}]
[--revert-ntfs-acl-support]
[--file-create-mode <integer>]
[--revert-file-create-mode]
[--file-filtering-enabled {yes | no}]
[--revert-file-filtering-enabled]
[--file-filter-type {deny | allow}]
[--revert-file-filter-type]
[--file-filter-extensions <string>]
[--clear-file-filter-extensions]
[--add-file-filter-extensions <string>]
[--remove-file-filter-extensions <string>]
[--revert-file-filter-extensions]
[--smb3-encryption-enabled {yes | no}]
[--revert-smb3-encryption-enabled]
[--verbose]

```

Options

<share>

Required. Specifies the name of the SMB share to modify.

--name <name>

Specifies the name for the SMB share.

--path <path>

Specifies a new path for the SMB share, starting in `/ifs`.

--zone <string>

Specifies the access zone that the SMB share is assigned to. If no access zone is specified, the system modifies the SMB share with the specified name assigned to the default `System` zone, if found.

--new-zone <string>

Specifies the new access zone that SMB share will be reassigned to.

--host-acl <host-acl>

An ACL expressing which hosts are allowed access. Specify `--host-acl` for each additional host ACL clause.

--revert-host-acl

Sets the value to the system default for `--host-acl`.

--clear-host-acl

Clears the value of an ACL that expresses which hosts are allowed access.

--add-host-acl <string>

Adds an ACL expressing which hosts are allowed access. Specify `--add-host-acl` for each additional host ACL clause to add.

--remove-host-acl <string>

Removes an ACL expressing which hosts are allowed access. Specify `--remove-host-acl` for each additional host ACL clause to remove.

--description *<string>*

The description for this SMB share.

--csc-policy, -C {**manual** | **documents** | **programs** | **none**}

Specifies the client-side caching policy for the shares.

--revert-csc-policy

Sets the value to the system default for `--csc-policy`.

{**--allow-variable-expansion** | **-a**} {**yes** | **no**}

Allows the automatic expansion of variables for home directories.

--revert-allow-variable-expansion

Sets the value to the system default for `--allow-variable-expansion`.

{**--auto-create-directory** | **-d**} {**yes** | **no**}

Automatically creates home directories.

--revert-auto-create-directory

Sets the value to the system default for `--auto-create-directory`.

{**--browsable** | **-b**} {**yes** | **no**}

The share is visible in the net view and the browse list.

--revert-browsable

Sets the value to the system default for `--browsable`.

--allow-execute-always {**yes** | **no**}

Allows users to execute files they have read rights for.

--revert-allow-execute-always

Sets the value to the system default for `--allow-execute-always`.

--directory-create-mask *<integer>*

Specifies the directory create mask bits.

--revert-directory-create-mask

Sets the value to the system default for `--directory-create-mask`.

--strict-locking {**yes** | **no**}

Specifies whether byte range locks contend against the SMB I/O.

--revert-strict-locking

Sets the value to the system default for `--strict-locking`.

--hide-dot-files {**yes** | **no**}

Hides files and directories that begin with a period ".".

--revert-hide-dot-files

Sets the value to the system default for `--hide-dot-files`.

--impersonate-guest {**always** | **"bad user"** | **never**}

Specifies the condition in which user access is done as the guest account.

--revert-impersonate-guest

Sets the value to the system default for `--impersonate-guest`.

--strict-flush {**yes** | **no**}

Handles the SMB flush operations.

--revert-strict-flush

Sets the value to system default for `--strict-flush`.

--access-based-enumeration {**yes** | **no**}

Specifies to only enumerate files and folders that the requesting user has access to.

--revert-access-based-enumeration

Sets the value to the system default for `--access-based-enumeration`.

--access-based-enumeration-root-only {yes | no}

Specifies access-based enumeration on only the root directory of the share.

--revert-access-based-enumeration-root-only

Sets the value to the system default for `--access-based-enumeration-root-only`.

--ca-timeout <integer>

The amount of time, in seconds, a persistent handle is retained after a client is disconnected or a server fails. The default is 120 seconds.

--revert-ca-timeout

Sets the value to the system default for `--ca-timeout`.

--strict-ca-lockout {yes | no}

If set to `yes`, prevents another client from opening a file if a client has an open but disconnected persistent handle for that file. If set to `no`, OneFS issues persistent handles, but discards them if any client other than the original opener attempts to open the file. This option is only relevant if `--continuously-available` was set to `yes` when the share was created. The default is `yes`.

--revert-strict-ca-lockout

Sets the value to the system default for `--strict-ca-lockout`.

--ca-write-integrity {none | write-read-coherent | full}

Specifies the level of write integrity on continuously available shares:

none	Continuously available writes are not handled differently than other writes to the cluster. If you specify <code>none</code> and a node fails, you may experience data loss without notification. Therefore, we do not recommend this option.
write-read-coherent	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data. This is the default setting.
full	Ensures that writes to the share are moved to persistent storage before a success message is returned to the SMB client that sent the data, and prevents OneFS from granting SMB clients write-caching and handle-caching leases.

--revert-ca-write-integrity

Sets the value to the system default for `--ca-write-integrity`.

--mangle-byte-start <integer>

Specifies the `wchar_t` starting point for automatic byte mangling.

--revert-mangle-byte-start

Sets the value to the system default for `--mangle-byte-start`.

--file-create-mask <integer>

Specifies the file create mask bits.

--revert-file-create-mask

Sets the value to the system default for `--file-create-mask`.

--create-permissions {"default acl" | "inherit mode bits" | "use create mask and mode"}

Sets the create permissions for new files and directories in a share.

--revert-create-permissions

Sets the value to the system default for `--create-permissions`.

--mangle-map <mangle-map>

The character mangle map. Specify `--mangle-map` for each additional character mangle map.

--revert-mangle-map

Sets the value to the system default for `--mangle-map`.

--clear-mangle-map

Clears the value for character mangle map.

--add-mangle-map <string>

Adds a character mangle map. Specify `--add-mangle-map` for each additional Add character mangle map.

--remove-mangle-map <string>

Removes a character mangle map. Specify `--remove-mangle-map` for each additional Remove character mangle map.

--impersonate-user <string>

The user account to be used as a guest account.

--revert-impersonate-user

Sets the value to the system default for `--impersonate-user`.

--change-notify {all | norecuse | none}

Specifies the level of change notification alerts on a share.

--revert-change-notify

Sets the value to the system default for `--change-notify`.

--oplocks {yes | no}

Supports oplocks.

--revert-oplocks

Sets the value for the system default of `--oplocks`.

--allow-delete-readonly {yes | no}

Allows the deletion of read-only files in the share.

--revert-allow-delete-readonly

Sets the value for the system default of `--allow-delete-readonly`.

--directory-create-mode <integer>

Specifies the directory create mode bits.

--revert-directory-create-mode

Sets the value for the system default of `--directory-create-mode`.

--ntfs-acl-support {yes | no}

Supports NTFS ACLs on files and directories.

--revert-ntfs-acl-support

Sets the value for the system default of `--ntfs-acl-support`.

--file-create-mode <integer>

Specifies the file create mode bits.

--revert-file-create-mode

Sets the value for the system default of `--file-create-mode`.

--file-filtering-enabled {yes | no}

If set to `yes`, enables file filtering at the share level. The default setting is `no`.

--revert-file-filtering-enabled

Sets the value for the system default of `--file-filtering-enabled`.

--file-filter-type {deny | allow}

If set to `allow`, allows the specified file types to be written to the share. The default setting is `deny`.

--revert-file-filter-type

Sets the value for the system default of `--file-filter-type`.

--file-filter-extensions <string>

Specifies the list of file types to deny or allow writes to the share, depending on the setting of `--file-filter-type`. File types may be specified in a list of comma separated values.

--clear-file-filter-extensions

Clears the list of file filtering extensions for the share.

--add-file-filter-extensions <string>

Adds entries to the list of file filter extensions. Repeat for each file extension to add.

--remove-file-filter-extensions <string>

Removes entries to the list of file filter extensions. Repeat for each file extension to remove.

--revert-file-filter-extensions

Sets the value for the system default of `--file-filter-extensions`.

`--smb3-encryption-enabled {yes | no}`

Enables SMBv3 encryption on the share.

`--revert-smb3-encryption-enabled`

Sets the value to the system default for `--smb3-encryption-enabled`.

`{--verbose | -v}`

Displays more detailed information.

isi smb shares permission create

Creates permissions for an SMB share.

Syntax

```
isi smb shares permission create <share> {<user> | --group <name>
| --gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
{--run-as-root | --permission-type {allow | deny}
--permission {full | change | read}}
[--zone <zone>]
[--verbose]
```

Options

`<share>`

Specifies the name of the SMB share.

`<user>`

Specifies a user by name.

`--group <name>`

Specifies a group by name.

`--gid <id>`

Specifies a group by UNIX group identifier.

`--uid <id>`

Specifies a user by UNIX user identifier.

`--sid <string>`

Specifies an object by its Windows security identifier.

`--wellknown <string>`

Specifies a well-known user, group, machine, or account name.

`{--permission-type | -d} {deny | allow}`

Specifies whether to allow or deny a permission.

`{--permission | -p} {read | full | change}`

Specifies the level of control to allow or deny.

`--run-as-root {yes | no}`

If set to `yes`, allows the account to run as root. The default setting is `no`.

`--zone <zone>`

Specifies an access zone.

`{--verbose | -v}`

Displays more detailed information.

isi smb shares permission delete

Deletes user or group permissions for an SMB share.

Syntax

```
isi smb shares permission delete <share> {<user> | --group <name>
  |--gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
  [--zone <string>]
  [--force]
  [--verbose]
```

Options

<share>

Required. Specifies the SMB share name.

<user>

Specifies a user by name.

--group <name>

Specifies a group by name.

--gid <id>

Specifies a group by UNIX group identifier.

--uid <id>

Specifies a user by UNIX user identifier.

--sid <string>

Specifies an object by its Windows security identifier.

--wellknown <string>

Specifies a well-known user, group, machine, or account name.

--zone <string>

Specifies an access zone.

{--force | -f}

Specifies that you want the command to execute without prompting for confirmation.

{--verbose | -v}

Displays more detailed information.

isi smb shares permission list

Displays a list of permissions for an SMB share.

Syntax

```
isi smb shares permission list <share>
  [--zone <zone>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
```

Options

<share>

Specifies the name of the SMB share to display.

--zone <zone>

Specifies the access zone to display.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

isi smb shares permission modify

Modifies permissions for an SMB share.

Syntax

```
isi smb shares permission modify <share> {<user> | --group <name>
| --gid <id> | --uid <id> | --sid <string> | --wellknown <string>}
{--run-as-root | --permission-type {allow | deny}
--permission {full | change | read}}
[--zone <zone>]
[--verbose]
```

Options

<share>

Specifies the name of the SMB share.

<user>

Specifies a user by name.

--group <name>

Specifies a group by name.

--gid <id>

Specifies a group by UNIX group identifier.

--uid <id>

Specifies a user by UNIX user identifier.

--sid <string>

Specifies an object by its Windows security identifier.

--wellknown <string>

Specifies a well-known user, group, machine, or account name.

{--permission-type | -d} {deny | allow}

Specifies whether to allow or deny a permission.

{--permission | -p} {read | full | change}

Specifies the level of control to allow or deny.

--run-as-root {*yes* | *no*}

If set to *yes*, allows the account to run as root. The default setting is *no*.

--zone <*zone*>

Specifies an access zone.

{**--verbose** | **-v**}

Displays more detailed information.

isi smb shares permission view

Displays a single permission for an SMB share.

Syntax

```
isi smb shares permission view <share> {<user> |  
--group <name> | --gid <integer>  
| --uid <integer> | --sid <string>  
| --wellknown <string>}  
[--zone <string>]
```

Options

<share>

Specifies the name of the SMB share.

<user>

Specifies a user name.

--group <*name*>

Specifies a group name.

--gid <*integer*>

Specifies a numeric group identifier.

--uid <*integer*>

Specifies a numeric user identifier.

--sid <*string*>

Specifies a security identifier.

--wellknown <*string*>

Specifies a well-known user, group, machine, or account name.

--zone <*string*>

Specifies an access zone.

isi smb shares view

Displays information about an SMB share.

Syntax

```
isi smb shares view <share>  
[--zone <string>]
```

Options

<share>

Specifies the name of the SMB share to view.

--zone <string>

Specifies the access zone that the SMB share is assigned to. If no access zone is specified, the system displays the SMB share with the specified name assigned to the default `SYSTEM` zone, if found.

isi snapshot aliases create

Assigns a snapshot alias to a snapshot or to the live version of the file system.

Syntax

```
isi snapshot aliases create <name> <target>
  [--verbose]
```

Options

<name>

Specifies the a name for the alias.

<target>

Assigns the alias to the specified snapshot or to the live version of the file system.

Specify as a snapshot ID or name. To target the live version of the file system, specify `LIVE`.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases delete

Deletes a snapshot alias.

Syntax

```
isi snapshot aliases delete {<alias> | --all}
  [--force]
  [--verbose]
```

Options

<alias>

Deletes the snapshot alias of the specified name.

Specify as a snapshot-alias name or ID.

--all

Deletes all snapshot aliases.

{--force | -f}

Runs the command without prompting you to confirm that you want to delete the snapshot alias.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases list

Displays a list of snapshot aliases.

Syntax

```
isi snapshot aliases list
  [--limit <integer>]
  [--sort {id | name | target_id | target_name | created}]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

id	Sorts output by the ID of the snapshot alias.
name	Sorts output by the name of the snapshot alias.
target_id	Sorts output by the ID of the snapshot that the snapshot alias is assigned to.
target_name	Sorts output by the name of the snapshot that the snapshot alias is assigned to.
created	Sorts output by the date the snapshot alias was created.

{--descending | -d}

Displays output in reverse order.

--format <output-format>

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases modify

Modifies a snapshot alias.

Syntax

```
isi snapshot aliases modify <alias>
  [--name <name> | --target <snapshot>]
  [--verbose]
```


Options

<alias>

Modifies the specified snapshot alias.
Specify as a snapshot-alias name or ID.

--name <name>

Specifies a new name for the snapshot alias.

--target <snapshot>

Reassigns the snapshot alias to the specified snapshot or the live version of the file system.
Specify as a snapshot ID or name. To target the live version of the file system, specify `LIVE`.

{--verbose | -v}

Displays more detailed information.

isi snapshot aliases view

Displays detailed information about a snapshot alias.

Syntax

```
isi snapshot aliases view <alias>
```

Options

<alias>

Displays detailed information about the specified snapshot alias.
Specify as a snapshot-alias name or ID.

isi snapshot locks create

Creates a snapshot lock.

NOTE: It is recommended that you do not create snapshot locks and do not use this command. If the maximum number of locks on a snapshot is reached, some applications, such as SyncIQ, might not function properly.

Syntax

```
isi snapshot locks create <snapshot>  
  [--comment <string>]  
  [--expires {<timestamp> | <duration>}]  
  [--verbose]
```

Options

<snapshot>

Specifies the name of the snapshot to apply this lock to.

{--comment | -c} <string>

Specifies a comment to describe the lock.
Specify as any string.

{--expires | -x} {<timestamp> | <duration>}

Specifies when the lock will be automatically deleted by the system.

If this option is not specified, the snapshot lock will exist indefinitely.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><time>
```

The following *<time>* values are valid:


Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

{**--verbose** | **-v**}

Displays a message confirming that the snapshot lock was deleted.

isi snapshot locks delete

Deletes a snapshot lock. Deleting a snapshot lock might result in data loss.

 **CAUTION: It is recommended that you do not delete snapshot locks and do not run this command. Deleting a snapshot lock that was created by OneFS might result in data loss.**

Syntax

```
isi snapshot locks delete <snapshot> <id>
  [--force]
  [--verbose]
```

Options

<snapshot>

Deletes a snapshot lock that has been applied to the specified snapshot.

Specify as a snapshot name or ID.

<id>

Modifies the snapshot lock of the specified ID.

{**--force** | **-f**}

Does not prompt you to confirm that you want to delete this snapshot lock.

{**--verbose** | **-v**}

Displays a message confirming that the snapshot lock was deleted.

isi snapshot locks list

Displays a list of all locks applied to a specific snapshot.

Syntax

```
isi snapshot locks list <snapshot>
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<snapshot>

Displays all locks belonging to the specified snapshot.
Specify as a snapshot name.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.
The following values are valid:

id	Sorts output by the ID of a snapshot lock.
comment	Sorts output alphabetically by the description of a snapshot lock.
expires	Sorts output by the length of time that a lock endures on the cluster before being automatically deleted.
count	Sorts output by the number of times that a lock is held.

{--descending | -d}

Displays output in reverse order.

--format <output-format>

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi snapshot locks modify

Modifies the expiration date of a snapshot lock.

CAUTION:

It is recommended that you do not modify the expiration date of snapshot locks and do not run this command. Modifying the expiration date of a snapshot lock that was created by OneFS might result in data loss.

Syntax

```
isi snapshot locks modify <snapshot> <id>
  {--expires {<timestamp> | <duration>} | --clear-expires}
  [--verbose]
```

Options

<snapshot>

Modifies a snapshot lock that has been applied to the specified snapshot.

Specify as a snapshot name or ID.

<id>

Modifies the snapshot lock of the specified ID.

{--expires | -x} {<timestamp> | <duration>}

Specifies when the lock will be automatically deleted by the system.

If this option is not specified, the snapshot lock will exist indefinitely.

Specify <timestamp> in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify <duration> in the following format:

```
<integer><time>
```

The following <time> values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--clear-expires

Removes the duration period for the snapshot lock. If specified, the snapshot lock will exist on the cluster indefinitely.

{--verbose | -v}

Displays a message confirming that the snapshot lock was modified.

Examples

The following command causes a snapshot lock applied to Wednesday_Backup to expire in three weeks:

```
isi snapshot locks modify Wednesday_Backup 1 --expires 3W
```

isi snapshot locks view

Displays information about a snapshot lock.

Syntax

```
isi snapshot locks view <snapshot> <id>
```

Options

<snapshot>

Specifies the snapshot to view locks for.
Specify as a snapshot name or ID.

<id>

Displays the specified lock.
Specify as a snapshot lock ID.

isi snapshot schedules create

Creates a snapshot schedule. A snapshot schedule determines when OneFS regularly generates snapshots on a recurring basis.

Syntax

```
isi snapshot schedules create <name> <path> <pattern> <schedule>
  [--alias <alias>]
  [--duration <duration>]
  [--verbose]
```

Options

<name>

Specifies a name for the snapshot schedule.

<path>

Specifies the path of the directory to include in the snapshots.

<pattern>

Specifies a naming pattern for snapshots created according to the schedule.

<schedule>

Specifies how often snapshots are created.
Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to *<integer>*. For example, you can specify "Every 1st month"

Specify *<day>* as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--alias <alias>

Specifies an alias for the latest snapshot generated based on the schedule. The alias enables you to quickly locate the most recent snapshot that was generated according to the schedule.

Specify as any string.

{--duration | -x} <duration>

Specifies how long snapshots generated according to the schedule are stored on the cluster before OneFS automatically deletes them.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

{--verbose | -v}

Displays a message confirming that the snapshot schedule was created.

isi snapshot schedules delete

Deletes a snapshot schedule. Once a snapshot schedule is deleted, snapshots will no longer be generated according to the schedule. However, snapshots previously generated according to the schedule are not affected.

Syntax

```
isi snapshot schedules delete {<schedule-name> | <all>}
  [--force]
  [--verbose]
```

Options

<schedule-name>

Deletes the specified snapshot schedule.

Specify as a snapshot schedule name or ID.

<all>

Deletes all snapshot schedules.

{--force | -f}

Does not prompt you to confirm that you want to delete this snapshot schedule.

{--verbose | -v}

Displays a message confirming that the snapshot schedule was deleted.

isi snapshot schedules list

Displays a list of all snapshot schedules.

Syntax

```
isi snapshot schedules list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

id	Sorts output by the ID of a snapshot schedule.
name	Sorts output alphabetically by the name of a snapshot schedule.
path	Sorts output by the absolute path of the directory contained by snapshots created according to a schedule.
pattern	Sorts output alphabetically by the snapshot naming pattern assigned to snapshots generated according to a schedule.
schedule	Sorts output alphabetically by the schedule. For example, "Every week" precedes "Yearly on January 3rd"
duration	Sorts output by the length of time that snapshots created according to the schedule endure on the cluster before being automatically deleted.
alias	Sorts output alphabetically by the name of the alias assigned to the most recent snapshot generated according to the schedule.
next_run	Sorts output by the next time that a snapshot will be created according to the schedule.
next_snapshot	Sorts output alphabetically by the name of the snapshot that is scheduled to be created next.

{--descending | -d}

Displays output in reverse order.

--format <output-format>

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi snapshot schedules modify

Modifies the attributes of an existing snapshot schedule.

If you modify a snapshot schedule, snapshots that have already been generated based on the schedule are not affected by the changes.

Syntax

```
isi snapshot schedules modify <schedule-name>
  {--name <name> | --alias <name> | --path <path>
  | --pattern <naming-pattern> | --schedule <schedule>
  | --duration <duration> | --clear-duration}...
  [--verbose]
```

Options

<schedule-name>

Modifies the specified snapshot schedule.

Specify as a snapshot schedule name or ID.

--name <name>

Specifies a new name for the schedule.

Specify as any string.

{--alias | -a} <name>

Specifies an alias for the latest snapshot generated based on the schedule. The alias enables you to quickly locate the most recent snapshot that was generated according to the schedule. If specified, the specified alias will be applied to the next snapshot generated by the schedule, and all subsequently generated snapshots.

Specify as any string.

--path <path>

Specifies a new directory path for this snapshot schedule. If specified, snapshots generated by the schedule will contain only this directory path.

Specify as a directory path.

--pattern <naming-pattern>

Specifies a pattern by which snapshots created according to the schedule are named.

--schedule <schedule>

Specifies how often snapshots are created.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]

- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

{--duration | -x} <duration>

Specifies how long snapshots generated according to the schedule are stored on the cluster before OneFS automatically deletes them.

Specify in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--clear-duration

Removes the duration period for snapshots created according to the schedule. If specified, generated snapshots will exist on the cluster indefinitely.

{--verbose | -v}

Displays a message confirming that the snapshot schedule was modified.

isi snapshot schedules pending list

Displays a list of snapshots that are scheduled to be generated by snapshot schedules.

Syntax

```
isi snapshot schedules pending list
  [--begin <timestamp>]
  [--end <timestamp>]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
```

```
[--no-footer]
[--verbose]
```

Options

{--begin | -b} <timestamp>

Displays only snapshots that are scheduled to be generated after the specified date.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

If this option is not specified, the output displays a list of snapshots that are scheduled to be generated after the current time.

{--end | -e} <time>

Displays only snapshots that are scheduled to be generated before the specified date.

Specify *<time>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

If this option is not specified, the output displays a list of snapshots that are scheduled to be generated before 30 days after the begin time.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format <output-format>

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi snapshot schedules view

Displays information about a snapshot schedule.

Syntax

```
isi snapshot schedules view <schedule-name>
```

Options

<schedule-name>

Displays information about the specified snapshot schedule.

Specify as a snapshot schedule name or ID.

isi snapshot settings modify

Modifies snapshot settings.


Syntax

```
isi snapshot settings modify
  {--service {enable | disable}
  | --autocreate {enable | disable}
  | --autodelete {enable | disable}
  | --reserve <integer>
  | --global-visible-accessible {yes | no}
  | --nfs-root-accessible {yes | no}
  | --nfs-root-visible {yes | no}
  | --nfs-subdir-accessible {yes | no}
  | --smb-root-accessible {yes | no}
  | --smb-root-visible {yes | no}
  | --smb-subdir-accessible {yes | no}
  | --local-root-accessible {yes | no}
  | --local-root-visible {yes | no}
  | --local-subdir-accessible {yes | no}}...
  [--verbose]
```

Options

--service {enable | disable}

Determines whether snapshots can be generated.

 **NOTE: Disabling snapshot generation might cause some OneFS operations to fail. It is recommended that you do not disable this setting.**

--autocreate {enable | disable}

Determines whether snapshots are automatically generated according to snapshot schedules.

Specifying `disable` does not prevent OneFS applications from generating snapshots.

--autodelete {enable | disable}


Determines whether snapshots are automatically deleted according to their expiration dates.

All snapshots that pass their expiration date while this option is disabled will immediately be deleted when the option is enabled again.

--reserve <integer>

Specifies the percentage of the file system to reserve for snapshot usage.

Specify as a positive integer between 1 and 100.

 **NOTE: This option limits only the amount of space available to applications other than SnapshotIQ. It does not limit the amount of space that snapshots are allowed to occupy. Snapshots can occupy more than the specified percentage of system storage space.**

--global-visible-accessible {yes | no}

Specifying `yes` causes snapshot directories and sub-directories to be visible and accessible through all protocols, overriding all other snapshot visibility and accessibility settings. Specifying `no` causes visibility and accessibility settings to be controlled through the other snapshot visibility and accessibility settings.

--nfs-root-accessible {yes | no}

Determines whether snapshot directories are accessible through NFS.

--nfs-root-visible {yes | no}

Determines whether snapshot directories are visible through NFS.

--nfs-subdir-accessible {yes | no}

Determines whether snapshot subdirectories are accessible through NFS.

--smb-root-accessible {yes | no}

Determines whether snapshot directories are accessible through SMB.

--smb-root-visible {yes | no}

Determines whether snapshot directories are visible through SMB.

--smb-subdir-accessible {yes | no}

Determines whether snapshot subdirectories are accessible through SMB.

--local-root-accessible {yes | no}

Determines whether snapshot directories are accessible through the local file system.

--local-root-visible {yes | no}

Determines whether snapshot directories are visible through the local file system.

--local-subdir-accessible {yes | no}

Determines whether snapshot subdirectories are accessible through the local file system.

{--verbose | -v}

Displays a message confirming which snapshot settings were modified.

isi snapshot settings view

Displays current SnapshotIQ settings.

Syntax

```
isi snapshot settings view
```

Options

There are no options for this command.

isi snapshot snapshots create

Creates a snapshot of a directory.

Syntax

```
isi snapshot snapshots create <path>  
  [--name <name>]  
  [--expires {<timestamp> | <duration>}]  
  [--alias <name>]  
  [--verbose]
```

Options

<path>

Specifies the path of the directory to include in this snapshot.

--name <name>

Specifies a name for the snapshot.

{--expires | -x} {<timestamp> | <duration>}

Specifies when OneFS will automatically delete this snapshot.

If this option is not specified, the snapshot will exist indefinitely.

Specify *<timestamp>* in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

{--alias | -a} <name>

Specifies an alias for this snapshot. A snapshot alias is an alternate name for a snapshot.

Specify as any string.

{--verbose | -v}

Displays a message confirming that the snapshot was created.

isi snapshot snapshots delete

Deletes a snapshot. If a snapshot is deleted, it can no longer be accessed by a user or the system.

Syntax

```
isi snapshot snapshots delete [--all | --snapshot <snapshot>
| --schedule <schedule> | --type <type>]
[--force]
[--verbose]
```

Options

--all

Deletes all snapshots.

--snapshot <snapshot>

Deletes the specified snapshot.

Specify as a snapshot name or ID.

--schedule <schedule>

Deletes all snapshots created according to the specified schedule.

Specify as a snapshot schedule name or ID.

--type <type>

Deletes all snapshots of the specified type.

The following types are valid:

alias	Deletes all snapshot aliases.
real	Deletes all snapshots.

{--force | -f}

Does not prompt you to confirm that you want to delete the snapshot.

{--verbose | -v}

Displays a message confirming that the snapshot was deleted.

Examples

The following command deletes newSnap1:

```
isi snapshot snapshots delete --snapshot newSnap1
```

isi snapshot snapshots list

Displays a list of all snapshots and snapshot aliases.

Syntax

```
isi snapshot snapshots list
  [--state <state>]
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--state <state>

Displays only snapshots and snapshot aliases that exist in the specified state.

The following states are valid:

all	Displays all snapshots and snapshot aliases that are currently occupying space on the cluster.
active	Displays only snapshots and snapshot aliases that have not been deleted.
deleting	Displays only snapshots that have been deleted but are still occupying space on the cluster. The space occupied by deleted snapshots will be freed the next time the snapshot delete job is run.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts command output by the specified attribute.

The following attributes are valid:

id	Sorts output by the ID of a snapshot.
name	Sorts output alphabetically by the name of a snapshot.
path	Sorts output by the absolute path of the directory contained in a snapshot.
has_locks	Sorts output by whether any snapshot locks have been applied to a snapshot.
schedule	If a snapshot was generated according to a schedule, sorts output alphabetically by the name of the snapshot schedule.
target_id	If a snapshot is an alias, sorts output by the snapshot ID of the target snapshot instead of the snapshot ID of the alias.
target_name	If a snapshot is an alias, sorts output by the name of the target snapshot instead of the name of the alias.

created	Sorts output by the time that a snapshot was created.
expires	Sorts output by the time at which a snapshot is scheduled to be automatically deleted.
size	Sorts output by the amount of disk space taken up by a snapshot.
shadow_bytes	Sorts output based on the amount of data that a snapshot references from shadow stores. Snapshots reference shadow store data if a file contained in a snapshot is cloned or a snapshot is taken of a cloned file.
pct_reserve	Sorts output by the percentage of the snapshot reserve that a snapshot occupies.
pct_filesystem	Sorts output by the percent of the file system that a snapshot occupies.
state	Sorts output based on the state of snapshots.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

isi snapshot snapshots modify

Modifies attributes of a snapshot or snapshot alias.

Syntax

```
isi snapshot snapshots modify <snapshot>
  {--name <name> | --expires {<timestamp> | <duration>}
  | --clear-expires | --alias <name>}...
  [--verbose]
```

Options

<snapshot>

Modifies the specified snapshot or snapshot alias.

Specify as the name or ID of a snapshot or snapshot alias.

--name <name>

Specifies a new name for the snapshot or snapshot alias.

Specify as any string.

{--expires | -x} {<timestamp> | <duration>}

Specifies when OneFS will automatically delete this snapshot.

Specify <timestamp> in the following format:

```
<yyyy>-<mm>-<dd>[T<HH>:<MM>[:<SS>]]
```

Specify *<duration>* in the following format:

```
<integer><time>
```

The following *<time>* values are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

You cannot modify the expiration date of a snapshot alias.

--clear-expires

Removes the expiration date from the snapshot, allowing the snapshot to exist on the cluster indefinitely.

You cannot modify the expiration date of a snapshot alias.

{--alias | -a} <name>

Specifies an alias for the snapshot. A snapshot alias is an alternate name for a snapshot. You cannot specify an alias for a snapshot alias.

Specify as any string.

{--verbose | -v}

Displays a message confirming that the snapshot or snapshot alias was modified.

isi snapshot snapshots view

Displays the properties of an individual snapshot.

Syntax

```
isi snapshot snapshots view <snapshot>
```

Options

<snapshot>

Displays information about the specified snapshot.

Specify as a snapshot name or ID.

isi snmp settings modify

Modify SNMP settings for a cluster.

Syntax

```
isi snmp settings modify
  [--service {yes | no}]
  [--system-location <string>]
  [--revert-system-location]
  [--system-contact <string>]
  [--revert-system-contact]
  [--snmp-v1-v2c-access {yes | no}]
  [--revert-snmp-v1-v2c-access]
  [--read-only-community <string>]
  [--revert-read-only-community]
```



```
[--snmp-v3-access {yes | no}]
[--revert-snmp-v3-access]
[--snmp-v3-read-only-user <string>]
[--revert-snmp-v3-read-only-user]
[--snmp-v3-password <string>]
[--revert-snmp-v3-password]
[--set-snmp-v3-password]
[--verbose]
```

Options

- service {yes | no}**
Enables or disables the SNMP service.
- system-location <string>**
The location of the SNMP system.
- revert-system-location**
Sets `--system-location` to the system default.
- system-contact <string>**
A valid email address for the system owner.
- revert-system-contact**
Sets `--system-contact` to the system default.
- snmp-v1-v2c-access {yes | no}**
Enables or disables the SNMP v1 and v2c protocols.
- revert-snmp-v1-v2c-access**
Sets `--snmp-v1-v2c-access` to the system default.
- {--read-only-community | -c} <string>**
The name of the read-only community.
- revert-read-only-community**
Sets `--read-only-community` to the system default.
- snmp-v3-access {yes | no}**
Enables or disables SNMP v3.
- revert-snmp-v3-access**
Sets `--snmp-v3-access` to the system default.
- {--snmp-v3-read-only-user | -u} <string>**
The read-only user for SNMP v3 read requests.
- revert-snmp-v3-read-only-user**
Sets `--snmp-v3-read-only-user` to the system default.
- {--snmp-v3-password | -p} <string>**
Modify the SNMP v3 password.
- revert-snmp-v3-password**
Sets `--snmp-v3-password` to the system default.
- set-snmp-v3-password**
Specify `--snmp-v3-password` interactively.

isi snmp settings view

View SNMP settings for the cluster.

Syntax

```
isi snmp settings view
```

Example

To view the currently-configured SNMP settings, run the following command:

```
isi snmp settings view
```

The system displays output similar to the following example:

```
System Location: unset
System Contact: unset@unset.invalid
Snmp V1 V2C Access: Yes
Read Only Community: I$ilonpublic
Snmp V3 Access: No
Snmp V3 Read Only User: general
SNMP Service Enabled: No
```

isi auth settings modify

Modify the SSH server settings.

Syntax

```
isi ssh settings modify
[--ciphers <string>]
[--host-key-algorithms <string>]
[--ignore-rhosts <boolean>]
[--kex-algorithms <string>]
[--login-grace-time <duration>]
[--log-level <string>]
[--macs <string>]
[--max-auth-tries <integer>]
[--max-sessions <integer>]
[--max-startups <string>]
[--permit-empty-passwords <boolean>]
[--permit-root-login <boolean>]
[--port <integer>]
[--print-motd <boolean>]
[--pubkey-accepted-key-types <string>]
[--strict-modes <boolean>]
[--subsystem <string>]
[--syslog-facility <string>]
[--tcp-keep-alive <boolean>]
[--user-auth-method (password | publickey | both | any)]
[--match <string>]
[{{--verbose | -v}}]
[{{--help | -h}}
```

Options

--ciphers <string>

Specifies the ciphers allowed for protocol version 2.

--host-key-algorithms <string>
 Specifies the protocol version 2 host key algorithms the server offers.

--ignore-rhosts <boolean>
 Enables ignoring .rhosts and .shosts files.

--kex-algorithms <string>
 Specifies the available KEX algorithms.

--login-grace-time <duration>
 Specifies the length of time before idle log in fails.

--log-level <string>
 Specifies the log level when logging messages from sshd.

--macs <string>
 Specifies the available MAC algorithms.

--max-auth-tries <integer>
 Specifies the maximum number of authentication attempts per connection.

--max-sessions <integer>
 Specifies the maximum number of open sessions permitted per network connection.

--max-startups <integer>
 Specify maximum number of unauthenticated connections.

--permit-empty-passwords <integer>
 Enables empty passwords if password authentication is allowed.

--permit-root-login <boolean>
 Enables root SSH login.

--port <integer>
 Specifies the port sshd should listen on.

--print-motd <boolean>
 Enables printing /etc/motd when a user logs in.

--pubkey-accepted-key-types <string>
 Specifies the accepted public key types.

--strict-modes <boolean>
 Specifies if sshd should check home directory permissions before accepting login.

--subsystem <string>
 Specifies an external subsystem.

--syslog-facility <string>
 Specifies the facility code when logging messages from sshd.

--tcp-keep-alive <boolean>
 Enables sending TCP keep alive messages.

--user-auth-method password publickey bothany
 Specifies how a user should authenticate.

--match <string>
 Specifies match blocks.

--verbose | -v
 Displays more detailed information.

--help | -h
 Displays help for this command.

isi statistics client

Displays the most active, by throughput, clients accessing the cluster for each supported protocol. You can specify options to track access by user, for example, more than one user on the same client host access the cluster.

Syntax

```
isi statistics client
  [--numeric]
  [--local-addresses <string>]
  [--local_names <string>]
  [--remote_addresses <integer>]
  [--remote_names <string>]
  [--user-ids <integer>]
  [--user-names <string>]
  [--protocols <value>]
  [--classes <string>]
  [--nodes <value>]
  [--degraded]
  [--nohumanize]
  [--interval <integer>]
  [--repeat <integer>]
  [--limit]
  [--long]
  [--totalby <column>]
  [--output <column>]
  [--sort <column>]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--numeric

If text identifiers of local hosts, remote clients, or users are in the list of columns to display (the default setting is for them to be displayed), display the unresolved numeric equivalent of these columns.

--local-addresses <string>

Specifies local IP addresses for which statistics will be reported.

--local-names <string>

Specifies local host names for which statistics will be reported.

--remote-addresses <string>

Specifies remote IP addresses for which statistics will be reported.

--remote-names <string>

Specifies remote client names for which statistics will be reported.

--user-ids <string>

Specifies user ids for which statistics will be reported. The default setting is all users.

--user-names <string>

Specifies user names for which statistics will be reported. The default setting is all users.

--protocols <value>

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list, for example `--protocols http,papi`. The following values are valid:

- all
- external
- ftp
- hdfs
- http

- internal
- irp
- jobd
- lsass_in
- lsass_out
- nlm
- nfs3
- nfs4
- papi
- siq
- smb1
- smb2

--classes <string>

Specify which operation classes to report statistics on. The default setting is all classes. The following values are valid:

- other** File-system information for other uncategorized operations
- write** File and stream writing
- read** File and stream reading
- namespace_read** Attribute stat and ACL reads; lookup directory reading
- namespace_writ** Renames; attribute setting; permission time and ACL writes
- e**

{--nodes | -n} <node>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`. The following values are valid:

- all
- <int>

{--degraded | -d}

Causes the report to continue if some nodes do not respond.

--nohumanize

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

{--interval | -i} <float>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To run the report to run indefinitely, specify -1.**

{--limit | -l} <integer>

Limits the number of statistics to display.

--long

Displays all possible columns.

--totalby <column>

Aggregates results according to specified fields. The following values are valid:

- Node
- {Proto | protocol}
- Class
- {UserId | user.id}
- {UserName | user.name}
- {LocalAddr | local_addr}
- {LocalName | local_name}

- {RemoteAddr | remote_addr}
- {RemoteName | remote_name}

--output <column>

Specifies which columns to display. The following values are valid:

{NumOps num_operations}	Displays the number of times an operation has been performed.
{Ops operation_rate}	Displays the rate at which an operation has been performed. Displayed in operations per second.
{InMax in_max}	Displays the maximum input (received) bytes for an operation.
{InMin in_min}	Displays the minimum input (received) bytes for an operation.
In	Displays the rate of input for an operation since the last time <code>isi statistics</code> collected the data. Displayed in bytes per second.
{InAvg in_avg}	Displays the average input (received) bytes for an operation.
{OutMax out_max}	Displays the maximum output (sent) bytes for an operation.
{OutMin out_min}	Displays the minimum output (sent) bytes for an operation.
Out	Displays the rate of output for an operation since the last time <code>isi statistics</code> collected the data. Displayed in bytes per second.
{OutAvg out_avg}	Displays the average output (sent) bytes for an operation.
{TimeMax time_max}	Displays the maximum elapsed time taken to complete an operation. Displayed in microseconds.
{TimeMin time_min}	Displays the minimum elapsed time taken to complete an operation. Displayed in microseconds.
{TimeAvg time_avg}	Displays the average elapsed time taken to complete an operation. Displayed in microseconds.
Node	Displays the node on which the operation was performed.
{Proto protocol}	Displays the protocol of the operation.
Class	Displays the class of the operation.
{UserID user.id}	Displays the numeric UID of the user issuing the operation request.
{UserName user.name}	Displays the resolved text name of the UserID. If resolution cannot be performed, UNKNOWN is displayed.
{LocalAddr local_addr}	Displays the local IP address of the user issuing the operation request.
{LocalName local_name}	Displays the local host name of the user issuing the operation request.
{RemoteAddr remote_addr}	Displays the remote IP address of the user issuing the operation request.
{RemoteName remote_name}	Displays the remote client name of the user issuing the operation request.

--sort <column>


Specifies how rows are ordered. The following values are valid:

- {NumOps | num_operations}

- {Ops | operation_rate}
- {InMax | in_max}
- {InMin | in_min}
- In
- {InAvg | in_avg}
- {OutMax | out_max}
- {OutMin | out_min}
- Out
- {OutAvg | out_avg}
- {TimeMax | time_max}
- {TimeMin | time_min}
- {TimeAvg | time_avg}
- Node
- {Proto | protocol}
- Class
- {UserID | user.id}
- {UserName | user.name}
- {LocalAddr | local_addr}
- {LocalName | local_name}
- {RemoteAddr | remote_addr}
- {RemoteName | remote_name}

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--noheader | -a}

Displays data without column headings.

{--no-footer | -z}

Displays data without footers.

{--verbose | -v}

Displays more detailed information.

isi statistics data-reduction

View statistics related to data reduction through compression and deduplication.

Syntax

```
isi statistics data-reduction <action>
  [--resolution <integer>]
  [--help | -h]
```

Options

--resolution <integer>

Specifies the minimum interval between series data points in seconds.

{--help | -h}

Displays help for this command.

isi statistics data-reduction view

View statistics related to data reduction through compression and deduplication.

Syntax

```
isi statistics data-reduction view
  [--resolution <integer>]
  [--help | -h]
```

Options

--resolution <integer>

Specifies the minimum interval between series data points in seconds.

{--help | -h}

Displays help for this command.

isi statistics drive

Displays performance information by drive.

Syntax

```
isi statistics drive
  [--type <value>]
  [--nodes <value>]
  [--degraded]
  [--nohumanize]
  [--interval <integer>]
  [--repeat <integer>]
  [--limit <integer>]
  [--long]
  [--output <column>]
  [--sort <column>]
  [--format] [--top]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--type <string>

Specifies the drive types for which statistics will be reported. The default setting is all drives. The following values are valid:

- sata
- sas
- ssd

{--nodes | -n} <node>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`. The following values are valid:

- all
- <int>

{--degraded | -d}

Sets the report to continue running if some nodes do not respond.

--nohumanize

Displays all data in base quantities, without dynamic conversion. If set, this parameter also disables the display of units within the data table.

{--interval | -I} <integer>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To set the report to run indefinitely, specify -1.**

{--limit | -l} <integer>

Limits the number of statistics to display.

--long

Displays all possible columns.

--output <column>

Specifies which columns to display. The following values are valid:

- {Timestamp | time}
- {Drive | drive_id}
- {Type | }
- {BytesIn | bytes_in}
- {SizeIn | xfer_size_in}
- {OpsOut | xfers_out}
- {BytesOut | bytes_out}
- {SizeOut | xfer_size_out}
- {TimeAvg | access_latency}
- {Slow | access_slow}
- {TimeInQ | iosched_latency}
- {Queued | iosched_queue}
- {Busy | used_bytes_percent}
- {Inodes | used_inodes}


--sort <column>

Specifies how the rows are ordered. The following values are valid:

- {Timestamp | time}
- {Drive | drive_id}
- {Type | }
- {BytesIn | bytes_in}
- {SizeIn | xfer_size_in}
- {OpsOut | xfers_out}
- {BytesOut | bytes_out}
- {SizeOut | xfer_size_out}
- {TimeAvg | access_latency}
- {Slow | access_slow}
- {TimeInQ | iosched_latency}
- {Queued | iosched_queue}
- {Busy | used_bytes_percent}
- {Inodes | used_inodes}

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--noheader | -a}

Displays data without column headings.

{ **--no-footer** | **-z** }

Displays data without footers.

{ **--verbose** | **-v** }

Displays more detailed information.

isi statistics heat

Displays the most active /ifs paths for various metrics.

Syntax

```
isi statistics heat
  [--events <string>]
  [--pathdepth <integer>]
  [--maxpath <integer>]
  [--classes <string>]
  [--numeric]
  [--nodes <value>]
  [--degraded]
  [--nohumanize]
  [--interval <integer>]
  [--repeat <integer>]
  [--limit <integer>]
  [--long]
  [--totalby <column>]
  [--output <column>]
  [--sort <column>]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--events <string>

Specifies which event types for the specified information are reported. The following values are valid:

blocked	Access to the LIN was blocked waiting for a resource to be released by another operation. Class is <code>other</code> .
contended	A LIN is experiencing cross-node contention; it is being accessed simultaneously through multiple nodes. Class is <code>other</code> .
deadlocked	The attempt to lock the LIN resulted in deadlock. Class is <code>other</code> .
getattr	A file or directory attribute has been read. Class is <code>namespace_read</code> .
link	The LIN has been linked into the file system; the LIN associated with this event is the parent directory and not the linked LIN. Class is <code>namespace_write</code> .
lock	The LIN is locked. Class is <code>other</code> .
lookup	A name is looked up in a directory; the LIN for the directory searched is the one associated with the event. Class is <code>namespace_read</code> .
read	A read was performed. Class is <code>read</code> .
rename	A file or directory was renamed. The LIN associated with this event is the directory where the rename took place for either the source directory or the destination directory, if they differ. Class is <code>namespace_write</code> .
setattr	A file or directory attribute has been added, modified, or deleted. Class is <code>namespace_write</code> .

unlink A file or directory has been unlinked from the file system, the LIN associated with this event is the parent directory of the removed item. Class is `namespace_write`.

write A write was performed. Class is `write`.

-pathdepth <integer>

Reduces paths to the specified depth.

--maxpath <integer>

Specifies the maximum path length to look up in the file system.

--classes <string>

Specifies which classes for the specified information will be reported. The default setting is all classes. The following values are valid:

write File and stream writing

read File and stream reading

namespace_write Renames; attribute setting; permission, time, and ACL writes

namespace_read Attribute, stat, and ACL reads; lookup, directory reading

other File-system information

--numeric

If text identifiers of local hosts, remote clients, or users are in the list of columns to display (the default setting is for them to be displayed), display the unresolved numeric equivalent of these columns.

{ --nodes | -n } <value>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list—for example, **--nodes 1,2**. The default value is `all`. The following values are valid:

- `all`
- `<int>`

{--degraded | -d}

Sets the report to continue running if some nodes do not respond.

--nohumanize

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

{--interval | -I} <integer>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To set the report to run indefinitely, specify -1.**

--limit <integer>

Displays only the specified number of entries after totaling and ordering.

--long

Displays all possible columns.

--totalby <column>

Aggregates results according to specified fields. The following values are valid:

- `Node`
- `{Event | event_name}`
- `{Class | class_name}`
- `LIN`
- `Path`

--output <column>

Specifies the columns to display. The following values are valid:

{Ops operation_rate}	Displays the rate at which an operation has been performed. Displayed in operations per second.
Node	Displays the node on which the operation was performed.
{Event event_name}	Displays the name of the event.
{Class class_name}	Displays the class of the operation.
LIN	Displays the LIN for the file or directory associated with the event.
Path	Displays the path associated with the event LIN.


--sort <column>

Specifies how rows are ordered. The following values are valid:

- {Ops | operation_rate}
- Node
- {Event | event_name}
- {Class | class_name}
- LIN
- Path

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--noheader | -a}

Displays data without column headings.

{--no-footer | -z}

Displays data without footers.

{--verbose | -v}

Displays more detailed information.

isi statistics list keys

Displays a list of all available keys.

Syntax

```
isi statistics list operations
  [--limit]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```


Options

{--limit | -l}<integer>

Limits the number of statistics to display.

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE:** If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--noheader | -a}`
Displays data without column headings.

`{--no-footer | -z}`
Displays data without footers.

`{--verbose | -v}`
Displays more detailed information.

isi statistics list operations

Displays a list of valid arguments for the `--operations` option.

Syntax

```
isi statistics list operations
  [--protocols <value>]
  [--limit]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

`--protocols <value>`


Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list, for example `--protocols http,papi`. The following values are valid:

- nfs3
- smb1
- nlm
- ftp
- http
- siq
- smb2
- nfs4
- papi
- jobd
- irp
- lsass_in
- lsass_out
- hdfs
- console
- ssh

`{--limit | -l}<integer>`
Limits the number of statistics to display.

`--format {table | json | csv | list | top}`

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE:** If you specify `--top` without `--repeat`, the report runs indefinitely.

`{--noheader | -a}`

Displays data without column headings.

{ --no-footer | -z }

Displays data without footers.

{ --verbose | -v }

Displays more detailed information.

isi statistics protocol

Displays statistics by protocol, such as NFSv3 and HTTP.

Syntax

```
isi statistics protocol
  [--classes <class>...]
  [--protocols <protocol>...]
  [--operations <operation>...]
  [--zero]
  [--nodes <value>]
  [--degraded]
  [--nohumanize]
  [--interval <integer>]
  [--repeat <integer>]
  [--limit]
  [--long]
  [--totalby <column>...]
  [--output <column>...]
  [--nodes <value>]
  [--sort <column>...]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--classes <class>

Specifies which operation classes to report statistics on. The following values are valid:

other	File-system information. Multiple values can be specified in a comma-separated list.
write	File and stream writing
read	File and stream reading
create	File link node stream and directory creation
delete	File link node stream and directory deletion
namespace_read	Attribute stat and ACL reading; lookup directory reading
namespace_write	Renames; attribute setting; permission time and ACL writes
file_state	Open, close; locking; acquire, release, break, check; notification
session_state	Negotiation inquiry or manipulation of protocol connection or session state

--protocols <value>

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list, for example `--protocols http,papi`. The following values are valid:

- nfs3
- smb1
- nlm
- ftp

- http
- siq
- smb2
- nfs4
- papi
- jobd
- irp
- lsass_in
- lsass_out
- hdfs
- all
- internal
- external

--operations <operation>

Specifies the operations on which statistics are reported. To view a list of valid values, run the `isi statistics list operations` command. Multiple values can be specified in a comma-separated list.

--zero

Shows table entries with no values.

{--nodes | -n} <node>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`. The following values are valid:

- all
- <int>

{--degraded | -d}

Causes the report to continue running if some nodes do not respond.

--nohumanize

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units in the data table.

{--interval | -i} <float>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To set the report to run indefinitely, specify -1.**

{--limit | -l} <integer>

Limits the number of statistics to display.

--long

Displays all possible columns.

--totalby <column>

Aggregates results according to specified fields. The following values are valid:

- Node
- {Proto | protocol}
- Class
- {Op | operation}

--output <column>

Specifies which columns to display. The following values are valid:

{timestamp | time} Displays the time at which the `isi statistics` tool last gathered data. Displayed in POSIX time (number of seconds elapsed since January 1, 1970).

Specify <time-and-date> in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

Specify <time> as one of the following values.

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
h	Specifies hours
s	Specifies seconds

{NumOps operation_count}	Displays the number of times an operation has been performed.
{Ops operation_rate }	Displays the rate at which an operation has been performed. Displayed in operations per second.
{InMax in_max}	Displays the maximum input (received) bytes for an operation.
{InMin in_min}	Displays the minimum input (received) bytes for an operation.
In	Displays the rate of input for an operation since the last time isi statistics collected the data. Displayed in bytes per second.
{InAvg in_avg}	Displays the average input (received) bytes for an operation.
{InStdDev in_standard_dev}	Displays the standard deviation of the input (received) bytes for an operation. Displayed in bytes.
{OutMax out_max}	Displays the maximum output (sent) bytes for an operation.
{OutMin out_min}	Displays the minimum output (sent) bytes for an operation.
Out	Displays the rate of output for an operation since the last time isi statistics collected the data. Displayed in bytes per second.
{OutAvg out_avg}	Displays the average output (sent) bytes for an operation.
{OutStdDev out_standard_dev}	Displays the standard deviation of the output (sent) bytes for an operation. Displayed in bytes.
{TimeMax time_max}	Displays the maximum elapsed time taken to complete an operation. Displayed in microseconds.
{TimeMin time_min}	Displays the minimum elapsed time taken to complete an operation. Displayed in microseconds.
{TimeAvg time_avg}	Displays the average elapsed time taken to complete an operation. Displayed in microseconds.
{TimeStdDev time_standard_dev}	Displays the elapsed time taken to complete an operation as a standard deviation from the mean elapsed time.
Node	Displays the node on which the operation was performed.
{Proto protocol}	Displays the protocol of the operation.

Class Displays the class of the operation.

{Op | operation} Displays the name of the operation


--sort <column>

Specifies how rows are ordered. The following values are valid:

- Class
- In
- InAvg | in_avg}
- InMax | in_max}
- InMin | in_min}
- InStdDev | in_standard_dev}
- Node
- NumOps | operation_count}
- Op | operation}
- Ops | operation_rate}
- Out
- OutAvg | out_avg}
- OutMax | out_max}
- OutMin | out_min}
- OutStdDev | out_standard_dev}
- Proto | protocol}
- TimeAvg | time_avg}
- TimeMax | time_max}
- TimeMin | time_min}
- TimeStamp | time}
- TimeStdDev | time_standard_dev}

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--noheader | -a}

Displays data without column headings.

{--no-footer | -z}

Displays data without footers.

{--verbose | -v}

Displays more detailed information.

isi statistics pstat

Displays a selection of cluster-wide and protocol data.

Syntax

```
isi statistics pstat
  [--protocol <protocol>]
  [--degraded]
  [--interval <integer>]
  [--repeat <integer>]
  [--format]
  [--verbose]
```

Options

--protocols <value>

Specifies which protocols to report statistics on. Multiple values can be specified in a comma-separated list, for example `--protocols http,papi`. The following values are valid:

- nfs3
- smb1
- nlm
- ftp
- http
- siq
- smb2
- nfs4
- papi
- jobd
- irp
- lsass_in
- lsass_out
- hdfs

{--degraded | -d}

Sets the report to continue running if some nodes do not respond.

{--interval | -i} <float>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To set the report to run indefinitely, specify -1.**

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--verbose | -v}

Displays more detailed information.

isi statistics query current

Displays current statistics.

Syntax

```
isi statistics query history
  [--keys <string>]
  [--substr]
  [--raw]
  [--nodes <value>]
  [--degraded]
  [--interval <number>]
  [--repeat <number>]
  [--limit]
  [--long]
  [--format]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--keys <string> ...

Specifies which statistics should be reported for requested nodes, where the value for <string> is a statistics key. Use the `isi statistics list keys` command for a complete listing of statistics keys.

--substr

Matches the statistics for `'.*<key>.*'` for every key specified with `--keys`.

--raw

Outputs complex objects as hex.

{ --nodes | -n } <node>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`. The following values are valid:

- `all`
- `<int>`

{ --degraded | -d }

Sets the report to continue running if some nodes do not respond.

{ --interval | -i } <float>

Reports data at the interval specified in seconds.

{ --repeat | -r } <integer>

Specifies how many times to run the report before quitting.

 **NOTE: To set the report to run indefinitely, specify `-1`.**

{ --limit | -l } <integer>

Limits the number of statistics to display.

--long

Displays all possible columns.

--format { table | json | csv | list | top }

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify `--top` without `--repeat`, the report runs indefinitely.**

{ --noheader | -a }

Displays data without column headings.

{ --no-footer | -z }

Displays data without footers.

{ --verbose | -v }

Displays more detailed information.

isi statistics query history

Displays available historical statistics. Not all statistics are configured to support a historical query.

Syntax

```
isi statistics query history
  [--keys <string>]
  [--substr]
  [--begin <integer>]
  [--end <integer>]
  [--resolution <number>]
  [--memory-only]
  [--raw]
  [--nodes <value>]
```

```
[--degraded]
[--nohumanize]
[--interval <number>]
[--repeat <number>]
[--limit]
[--format]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--keys <string> ...

Specifies which statistics should be reported for requested nodes, where the value for *<string>* is a statistics key. Use the `isi statistics list keys` command for a complete listing of statistics keys.

--substr

Matches the statistics for `'.*<key>.*'` for every key specified with `--keys`.

--begin <time>

Specifies begin time in UNIX Epoch timestamp format.

--end <time>

Specifies end time in UNIX Epoch timestamp format.

--resolution <integer>

Specifies the minimum interval between series data points in seconds.

--memory-only

Retrieves only the statistics in memory, not those persisted to disk.

--raw

Outputs complex objects as hex.

{ --nodes | -n } <node>

Specifies which nodes to report statistics on. Multiple values can be specified in a comma-separated list, for example, `--nodes 1,2`. The default value is `all`. The following values are valid:

- `all`
- `<int>`

{--degraded | -d}

Sets the report to continue running if some nodes do not respond.

--nohumanize

Displays all data in base quantities, without dynamic conversion. If set, this option also disables the display of units within the data table.

{--interval | -i} <float>

Reports data at the interval specified in seconds.

{--repeat | -r} <integer>

Specifies how many times to run the report before quitting.


 **NOTE: To set the report to run indefinitely, specify -1.**

{--limit | -l} <integer>

Limits the number of statistics to display.

--format {table | json | csv | list | top}

Displays output in table, JavaScript Object Notation (JSON), comma-separated value (CSV), list format, or top-style display where data is continuously overwritten in a single table.

 **NOTE: If you specify --top without --repeat, the report runs indefinitely.**

{--noheader | -a}

Displays data without column headings.

{--no-footer | -z}

Displays data without footers.

{--verbose | -v}

Displays more detailed information.

isi status

Displays information about the current status of the nodes on the cluster.

Syntax

```
isi status
  [--all-nodes | -a]
  [--node | -n <integer>]
  [--all-nodepools | -p]
  [--nodepool | -l <string>]
  [--quiet | -q]
  [--verbose | -v]
```

Options

--all-nodes | -a

Display node-specific status for all nodes on a cluster.

--node | -n <integer>

Display node-specific status for the node specified by its logical node number (LNN).

--all-nodepools | -p

Display node pool status for all node pools in the cluster.

--nodepool | -l <string>

Display node pool status for the specified node pool.

--quiet | -q

Display less detailed information.

--verbose | -v

Display more detailed information for the `--nodepool` or `--all-nodepools` options.

isi storagepool compatibilities active create

Creates a compatibility to enable an unprovisioned node to join a node pool.

Syntax

```
isi storagepool compatibilities active create <class-1> <class-2>
  [--assess {yes|no}]
  [--verbose]
  [--force]
```

Options

<class-1>

An existing node pool class, one of S200 or X400.

<class-2>

The node class that is compatible with the existing node pool, one of S210 or X410. Note that S210 nodes are only compatible with S200 node pools, and X410 nodes are only compatible with X400 node pools.

{--assess | -a} {yes | no}

Checks whether the compatibility is valid without actually creating the compatibility.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Performs the action without asking for confirmation.

Examples

The following command creates a compatibility between S200 and S210 nodes without asking for confirmation:

```
isi storagepool compatibilities active create S200 S210 --force
```

isi storagepool compatibilities active delete

Deletes a node compatibility. If fewer than three compatible nodes had been added to an existing node pool, they are removed and become unprovisioned.

Syntax

```
isi storagepool compatibilities active delete <ID>  
  [--assess {yes | no}]  
  [--verbose]  
  [--force]
```

Options

<ID>

The ID number of the compatibility. You can use the `isi storagepool compatibilities active list` command to view the ID numbers of active compatibilities.

{--assess | -a} {yes | no}

Checks the results without actually deleting the compatibility.

{--verbose | -v}

Displays more detailed information.

{--force | -f}

Performs the action without asking for confirmation.

Example

The following command provides information about the results of deleting a compatibility without actually performing the action:

```
isi storagepool compatibilities active delete 1 --assess yes
```

Provided that a compatibility with the ID of 1 exists, OneFS displays information similar to the following example:

```
Deleting compatibility with id 1 is possible.  
This delete will cause these nodepools to split:  
1: Nodepool s200_0b_0b will be split. A tier will be created and all  
resultant nodepools from this split will be incorporated into it. All  
filepool policies targeted at the splitting pool will be redirected  
towards this new tier. That tier's name is s200_0b_0b-tier
```

isi storagepool compatibilities active list

Lists node compatibilities that have been created.

Syntax

```
isi storagepool compatibilities active list
[--limit <integer>]
[--format {table | json |
csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of compatibilities that are listed.

{--format | -f}

Lists active compatibilities in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

The following command lists active compatibilities:

```
isi storagepool compatibilities active list
```

Command output appears similar to the following example:

```
ID   Class 1   Class 2
-----
1    S200     S210
2    X400     X410
-----
Total: 2
```

isi storagepool compatibilities active view

Displays the details of an active node compatibility.

Syntax

```
isi storagepool compatibilities active view <ID>
```

Options

<ID>

The ID number of the compatibility to view. You can use the `isi storagepool compatibilities active list` command to display the ID numbers of active compatibilities.

Example

The following command displays information about an active compatibility with ID number 1:

```
isi storagepool compatibilities active view 1
```

Output from the command would be similar to the following:

```
ID: 1
Class 1: S200
Class 2: S210
```

isi storagepool compatibilities available list

Lists compatibilities that are available, but not yet created.

Syntax

```
isi storagepool compatibilities available list <name>
  [--limit <integer>]
  [--format {table | json |
  csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of available compatibilities that are listed.

{--format | -f}

Lists available compatibilities in the specified format. The following values are valid:

```
table
json
csv
list
```

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Example

The following command lists available compatibilities:


```
isi storagepool compatibilities available list
```

If available compatibilities exist, command output appears similar to the following example:

```
Class 1   Class 2
-----
S200     S210
X400     X410
-----
Total: 2
```

isi storagepool compatibilities class active create

Creates a compatibility to enable an unprovisioned node to join a node pool.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities class active create <class-1> <class-2>
  [--assess {yes | no}]
  [--verbose]
  [--force]
```

Options

<class-1>

An existing node pool class, one of S200, X200, X400, or N400.

<class-2>

The node class that is compatible with the existing node pool, one of S210, X210, X410, or N410. Note that S210, X210, X410, and NL410 nodes are compatible only with similarly configured S200, X200, X400, and NL400 node pools, respectively. Also note that, in CLI commands, NL400 and NL410 nodes are expressed as N400 and N410.

--assess {yes | no}

Checks whether the compatibility is valid without actually creating the compatibility.

--verbose

Displays more detailed information.

--force

Performs the action without asking for confirmation.


Examples

The following command creates a compatibility between S200 and S210 nodes without asking for confirmation:

```
isi storagepool compatibilities class active create S200 S210 --force
```

isi storagepool compatibilities class active delete

Deletes a node class compatibility. If fewer than three compatible nodes were added to an existing node pool, they are removed from the node pool and become unprovisioned.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities class active delete <ID>
  [--assess {yes | no}]
  [--verbose]
  [--force]
```

Options

<ID>

The ID number of the compatibility. You can use the `isi storagepool compatibilities class active list` command to view the ID numbers of active compatibilities.

--assess {yes | no}

Checks the results without actually deleting the compatibility.

--verbose

Displays more detailed information.

--force

Performs the action without asking for confirmation.

Example

The following command provides information about the results of deleting a compatibility without actually performing the action:


```
isi storagepool compatibilities class active delete 1 --assess yes
```

Provided that a compatibility with the ID of 1 exists, OneFS displays information similar to the following example:

```
Deleting compatibility with id 1 is possible.
This delete will cause these nodepools to split:
1: Nodepool s200_0b_0b will be split. A tier will be created and all
resultant nodepools from this split will be incorporated into it. All
filepool policies targeted at the splitting pool will be redirected
towards this new tier. That tier's name is s200_0b_0b-tier
```

isi storagepool compatibilities class active list

Lists node class compatibilities that have been created.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities class active list
  [--limit <integer>]
  [--format {table | json |
  csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--limit<integer>

Limits the number of compatibilities that are listed.

--format

Lists active compatibilities in the specified format. The following values are valid:

- table
- json
- csv
- list

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following command lists active node class compatibilities:

```
isi storagepool compatibilities class active list
```


Command output appears similar to the following example:

```
ID   Class 1   Class 2
-----
1    S200     S210
2    X200     X210
3    X400     X410
4    N400     N410
-----
Total: 4
```

 **NOTE:** In CLI commands and output, NL400 and NL410 nodes are expressed as N400 and N410, respectively.

isi storagepool compatibilities class active view

Displays the details of an active node class compatibility.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities class active view <ID>
```

Options

<ID>

The ID number of the compatibility to view. You can use the `isi storagepool compatibilities class active list` command to display the ID numbers of active node class compatibilities.

Example

The following command displays information about an active compatibility with ID number 1:


```
isi storagepool compatibilities class active view 1
```

Output from the command will be similar to the following:

```
ID: 1
Class 1: S200
Class 2: S210
```

isi storagepool compatibilities class available list

Lists node class compatibilities that are available, but not yet created.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities class available list <name>
  [--limit <integer>]
  [--format {table | json |
  csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--limit<integer>

Limits the number of available compatibilities that are listed.

--format

Lists available compatibilities in the specified format. The following values are valid:

- table
- json
- csv
- list

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following command lists available compatibilities:

```
isi storagepool compatibilities class available list
```

If compatibilities are available, command output similar to the following example appears:

```
Class 1   Class 2
-----
S200     S210
X400     X410
```

Total: 2

isi storagepool compatibilities ssd active create

Creates an SSD compatibility, which can help to provision nodes with different SSD capacities to an existing compatible node pool. Without an SSD compatibility, compatible nodes having different SSD capacities cannot join the same node pool. If you have fewer than three nodes with a different SSD capacity, the nodes would remain unprovisioned, and therefore not functional.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities ssd active create <class-1>
  [--class-2 <string>]
  [--count {yes | no}]
  [--assess {yes | no}]
  [--verbose]
  [--force]
```

Options

<class-1>

The node class that the SSD compatibility will be created for. For example, you can create an SSD compatibility for S200 nodes that have larger-capacity SSDs than the nodes in an existing S200 node pool. In this way, OneFS can autoprovision the newer S200 nodes to the existing S200 node pool. You can use the `isi storagepool compatibilities ssd available list` command to display valid node class values. For example, S200.

--class-2 <string>

The second node class that will be made SSD-compatible with the first node class. For example, you can create an SSD compatibility for S210 nodes that have larger-capacity SSDs than the nodes in an existing S200 node pool. Because S210 nodes can be made compatible with S200 nodes, they can be autoprovisioned to an S200 node pool. However, in this case, you also need to create a node class compatibility between S200 and S210 nodes.

--count {yes | no}

Specifies whether to create an SSD count compatibility.

--assess {yes | no}

Checks whether the SSD compatibility is valid without actually creating the compatibility.

--verbose

Displays more detailed information.

--force

Performs the action without asking for confirmation.

Examples

The following command creates an SSD class compatibility and SSD count compatibility between S200 and S210 nodes:

```
isi storagepool compatibilities ssd active create S200 --class-2 S210 --count yes
```

OneFS displays an advisory message similar to the following, and requires you to confirm the operation:

```
You are attempting to create an SSD compatibility for node class 1. You are also attempting
to create an SSD compatibility for node class 2. Creating an SSD compatibility will merge
all automatic node pools with nodes from the compatibility's node class with the same ssd
count and hdd configuration and compatible RAM into a single node pool. This will require
all of these automatic node pools to have the same L3 setting, requested protection, and
tier membership. Any file pool policies currently targeting any of the merging node pools
will automatically be re-targeted towards the resultant merged pool. If there exists
enough unprovisioned nodes belonging to this compatibility's node class to form a node
```


pool, that node pool will be formed. This may potentially be very costly from a performance standpoint the next time the smartpools job runs. If this is a concern, please contact EMC Isilon Technical Support for more information.

Continue with creation? (yes/[no]):

Type **yes**, then press ENTER to continue. Type **no**, then press ENTER to cancel the process.

isi storagepool compatibilities ssd active delete

Deletes an SSD compatibility. If fewer than three nodes of a particular class were added to a node pool when the SSD compatibility was created, these nodes are removed from the node pool and become unprovisioned.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities ssd active delete <ID>
  [--id-2 <integer>]
  [--assess {yes | no}]
  [--verbose]
  [--force]
```

Options

<ID>

The ID number of the compatibility. You can use the `isi storagepool compatibilities ssd active list` command to view the ID numbers of active SSD compatibilities.

--id-2 <integer>

The ID number of the second SSD compatibility to delete. You can use the `isi storagepool compatibilities ssd active list` command to view the ID numbers of active SSD compatibilities. The `--id-2` setting is optional, unless the node pool with the SSD compatibility also has an associated node class compatibility. In this case, the setting is required, and deleting the second SSD compatibility will unprovision some of the nodes from the node pool.

--assess {yes | no}

Checks the results without actually deleting the SSD compatibility.

--verbose

Displays more detailed information.

--force

Performs the action without asking for confirmation.

Example

The following command provides information about the results of deleting an SSD compatibility without actually performing the action:


```
isi storagepool compatibilities ssd active delete 1 --id-2 2 --assess yes
```

Provided that an SSD compatibility between ID 1 and ID 2 exists, OneFS displays information similar to the following example:

```
Deleting ssd compatibility with id 1 is possible.
Deleting ssd compatibility with id 2 is possible.
This delete will cause these nodepools to split:
1: Nodepool s200_9.8kb_9.8kb-ssd_0b will be split. A tier will be created and all
resultant nodepools from this split will be incorporated into it. All
filepool policies targeted at the splitting pool will be redirected
towards this new tier. That tier's name is s200_9.8kb_9.8kb-ssd_0b-tier
```

isi storagepool compatibilities ssd active list

Lists SSD compatibilities that have been created.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities ssd active list
[--limit <integer>]
[--format {table | json |
csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--limit<integer>

Limits the number of SSD compatibilities that are listed.

--format

Lists active SSD compatibilities in the specified format. The following values are valid:

- table
- json
- csv
- list

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following command lists active SSD compatibilities:

```
isi storagepool compatibilities ssd active list
```


Command output appears similar to the following example:

```
ID    Class
-----
1     S200
2     S210
3     N400
4     N410
-----
Total: 4
```

 **NOTE:** In CLI commands and output, NL400 and NL410 nodes are expressed as N400 and N410, respectively.

isi storagepool compatibilities ssd active view

Displays the details of an active SSD compatibility.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities ssd active view <ID>
```

Options

<ID>

The ID number of the SSD compatibility to view. You can use the `isi storagepool compatibilities ssd active list` command to display the ID numbers of active SSD compatibilities.

Example

The following command displays information about an active compatibility with ID number 1:


```
isi storagepool compatibilities ssd active view 1
```

Output from the command will be similar to the following:

```
ID: 1  
Class: S200
```

isi storagepool compatibilities ssd available list

Lists SSD compatibilities that are available, but not yet created.

 **NOTE:** This command is not applicable for IsilonSD Edge.

Syntax

```
isi storagepool compatibilities ssd available list  
  [--limit <integer>]  
  [--format {table | json |  
  csv | list}]  
  [--no-header]  
  [--no-footer]  
  [--verbose]
```

Options

--limit<integer>

Limits the number of SSD compatibilities that are listed.

--format

Lists active SSD compatibilities in the specified format. The following values are valid:

- table
- json
- csv
- list

--no-header

Displays table and CSV output without headers.

--no-footer

Displays table output without footers.

--verbose

Displays more detailed information.

Example

The following command lists available SSD compatibilities:

```
isi storagepool compatibilities ssd available list
```

If available SSD compatibilities exist, command output similar to the following example appears:

```
Class 1
-----
S200
S210
-----
Total: 2
```

isi storagepool health

Displays the health information of storage pools.

Syntax

```
isi storagepool health
```

Options

{--verbose | -v}

Displays more detailed information.

isi storagepool list

Displays node pools and tiers in the cluster.

Syntax

```
isi storagepool list
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--format

Displays node pools and tiers in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi storagepool nodepools create

Creates a manually managed node pool. This command should only be used by experienced OneFS administrators or the with assistance of technical support personnel.

Syntax

```
isi storagepool nodepools create <name>
  [--lnns <lnns>]
  [--verbose]
```

Options

<name>

Specifies the name for the node pool. Names must begin with a letter or an underscore and may contain only letters, numbers, hyphens, underscores, or periods.

{--lnns <lnns> | -n <lnns>}

Specifies the nodes in this pool. Nodes can be a comma-separated list or range of LNNs—for example, **1,4,10,12,14,15** or **1-6**.

{--verbose | -v}

Displays more detailed information.

isi storagepool nodepools delete

Deletes a node pool and autoprovisions the affected nodes into the appropriate node pool. This command is used only for manually managed node pools and should be executed by experienced OneFS administrators or with direction from technical support personnel.

Syntax

```
isi storagepool nodepools delete <name>
  [--force]
  [--verbose]
```

Options

<name>

Specifies the name of the node pool to be deleted.

{--force | -f}

Suppresses any prompts, warnings, or confirmation messages that would otherwise appear.

{--verbose | -v}

Displays more detailed information.

isi storagepool nodepools list

Displays a list of node pools.

Syntax

```
isi storagepool nodepools list
  [--limit <integer>
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Specifies the number of node pools to display.

--format

Displays tiers in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi storagepool nodepools modify

Modifies a node pool.

Syntax

```
isi storagepool nodepools modify <name>
  [--protection-policy <string>]
  [--lnns <integer>]
  [--clear-lnns]
  [--add-lnns <integer>]
  [--remove-lnns <integer>]
  [--tier <string>]
  [--clear-tier]
  [--l3 {yes | no}]
  [--set-name <string>]
```

Options

<string>

Name of the node pool to be modified.

--protection-policy <string>

Requested protection for the node pool. Possible protection policy values are:

- +1n
- +2d:1n
- +2n
- +3d:1n
- +3d:1n1d
- +3n
- +4d:1n
- +4d:2n
- +4n
- Mirror values: 2x, 3x, 4x, 5x, 6x, 7x, 8x

OneFS calculates the optimal protection policy (referred to as suggested protection). If the value you set is lower than the suggested protection, OneFS displays an alert.

--lnns <integer>

Nodes for the manually managed node pool. Specify `--lnns` for each additional node for the manually managed node pool.

--clear-lnns

Clear value for nodes for the manually managed node pool.

--add-lnns <integer>

Add nodes for the manually managed node pool. Specify `--add-lnns` for each additional node to add.

--remove-lnns <integer>

Remove nodes for the manually managed node pool. Specify `--remove-lnns` for each additional node to remove.

--tier <string>

Set parent for the node pool. Node pools can be grouped into a tier to service particular file pools.

--clear-tier

Remove the specified node pool from its parent tier.

--l3 {yes | no}

Use SSDs in the specified node pool as L3 cache. Note that, on Isilon HD400 node pools, L3 cache is on by default and you cannot disable it. If you try to disable L3 cache on an HD400 node pool, OneFS generates the following error message: `Disabling L3 not supported for the given node type.`

--set-name <string>

New name for the manually managed node pool.

Examples

The following command specifies that SSDs in a node pool named `hq_datastore` are to be used as L3 cache:

```
isi storagepool nodepools modify hq_datastore --l3 yes
```

The following command adds the node pool `hq_datastore` to an existing tier named `archive-1`:

```
isi storagepool nodepools modify hq_datastore --tier archive-1
```

isi storagepool nodepools view

Displays details for a node pool.

Syntax

```
isi storagepool nodepools view <name>
[--verbose]
```

Options

<name>

Specifies the name of the storage pool.

{--verbose | -v}

Displays more detailed information.

isi storagepool settings modify

Modifies global SmartPools settings.

Syntax

```
isi storagepool settings modify
[--automatically-manage-protection {all | files_at_default | none}]
[--automatically-manage-io-optimization {all | files_at_default | none}]
[--protect-directories-one-level-higher {yes | no}]
[--global-namespace-acceleration-enabled {yes | no}]
[--virtual-hot-spare-deny-writes {yes | no}]
[--virtual-hot-spare-hide-spare {yes | no}]
[--virtual-hot-spare-limit-drives <integer>]
[--virtual-hot-spare-limit-percent <integer>]
[--snapshot-disk-pool-policy-id <integer>]
[--spillover-target <string>| --no-spillover | --spillover-anywhere]
[--ssd-l3-cache-default-enabled {yes | no}]
[--ssd-qab-mirrors {one | all}]
[--ssd-system-btree-mirrors {one | all}]
[--ssd-system-delta-mirrors {one | all}]
[--verbose]
```

Required Privileges

ISI_PRIV_SMARTPOOLS

Options

--automatically-manage-protection {all | files_at_default | none}

Specifies whether SmartPools manages files' protection settings.

--automatically-manage-io-optimization {all | files_at_default | none}

Specifies whether SmartPools manages I/O optimization settings for files.

--protect-directories-one-level-higher {yes | no}

Protects directories at one level higher.

--global-namespace-acceleration-enabled {yes | no}

Enables or disables global namespace acceleration.

--virtual-hot-spare-deny-writes {yes | no}

Denies new data writes to the virtual hot spare.

--virtual-hot-spare-hide-spare {yes | no}

Reduces the amount of available space for the virtual hot spare.

--virtual-hot-spare-limit-drives <integer>

Specifies the maximum number of virtual drives.

--virtual-hot-spare-limit-percent <integer>

Limits the percentage of node resources that is allocated to virtual hot spare.

--spillover-target <string>

Specifies the target for spillover.

--no-spillover

Globally disables spillover.

--spillover-anywhere

Globally sets spillover to anywhere.

--ssd-l3-cache-default-enabled {yes | no}

Enables or disables SSDs on new node pools to serve as L3 cache.

--ssd-qab-mirrors {one | all}

Specifies that either one QAB (quota accounting block) mirror, or all QAB mirrors, be stored on SSDs. The default is for one mirror to be stored on SSDs. By specifying `all`, system access to the QAB is likely to be faster.

--ssd-system-btree-mirrors {one | all}

Specifies that either one system B-tree mirror, or all system B-tree mirrors, be stored on SSDs. The default is for one mirror to be stored on SSDs. By specifying `all`, system access to the B-tree is likely to be faster.

--ssd-system-delta-mirrors {one | all}

Specifies that either one system delta mirror, or all system delta mirrors, be stored on SSDs. The default is for one mirror to be stored on SSDs. By specifying `all`, access to the system delta is likely to be faster.

--verbose

Enables verbose messaging.

Examples

The following command specifies that SSDs on newly created node pools are to be used as L3 cache:

```
isi storagepool settings modify --ssd-l3-cache-default-enabled yes
```

The following command specifies that 20 percent of node resources can be used for the virtual hot spare:

```
isi storagepool settings modify --virtual-hot-spare-limit-percent 20
```

isi storagepool settings modify

Modify global SmartPools settings.

Syntax

```
isi storagepool settings modify
  [--automatically-manage-protection {all | files_at_default | none}]
  [--automatically-manage-io-optimization {all | files_at_default | none}]
  [--protect-directories-one-level-higher <boolean>]
  [--global-namespace-acceleration-enabled <boolean>]
  [--virtual-hot-spare-deny-writes <boolean>]
```

```
[--virtual-hot-spare-hide-spare <boolean>]
[--virtual-hot-spare-limit-drives <integer>]
[--virtual-hot-spare-limit-percent <integer>]
[--snapshot-disk-pool-policy-id <integer>]
[--spillover-target <string>| --no-spillover | --spillover-anywhere]
[--ssd-l3-cache-default-enabled <boolean>]
[{-verbose | -v}]
[{-help | -h}]
```

Required Privileges

ISI_PRIV_SMARTPOOLS

Options

- automatically-manage-protection {all | files_at_default | none}**
Set whether SmartPools manages files' protection settings.
- automatically-manage-io-optimization {all | files_at_default | none}**
Set whether SmartPools manages files' I/O optimization settings
- protect-directories-one-level-higher <boolean>**
Protect directories at one level higher.
- global-namespace-acceleration-enabled <boolean>**
Enable or disable global namespace acceleration.
- virtual-hot-spare-deny-writes <boolean>**
Deny new data writes.
- virtual-hot-spare-hide-spare <boolean>**
Reduce the amount of available space.
- virtual-hot-spare-limit-drives <integer>**
Specify the maximum number of virtual drives.
- virtual-hot-spare-limit-percent <integer>**
Limit the percent of node resources allocated to virtual hot spare.
- spillover-target <string>**
Specifies the target for spillover.
- no-spillover**
Globally disables spillover.
- spillover-anywhere**
Globally sets spillover to anywhere.
- ssd-l3-cache-default-enabled**
Enable or disable SSDs on new node pools to serve as L3 cache.

Examples

The following command specifies that SSDs on newly created node pools are to be used as L3 cache:

```
isi storagepool settings modify --ssd-l3-cache-default on
```

The following command specifies that 20 percent of node resources can be used for virtual hot spare purposes:

```
isi storagepool settings modify --virtual-hot-spare-limit-percent 20
```

isi storagepool settings view

Displays global SmartPools settings.

Syntax

```
isi storagepool settings view
```

Options

There are no options for this command.

Example

The following command displays the global SmartPools settings on your cluster:

```
isi storagepool settings view
```

The system displays output similar to the following example:

```
Automatically Manage Protection: files_at_default
Automatically Manage Io Optimization: files_at_default
Protect Directories One Level Higher: Yes
  Global Namespace Acceleration: disabled
  Virtual Hot Spare Deny Writes: Yes
  Virtual Hot Spare Hide Spare: Yes
  Virtual Hot Spare Limit Drives: 1
  Virtual Hot Spare Limit Percent: 0
  Global Spillover Target: anywhere
  Spillover Enabled: Yes
  SSD L3 Cache Default Enabled: Yes
  SSD Qab Mirrors: one
  SSD System Btree Mirrors: one
  SSD System Delta Mirrors: one
```

isi storagepool tiers create

Creates a tier.

Syntax

```
isi storagepool tiers create <name>
  [--children <string>]
  [--verbose]
```

Options

<name>


Specifies the name for the storage pool tier. Specify as any string.

--children <string>

Specifies a node pool to be added to the tier. For each node pool that you intend to add, include a separate --children argument.

--verbose

Displays more detailed information.

 **NOTE:** Names must begin with a letter or underscore and must contain only letters, numbers, hyphens, underscores, or periods.

Example

The following command creates a tier and adds two node pools to the tier:

```
isi storagepool tiers create ARCHIVE_1 --children hq_datastore1
--children hq_datastore2
```

isi storagepool tiers delete

Deletes a tier.

Syntax

```
isi storagepool tiers delete {<name> | --all}
[--verbose]
```

Options

{<name> | --all}

Specifies the tier to delete. The acceptable values are the name of the tier or all.

{--verbose | -v}

Displays more detailed information.

isi storagepool tiers list

Displays a list of tiers.

Syntax

```
isi storagepool tiers list
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--format

Displays tiers in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

`{--no-footer | -z}`

Displays table output without footers.

`{--verbose | -v}`

Displays more detailed information.

isi storagepool tiers modify

Renames a tier.

Syntax

```
isi storagepool tiers modify <name>
  [--set-name <string>]
  [--verbose]
```

Options

`<name>`

Specifies the tier to be renamed.

`{--set-name | -s} <string>`

Sets the new name for the tier.

`{--verbose | -v}`

Displays more detailed information.

i **NOTE:** Names must begin with a letter or underscore and must contain only letters, numbers, hyphens, underscores, or periods.

isi storagepool tiers view

Displays details for a tier.

Syntax

```
isi storagepool tiers view <name>
```

Options

`<name>`

Specifies the name of the tier.

`{--verbose | -v}`

Displays more detailed information.

isi storagepool unprovisioned view

Displays unprovisioned nodes and drives in an Isilon cluster.

Syntax

```
isi storagepool unprovisioned view
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Limits the number of unprovisioned nodes and drives to display.

--format

Displays the list of unprovisioned nodes and drives in the specified format. The following values are valid:

table

json

csv

list

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi swift accounts create

Create a new Swift account.

Syntax

```
isi swift accounts create <name> <swiftuser> <swiftgroup>
  [--zone <string>]
  [--users <string>]
  [--verbose]
```

Options

<name>

Specifies the name of the Swift account.

<swiftuser>

Specifies the file system user who owns files in the Swift account.

<swiftgroup>

Specifies the file system group who owns files in the Swift account. *<swiftgroup>* is associated with the directory that represents an account whereas the containers and objects are associated with the primary group of the user. If an administrator wants to assign *<swiftgroup>* to containers and objects as well, they must specify *<swiftgroup>* as the primary group of the Swift user.

--zone <string>

Specifies the access zone that is associated with the Swift account.

--users <string>

Specifies the users who are assigned access to the Swift account. Specify **--users** for each additional user who must be assigned access to the Swift account.

{--verbose | -v}

Displays detailed information.

isi swift accounts delete

Deletes a Swift account.

Syntax

```
isi swift accounts delete <name>
  [--zone <string>]
  [--verbose]
  [--force]
```

Options

<name>

Specifies the Swift account name.

--zone <string>

Specifies the access zone that is associated with the Swift account.

{--verbose | -v}

Displays detailed information.

{--force | -f}

Does not ask for a confirmation before deletion.

isi swift accounts list

Lists all of the Swift accounts.

Syntax

```
isi swift accounts list
  [--zone <string>]
  [--limit <integer>]
  [--sort (name | zone | swiftuser | swiftgroup | users)]
  [--descending]
  [--format (table | json | csv | list)]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--zone <string>

Specifies the access zone that is associated with the Swift account.

{--limit | -l} <integer>

Specifies the number of accounts to display.

--sort (name | zone | swiftuser | swiftgroup | users)

Sorts data using one of account name, access zone, Swift user, Swift group, or the users who are assigned access to the Swift account.

{--descending | -d}

Sorts data in descending order.

--format (table | json | csv | list)

Displays accounts in table, JSON, CSV, or list format.

{--no-header | -a}

Does not display headers in CSV or table formats.

{--no-footer | -z}

Does not display table summary footer information.

{--verbose | -v}

Displays detailed information.

isi swift accounts modify

Modifies a Swift account.

Syntax

```
isi swift accounts modify <name>
  [--zone <string>]
  [--swiftuser <string>]
  [--swiftgroup <string>]
  [--users <string> | --clear-users | --add-users <string> | --remove-users <string>]
  [--verbose]
```

Options

<name>

Specifies the name of the Swift account.

--zone <string>

Specifies the access zone that is associated with the Swift account.

--swiftuser <string>

Specifies the file system user who owns files in the Swift account.

--swiftgroup <string>

Specifies the file system group that owns files in the Swift account. *<swiftgroup>* is associated with the directory that represents an account whereas the containers and objects are associated with the primary group of the user. If an administrator wants to assign *<swiftgroup>* to containers and objects as well, they must specify *<swiftgroup>* as the primary group of the Swift user.

--users <string>

Specifies the users who are assigned access to the Swift account. Specify *--users* for each additional user who must be assigned access to the Swift account.

--clear-users

Clears the values specified for the users who are assigned access to the Swift account.

--add-users <string>

Adds users and assigns them access to the Swift account. Specify *--add-users* for each additional user that you want to add.

--remove-users <string>

Removes users that are assigned access to the Swift account. Specify `--remove-users` for each additional user that you want to remove.

{--verbose | -v}

Displays detailed information.

isi swift accounts view

Displays the details associated with a Swift account.

Syntax

```
isi swift accounts view <name>
[--zone <string>]
```

Options

<name>

Specifies the name of the Swift account.

--zone <string>

Specifies the access zone that is associated with the Swift account.

isi sync certificates peer delete

Delete a trusted SynclQ TLS certificate.

Syntax

```
isi sync certificates peer delete <id>
[--force]
[--verbose]
```

Options

<id>

System certificate identifier or certificate name.

{--force | -f}

Do not prompt for confirmation of delete.

{--verbose | -v}

Displays more detailed information.

isi sync certificates peer import

Import a trusted SynclQ TLS certificate.

Syntax

```
isi sync certificates peer import <certificate-path>
[--name <string>]
```

```
[--description <string>]
[--verbose]
```

Options

<certificate-path>

Local path to the TLS certificate file in PEM, DER, or PKCS#12 format. This certificate file is copied into the system certificate store, and you can remove it after import.

--name <string>

Administrator-configured certificate identifier.

--description <string>

Description field for administrative convenience.

{--verbose | -v}

Displays more detailed information.

isi sync certificates peer list

View a list of trusted SynclQ TLS certificates.

Syntax

```
isi sync certificates peer list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l}<integer>

The number of SynclQ certificate peers to list.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync certificates peer modify

Modify a trusted SynclQ TLS certificate.

Syntax

```
isi sync certificates peer modify <id>
  [--name <string>]
```

```
[--description <string>]
[--verbose]
```

Options

- <id>**
System certificate identifier or certificate name.
- name <string>**
Administrator-configured certificate identifier.
- description <string>**
Description field for administrative convenience.
- {--verbose | -v}**
Displays more detailed information.

isi sync certificates peer view

View a trusted SynclQ TLS certificate.

Syntax

```
isi sync certificates peer view <id>
  [--format (list | json)]
```

Options

- <id>**
System certificate identifier or certificate name.
- format (list | json)**
Displays output in list (default) or JavaScript Object Notation (JSON) format.
- {--no-header | -a}**
Displays table and CSV output without headers.
- {--no-footer | -z}**
Displays table output without footers.
- {--verbose | -v}**
Displays more detailed information.

isi sync certificates server delete

Delete a SynclQ TLS certificate.

Syntax

```
isi sync certificates server delete <id>
  [--force]
  [--verbose]
```

Options

- <id>**
System certificate identifier or certificate name.

{--force | -f}

Do not prompt for confirmation of delete.

{--verbose | -v}

Displays more detailed information.

isi sync certificates server import

Import a SynclQ TLS certificate.

Syntax

```
isi sync certificates server import <certificate-path> <certificate-key-path>
  [--name <string>]
  [--certificate-key-password <string>]
  [--verbose]
```

Options

<certificate-path>

Local path to the TLS certificate file in PEM, DER, or PKCS#12 format. This certificate file is copied into the system certificate store, and you can remove it after import.

<certificate-key-path>

Local path to the TLS certificate key file in PEM, DER, or PKCS#12 format. This certificate key file is copied into the system certificate store, and you should remove it after import.

--name <string>

Administrator-configured certificate identifier.

--certificate-key-password <string>

The password for the certificate key, if the private key is password encrypted.

{--verbose | -v}

Displays more detailed information.

isi sync certificates server list

View a list of SynclQ TLS certificates.

Syntax

```
isi sync certificates server list
  [--limit <integer>]
  [--format (table | json | csv | list)]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l}<integer>

The number of SynclQ certificate servers to list.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync certificates server modify

Modify a SynclQ TLS certificate.

Syntax

```
isi sync certificates server modify <id>
  [--name <string>]
  [--description <string>]
  [--verbose]
```

Options

<id>

System certificate identifier or certificate name.

--name <string>

Administrator-configured certificate identifier.

--description <string>

Description field for administrative convenience.

{--verbose | -v}

Displays more detailed information.

isi sync certificates server view

View a SynclQ TLS certificate.

Syntax

```
isi sync certificates peer view <id>
  [--format (list | json)]
```

Options

<id>

System certificate identifier or certificate name.

--format (list | json)

Displays output in list (default) or JavaScript Object Notation (JSON) format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync jobs cancel

Cancels a running or paused replication job.

Syntax

```
isi sync jobs cancel {<policy-name> | --all}
[--verbose]
```

Options

<policy-name>

Cancels a job that was created according to the specified replication policy.

Specify as a replication policy name or ID.

--all

Cancels all currently running replication jobs.

--verbose

Displays more detailed information.

isi sync jobs list

Displays information about the most recently completed and next scheduled replication jobs of replication policies.

Syntax

```
isi sync jobs list
[--state <state>]
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

If no options are specified, displays information about replication jobs for all policies.

--state <state>

Displays only jobs in the specified state.

The following values are valid:

scheduled	Displays jobs that are scheduled to run.
running	Displays running jobs.
paused	Displays jobs that were paused by a user.
finished	Displays jobs that have completed successfully.
failed	Displays jobs that failed during the replication process.
canceled	Displays jobs that were cancelled by a user.
needs_attention	Displays jobs that require user intervention before they can continue.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync jobs pause

Pauses a running replication job.

Syntax

```
isi sync jobs pause {<policy-name> | --all}
[--verbose]
```

Options

<policy-name>

Pauses a job that was created according to the specified replication policy.

Specify as a replication policy name.

--all

Pauses all currently running replication jobs.

{--verbose | -v}

Displays more detailed information.

isi sync jobs reports list

Displays information about running replication jobs targeting the local cluster.

Syntax

```
isi sync jobs reports list
[--limit <integer>]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync jobs reports view

Displays information about a running replication job targeting the local cluster.

Syntax

```
isi sync jobs reports view <policy>
```

Options

<policy>

Displays information about a replication job created according to the specified replication policy.

Specify as a replication policy name or ID.

isi sync jobs resume

Resumes paused replication jobs.

Syntax

```
isi sync jobs resume {<policy-name> | --all}  
[--verbose]
```

Options

<policy-name>

Resumes a paused job that was created by the specified policy.

Specify as a replication policy name.

--all

Resumes all currently running replication jobs.

{--verbose | -v}

Displays more detailed information.

isi sync jobs start

Starts a replication job for a replication policy.

Syntax

```
isi sync jobs start <policy-name>  
[--test]  
[--source-snapshot <snapshot>]  
[--verbose]
```

Options

<policy-name>

Starts a replication job for the specified replication policy.

--test

Creates a replication policy report that reflects the number of files and directories that would be replicated if the specified policy was run. You can test only policies that have not been run before.

--source-snapshot <snapshot>

Replicates data according to the specified SnapshotIQ snapshot. If specified, a snapshot is not generated for the replication job. Replicating data according to snapshots generated by the SynclQ tool is not supported.

Specify as a snapshot name or ID. The source directory of the policy must be contained in the specified snapshot. This option is valid only if the last replication job completed successfully or if you are performing a full or differential replication. If the last replication job completed successfully, the specified snapshot must be more recent than the snapshot referenced by the last replication job.

{--verbose | -v}

Displays more detailed information.

isi sync jobs view

Displays information about a running replication job.

Syntax

```
isi sync jobs view <policy>
```

Options

<policy>

Displays information about a running replication job created according to the specified policy.

Specify as a replication policy name or ID.

isi sync policies create

Creates a replication policy.

Syntax

```
isi sync policies create <name> <action>
  <source-root-path> <target-host> <target-path>
  [--enabled (yes | no)]
  [--description <string>]
  [--check-integrity (yes | no)]
  [--source-include-directories <string>]
  [--source-exclude-directories <string>]
  [--source-subnet <subnet> | --source-pool <pool>]
  [--target-snapshot-archive (on | off)]
  [--target-snapshot-pattern <naming-pattern>]
  [--target-snapshot-expiration <duration>]
  [--target-snapshot-alias <naming-pattern>]
  [--target-detect-modifications (on | off)]
  [--source-snapshot-archive (on | off)]
  [--source-snapshot-pattern <naming-pattern>]
  [--source-snapshot-expiration <duration>]
  [--snapshot-sync-pattern <pattern>]
  [--snapshot-sync-existing (yes | no)]
  [--schedule (<schedule> | when-source-modified
  | when-snapshot-taken)]
  [--job-delay <duration>]
  [--skip-when-source-unmodified (true | false)]
  [--rpo-alert <duration>]
```

```

[--log-level <level>]
[--log-removed-files (yes | no)]
[--workers-per-node <integer>]
[--report-max-age <duration>]
[--report-max-count <integer>]
[--restrict-target-network (on | off)]
[--target-compare-initial-sync (on | off)]
[--accelerated-failback (yes | no)]
[--priority (0 | 1 | normal | high)]
[--cloud-deep-copy (deny | allow | force)]
[--bandwidth-reservation <integer>]
[--target-certificate-id <string>]
[--ocsp-issuer-certificate-id <string>]
[--ocsp-address <string>]
[--encryption-cipher-list <string>]
[--linked-service-policies <string>]
[--delete-quotas (yes | no)]
[--password (<password>)]
[--set-password]
[--verbose]

```

Options

<name>

Specifies a name for the replication policy.

Specify as any string.

<action>

Specifies the type of replication policy.

The following types of replication policy are valid:

copy	Creates a copy policy that adds copies of all files from the source to the target.
sync	Creates a synchronization policy that synchronizes data on the source cluster to the target cluster and deletes all files on the target cluster that are not present on the source cluster.

<source-root-path>

Specifies the directory on the local cluster that files are replicated from.

Specify as a full directory path.

<target-host>

Specifies the cluster that the policy replicates data to.

Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**

This will replicate data to another directory on the local cluster.

NOTE: SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

<target-path>

Specifies the directory on the target cluster that files are replicated to.

Specify as a full directory path.

--enabled (yes | no)

Determines whether the policy is enabled or disabled.

The default value is *yes*.

--description <string>

Specifies a description of the replication policy.

--check-integrity {yes | no}

Specifies whether to perform a checksum on each file data packet that is affected by the SyncIQ job. If this option is set to `yes`, and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is `yes`.

{--source-include-directories | -i} <path>

Includes only the specified directories in replication.

Specify as any directory path contained in the root directory. You can specify multiple directories by specifying `--source-include-directories` multiple times within a command. For example, if the root directory is `/ifs/data`, you could specify the following:

```
--source-include-directories /ifs/data/music --source-include-directories /ifs/data/movies
```

{--source-exclude-directories | -e} <path>

Does not include the specified directories in replication. Specify as any directory path contained in the root directory. If `--source-include-directories` is specified, `--source-exclude-directories` directories must be contained in the included directories. You can specify multiple directories by specifying `--source-exclude-directories` multiple times within a command. For example, you could specify the following:

```
--source-exclude-directories /ifs/data/music --source-exclude-directories /ifs/data/movies \  
--exclude /ifs/data/music/working
```

--source-subnet <subnet>

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster. If you specify this option, you must also specify `--source-pool`.

--source-pool <pool>

Restricts replication jobs to running only on nodes in the specified pool on the local cluster. If you specify this option, you must also specify `--source-subnet`.

--target-snapshot-archive {on | off}

Determines whether archival snapshots are generated on the target cluster. If this option is set to `off`, SyncIQ will still maintain exactly one snapshot at a time on the target cluster to facilitate failback. You must activate a SnapshotIQ license on the target cluster to generate archival snapshots on the target cluster.

--target-snapshot-pattern <naming-pattern>

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.

The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

--target-snapshot-expiration <duration>

Specifies an expiration period for archival snapshots on the target cluster.

If this option is not specified, archival snapshots will remain indefinitely on the target cluster.

Specify in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days

H Specifies hours

--target-snapshot-alias < naming-pattern >


Specifies a naming pattern for the most recent archival snapshot generated on the target cluster.

The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

--target-detect-modifications {on | off}

Determines whether SyncIQ checks the target directory for modifications before replicating files.

 **CAUTION: Specifying `off` could result in data loss. It is recommended that you consult Isilon Technical Support before specifying `off`.**

--source-snapshot-archive {on | off}

Determines whether archival snapshots are retained on the source cluster. If this option is set to `off`, SyncIQ will still maintain one snapshot at a time for the policy to facilitate replication.

--source-snapshot-pattern < naming-pattern >

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster.

For example, the following pattern is valid:

```
SIQ-source-#{PolicyName}-%Y-%m-%d_%H-%M
```

--source-snapshot-expiration < duration >

Specifies an expiration period for archival snapshots retained on the source cluster.

If this option is not specified, archival snapshots will exist indefinitely on the source cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--snapshot-sync-pattern < string >

The naming pattern that a snapshot must match to trigger a replication job, when the schedule is set to **when-snapshot-taken**. The default value is asterisk (*).

--snapshot-sync-existing {yes | no}

If set to Yes, snapshot-triggered replication jobs will include replications taken before the policy creation time. The default is No. If set to yes, set `--schedule when-snapshot-taken`.

{--schedule | -S} {< schedule > | when-source-modified | when-snapshot-taken}

Specifies how often data will be replicated. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified. Specifying `when-snapshot-taken` causes OneFS to replicate data every time that a snapshot is taken of the source directory.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify *<interval>* in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]

- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

--job-delay <duration>

Specifies the amount of time after the source directory is modified that SyncIQ waits before starting a replication job. If the `--schedule` of this replication policy is set to `when-source-modified`, and the contents of the source directory are modified, SyncIQ will wait the specified amount of time before starting a replication job.

The default value is 0 seconds.

--skip-when-source-unmodified {true | false}

Causes the policy not to be run if the contents of the source directory have not been modified since the last time the policy has been run. If `--schedule` of this replication policy is set to <schedule>, and the policy is scheduled to run before changes have been made to the contents of the source directory, the policy will not be run.

--rpo-alert <duration>

Creates a OneFS event if the specified Recovery Point Objective (RPO) is exceeded. For example, assume you set an RPO of 5 hours; a job starts at 1:00 PM and completes at 3:00 PM; a second job starts at 3:30 PM; if the second job does not complete by 6:00 PM, SyncIQ will create a OneFS event.

The default value is 0, which will not generate events. This option is valid only if `--schedule` is set to <schedule>.



NOTE: This option is valid only if RPO alerts have been globally enabled through SyncIQ settings. The events have an event ID of 400040020.

--log-level <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug
- trace

The default value is info.

--log-removed-files {yes | no}

Determines whether SyncIQ retains a log of all files that are deleted when a synchronization policy is run. This parameter has no effect for copy policies.

The default value is `no`.

{--workers-per-node | -w} <integer>

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is 3.

 **NOTE: This option has been deprecated and will not be recognized if configured.**

--report-max-age <duration>

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--report-max-count <integer>

Specifies the maximum number of reports to retain for the replication policy.

--restrict-target-network (on | off)

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

--target-compare-initial-sync (on | off)

Determines whether the full or differential replications are performed for this policy. Full or differential replications are performed the first time a policy is run and after a policy has been reset. If set to `on`, performs a differential replication. If set to `off`, performs a full replication.

If differential replication is enabled the first time a replication policy is run, the policy will run slower without any benefit.

The default value is `off`.

--accelerated-failback (enable | disable)

If enabled, SyncIQ will perform failback configuration tasks the next time that a job is run, rather than waiting to perform those tasks during the failback process. Performing these tasks ahead of time will increase the speed of failback operations.

--priority (0 | 1)

Determines whether the policy has priority.

The default value is 0, which means that the policy does not have priority. If set to 1, the policy is high-priority.

--cloud-deep-copy (deny | allow | force)

Determines how the policy replicates CloudPools smartlinks. If set to `deny`, SyncIQ replicates all CloudPools smartlinks to the target cluster as smartlinks; if the target cluster does not support the smartlinks, the job will fail. If set to `force`, SyncIQ replicates all smartlinks to the target cluster as regular files. If set to `allow`, SyncIQ will attempt to replicate smartlinks to the target cluster as smartlinks; if the target cluster does not support the smartlinks, SyncIQ will replicate the smartlinks as regular files.

--bandwidth-reservation <integer>

The desired bandwidth reservation for this policy, in kb/s. This feature does not activate unless a SynclQ bandwidth rule is in effect.

--target-certificate-id <string>

The identifier of the target cluster certificate being used for encryption.

--ocsp-issuer-certificate-id <string>

The identifier of the certificate authority that issued the certificate whose revocation status is being checked.

--ocsp-address <string>

The address of the OCSP responder to which you want to connect.

--encryption-cipher-list <string>

The cipher list being used with cluster encryption. For SynclQ targets, this is a list of supported ciphers. For SynclQ sources, the list of ciphers is used in order.

--linked-service-policies <string>...

A list of SynclQ policy identifiers whose source root directories will be used to filter service replication. Specify this option again for each additional service policy identifier.

--delete-quotas (yes | no)

If set to Yes, forcibly removes quotas from the target when they are removed from the source.

--password <password>

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--set-password`.

--set-password

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--password`.

{**--verbose** | **-v**}

Displays a message confirming that the snapshot schedule was created.

isi sync policies delete

Deletes a replication policy.

The command will not succeed until SynclQ can communicate with the target cluster; until then, the policy will still appear in the output of the `isi sync policies list` command. After the connection between the source cluster and target cluster is reestablished, SynclQ will delete the policy the next time that the job is scheduled to run; if the policy is configured to run only manually, you must manually run the policy again. If SynclQ is permanently unable to communicate with the target cluster, specify the `--local-only` option. This will delete the policy from the local cluster only and not break the target association on the target cluster.

Syntax

```
isi sync policies delete {<policy> | --all}
  [--local-only]
  [--force]
  [--verbose]
```

Options

<policy>


Deletes the specified replication policy.

--all

Deletes all replication policies.

--local-only

Does not break the target association on the target cluster. Not deleting a policy association on the target cluster will cause the target directory to remain in a read-only state.

 **NOTE:** If SyncIQ is unable to communicate with the target cluster, you must specify this option to successfully delete the policy.

`{--force | -f}`

Deletes the policy, even if an associated job is currently running. Also, does not prompt you to confirm the deletion.

 **CAUTION:** Forcing a policy to delete might cause errors if an associated replication job is currently running.

`{--verbose | -v}`

Displays a confirmation message.

isi sync policies disable

Temporarily disables a replication policy. If a replication policy is disabled, the policy will not create replication jobs. However, if a replication job is currently running for a replication policy, disabling the policy will not pause or stop the job.

Syntax

```
isi sync policies disable {<policy> | --all}
[--verbose]
```

Options

`<policy>`

Disables the specified replication policy. Specify as a replication policy name or a replication policy ID.

`--all`

Disables all replication policies on the cluster.

`--verbose`

Displays more detailed information.

isi sync policies enable

Enables a disabled replication policy.

Syntax

```
isi sync policies enable {<policy> | --all}
[--verbose]
```

Options

`<policy>`

Enables the specified replication policy. Specify as a replication policy name or a replication policy ID.

`--all`

Enables all replication policies on the cluster.

`--verbose`

Displays more detailed information.

isi sync policies list

Displays a list of replication policies.

Syntax

```
isi sync policies list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, displays a table of all replication policies.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

name	Sorts output by the name of the replication policy.
target_path	Sorts output by the path of the target directory.
action	Sorts output by the type of replication policy.
description	Sorts output by the policy description.
enabled	Sorts output by whether the policies are enabled or disabled.
target_host	Sorts output by the target cluster.
check_integrity	Sorts output by whether the policy is configured to perform a checksum on each file data packet that is affected by a replication job.
source_root_path	Sorts output by the path of the source directory.
source_include_directories	Sorts output by directories that have been explicitly included in replication.
source_exclude_directories	Sorts output by directories that have been explicitly excluded in replication.
file_matching_pattern	Sorts output by the predicate that determines which files are replicated.
target_snapshot_archive	Sorts output by whether archival snapshots are generated on the target cluster.
target_snapshot_pattern	Sorts output by the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.
target_snapshot_expiration	Sorts output by the expiration period for archival snapshots on the target cluster.
target_detect_modifications	Sorts output by whether full or differential replications are performed for this policy.
source_snapshot_archive	Sorts output by whether archival snapshots are retained on the source cluster.

source_snapshot_pattern	Sorts output by the naming pattern for the most recent archival snapshot generated on the source cluster.
source_snapshot_expiration	Sorts output by the expiration period for archival snapshots retained on the source cluster.
schedule	Sorts output by the schedule of the policy.
log_level	Sorts output by the amount of data that is recorded in logs.
log_removed_files	Sorts output by whether OneFS retains a log of all files that are deleted when the replication policy is run.
workers_per_node	Sorts output by the number of workers per node that are generated by OneFS to perform each replication job for the policy.
report_max_age	Sorts output by how long replication reports are retained before they are automatically deleted by OneFS.
report_max_count	Sorts output by the maximum number of reports that are retained for the replication policy.
force_interface	Sorts output by whether data is sent over only the default interface of the subnet specified by the <code>--source-network</code> option of the <code>isi sync policies create</code> or <code>isi sync policies modify</code> commands.
restrict_target_network	Sorts output by whether replication jobs are restricted to connecting to nodes in a specified zone on the target cluster.
target_compare_initial_sync	Sorts output by whether full or differential replications are performed for the policies.
last_success	Sorts output by the last time that a replication job completed successfully.
password_set	Sorts output by whether the policy specifies a password for the target cluster.
source_network	Sorts output by the subnet on the local cluster that the replication policy is restricted to.
source_interface	Sorts output by the pool on the local cluster that the replication policy is restricted to.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync policies modify

Modifies existing replication policies.

Syntax

```
isi sync policies modify <policy>
  [--name <new-policy-name>]
  [--action <policy-type>]
  [--target-host <target-cluster>]
  [--target-path <target-path>]
  [--source-root-path <root-path>]
  [--description <string>]
```

```

[--password <password>]
[--set-password]
[--source-include-directories <string>]
[--clear-source-include-directories]
[--add-source-include-directories <string>]
[--remove-source-include-directories <string>]
[--source-exclude-directories <string>]
[--clear-source-exclude-directories]
[--add-source-exclude-directories <string>]
[--remove-source-exclude-directories <string>]
[--begin-filter <predicate> --operator <value> ... --end-filter]
[--schedule {<schedule> | when-source-modified}]
[--skip-when-source-unmodified {true | false}]
[--rpo-alert <duration>]
[--job-delay <duration>]
[--clear-job-delay]
[--snapshot-sync-pattern <pattern>]
[--snapshot-sync-existing {yes | no}]
[--enabled {true | false}]
[--check-integrity {true | false}]
[--log-level <level>]
[--log-removed-files {yes | no}]
[--workers-per-node <integer>]
[--target-snapshot-archive {on | off}]
[--target-snapshot-pattern <naming-pattern>]
[--target-snapshot-expiration <duration>]
[--target-snapshot-alias <naming-pattern>]
[--target-detect-modifications {on | off}]
[--source-snapshot-archive {on | off}]
[--source-snapshot-pattern <naming-pattern>]
[--source-snapshot-expiration <duration>]
[--report-max-age <duration>]
[--report-max-count <integer>]
[--restrict-target-network {on | off}]
[--source-subnet <subnet> --source-pool <pool>]
[--clear-source-network]
[--target-compare-initial-sync {on | off}]
[--accelerated-failback {yes | no}]
[--priority {0 | 1}]
[--cloud-deep-copy {deny | allow | force}]
[--verbose]
[--force]

```

Options

<policy>

Identifies the policy to modify, either by current policy ID or name.

{--name | -n} <new-policy-name>

Specifies a new name for this replication policy.

--action <policy-type>

Specifies the type of replication policy.

The following types of replication policy are valid:

copy	Creates a copy policy that adds copies of all files from the source to the target.
sync	Creates a synchronization policy that synchronizes data on the source cluster to the target cluster and deletes all files on the target cluster that are not present on the source cluster.

{--target-host | -C} <target-cluster>

Specifies the cluster that the policy replicates data to.

Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.

- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**

This will replicate data to another directory on the local cluster.

i **NOTE: SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.**

{--target-path | -p} <target-path>

Specifies the directory on the target cluster that files are replicated to.

Specify as a full directory path.

--source-root-path <root-path>

Specifies the directory on the local cluster that files are replicated from.

Specify as a full directory path.

--description <string>

Specifies a description of this replication policy.

--password <password>

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--set-password`.

--set-password

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--password`.

{--source-include-directories | -i} <path>

Includes only the specified directories in replication.

Specify as any directory path contained in the root directory. You can specify multiple directories by specifying `--source-include-directories` multiple times within a command. For example, if the root directory is `/ifs/data`, you could specify the following:

```
--source-include-directories /ifs/data/music --source-include-directories /ifs/data/movies
```

--clear-source-include-directories

Clears the list of included directories.

--add-source-include-directories <path>

Adds the specified directory to the list of included directories.

--remove-source-include-directories <path>

Removes the specified directory from the list of included directories.

{--source-exclude-directories | -e} <path>

Does not include the specified directories in replication.

Specify as any directory path contained in the root directory. If `--source-include-directories` is specified, `--source-exclude-directories` directories must be contained in the included directories. You can specify multiple directories by specifying `--source-exclude-directories` multiple times within a command. For example, you could specify the following:

```
--source-exclude-directories /ifs/data/music --source-exclude-directories /ifs/data/movies --exclude /ifs/data/music/working
```

--clear-source-exclude-directories

Clears the list of excluded directories.

--add-source-exclude-directories <path>

Adds the specified directory to the list of excluded directories.

--remove-source-exclude-directories <path>

Removes the specified directory from the list of excluded directories.

--begin-filter <predicate> --operator <value> [<predicate> --operator <operator> <link>]... --end-filter

Specifies the file-matching criteria that determines which files are replicated. Specify <predicate> as one or more of the following options:

The following options are valid for both copy and synchronization policies:

--size<integer>[<B | KB | MB | GB | TB | PB>]

Selects files according to the specified size.

--file-type <value>

Selects only the specified file-system object type.

The following values are valid:

f	Specifies regular files
d	Specifies directories
l	Specifies soft links

--name <value>

Selects only files whose names match the specified string.

You can include the following wildcards:

- . *
- . []
- . ?

The following options are valid only for copy policies:

--accessed-after '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that have been accessed since the specified time. This predicate is valid only for copy policies.

--accessed-before '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that have not been accessed since the specified time. This predicate is valid only for copy policies.

--accessed-time '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that were accessed at the specified time. This predicate is valid only for copy policies.

--birth-after '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that were created after the specified time. This predicate is valid only for copy policies.

--birth-before '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that were created before the specified time. This predicate is valid only for copy policies.

--birth-time '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that were created at the specified time. This predicate is valid only for copy policies.

--changed-after '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that have been modified since the specified time. This predicate is valid only for copy policies.

--changed-before '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that have not been modified since the specified time. This predicate is valid only for copy policies.

--changed-time '{<mm>/<dd>/<yyyy> [<HH>:<mm>] | <integer> {days | weeks | months | years} ago}'

Selects files that were modified at the specified time. This predicate is valid only for copy policies.

--no-group

Selects files based on whether they are owned by a group.

--no-user

Selects files based on whether they are owned by a user.

--posix-regex-name <value>

Selects only files whose names match the specified POSIX regular expression. IEEE Std 1003.2 (POSIX.2) regular expressions are supported.

--user-id <id>

Selects files based on whether they are owned by the user of the specified ID.

--user-name <name>

Selects files based on whether they are owned by the user of the specified name.

--group-id <id>

Selects files based on whether they are owned by the group of the specified ID.

--group-name <name>

Selects files based on whether they are owned by the group of the specified name.

The following <operator> values are valid:

Operator	Description
eq	Equal. This is the default value.
ne	Not equal
lt	Less than
le	Less than or equal to
gt	Greater than
ge	Greater than or equal to
not	Not

You can use the following <link> values to combine and alter the options available for predicates:

--and

Selects files that meet the criteria of the options that come before and after this value.

--or

Selects files that meet either the criterion of the option that comes before this value or the criterion of the option that follows this value.

{--schedule | -S} {<schedule> | when-source-modified}

Specifies how often data will be replicated. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified.

Specify <schedule> in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month

- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

To configure a policy to be run only manually, specify the following option:

```
--schedule ""
```

--skip-when-source-unmodified {true | false}

Causes the policy not to be run if the contents of the source directory have not been modified since the last time the policy has been run. If `--schedule` of this replication policy is set to <schedule>, and the policy is scheduled to run before changes have been made to the contents of the source directory, the policy will not be run.

--rpo-alert <duration>

Creates a OneFS event if the specified Recovery Point Objective (RPO) is exceeded. For example, assume you set an RPO of 5 hours; a job starts at 1:00 PM and completes at 3:00 PM; a second job starts at 3:30 PM; if the second job does not complete by 6:00 PM, SyncIQ will create a OneFS event.

The default value is 0, which will not generate events. This option is valid only if `--schedule` is set to <schedule>.

 **NOTE: This option is valid only if RPO alerts have been globally enabled through SyncIQ settings. The events have an event ID of 400040020.**

--job-delay <duration>

Specifies the amount of time after the source directory is modified that SyncIQ waits before starting a replication job. If the `--schedule` of this replication policy is set to `when-source-modified`, and the contents of the source directory are modified, SyncIQ will wait the specified amount of time before starting a replication job.

The default value is 0 seconds.

--clear-job-delay

Clears the amount of time after the source directory is modified that SyncIQ waits before starting a replication job.


--snapshot-sync-pattern <pattern>

Specifies the naming pattern for snapshots to be synced. If the `--schedule` of this replication policy is set to `when-snapshot-taken`, and a snapshot is taken of the source directory, and the snapshot name matches the specified naming pattern, SyncIQ will replicate the snapshot to the target cluster.

The default value is "*", which causes all snapshots of the source directory to be replicated if the `--schedule` of the policy is set to `when-snapshot-taken`.

--snapshot-sync-existing {yes | no}

Determines whether the policy replicates the data contained snapshots taken before the policy was created.

 **NOTE: Because this setting cannot be modified after the policy is initially created, this option cannot be specified with `isi sync policies modify`.**

--enabled {true | false}

Determines whether the policy is enabled or disabled.

--check-integrity {true | false}

Specifies whether to perform a checksum on each file data packet that is affected by the SyncIQ job. If this option is set to `true` and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is `true`.

--log-level <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug
- trace

The default value is `info`.

--log-removed-files {yes | no}

Determines whether SyncIQ retains a log of all files that are deleted when a synchronization policy is run. If the policy is a copy policy, this parameter has no effect.

The default value is `no`.

{--workers-per-node | -w} <integer>

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is 3.

 **NOTE: This option has been deprecated and will not be recognized if configured.**

--target-snapshot-archive {on | off}

Determines whether archival snapshots are generated on the target cluster. If this option is set to `off`, SyncIQ will still maintain exactly one snapshot at a time on the target cluster to facilitate failback. You must activate a SnapshotIQ license on the target cluster to generate archival snapshots on the target cluster.

--target-snapshot-pattern <naming-pattern>

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster.

The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

--target-snapshot-expiration <duration>

Specifies an expiration period for archival snapshots on the target cluster.

If this option is not specified, archival snapshots will remain indefinitely on the target cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

- | | |
|----------|------------------|
| Y | Specifies years |
| M | Specifies months |
| W | Specifies weeks |
| D | Specifies days |
| H | Specifies hours |

--target-snapshot-alias <naming-pattern>


Specifies a naming pattern for the most recent archival snapshot generated on the target cluster.

The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

--target-detect-modifications {on | off}

Determines whether SyncIQ checks the target directory for modifications before replicating files.

 **CAUTION: Specifying `off` could result in data loss. It is recommended that you consult Isilon Technical Support before specifying `off`.**

--source-snapshot-archive {on | off}

Determines whether archival snapshots are retained on the source cluster. If this option is set to `off`, SyncIQ will still maintain one snapshot at a time for the policy to facilitate replication.

--source-snapshot-pattern <naming-pattern>

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster.

For example, the following pattern is valid:

```
SIQ-source-#{PolicyName}-%Y-%m-%d_%H-%M
```

--source-snapshot-expiration <duration>

Specifies an expiration period for archival snapshots retained on the source cluster.

If this option is not specified, archival snapshots will exist indefinitely on the source cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--report-max-age <duration>

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--report-max-count <integer>

Specifies the maximum number of reports to retain for the replication policy.

--restrict-target-network {on | off}

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

--source-subnet <subnet>

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster.

--source-pool <pool>

Restricts replication jobs to running only on nodes in the specified pool on the local cluster.

--clear-source-network

Runs replication jobs on any nodes in the cluster, instead of restricting the jobs to a specified subnet.

--target-compare-initial-sync {on | off}

Determines whether the full or differential replications are performed for this policy. Full or differential replications are performed the first time a policy is run and after a policy has been reset. If set to `on`, performs a differential replication. If set to `off`, performs a full replication.

If differential replication is enabled the first time a replication policy is run, the policy will run slower without any benefit.

The default value is `off`.

--accelerated-failback {enable | disable}

If enabled, SyncIQ will perform failback configuration tasks the next time that a job is run, rather than waiting to perform those tasks during the failback process. Performing these tasks ahead of time will increase the speed of failback operations.

--priority {0 | 1}

Determines whether the policy has priority.

--cloud-deep-copy {deny | allow | force}

Determines how the policy replicates CloudPools smartlinks. If set to `deny`, SyncIQ replicates all CloudPools smartlinks to the target cluster as smartlinks; if the target cluster does not support the smartlinks, the job will fail. If set to `force`, SyncIQ replicates all smartlinks to the target cluster as regular files. If set to `allow`, SyncIQ will attempt to replicate smartlinks to the target cluster as smartlinks; if the target cluster does not support the smartlinks, SyncIQ will replicate the smartlinks as regular files.

{--verbose | -v}

Displays a confirmation message.

{--force | -f}

Does not prompt you to confirm modifications.

isi sync policies reset

Resets a replication policy after the policy encounters an error and the cause of the error cannot be identified or fixed. If you fix the cause of the error, run `isi sync policies resolve` instead.

Resetting a replication policy causes either a full replication or a differential replication to be performed the next time the policy is run.

Syntax

```
isi sync policies reset {<policy> | --all}
  [--verbose]
```

Options

<policy>

Resets the specified replication policy.

Specify as a replication policy name or ID

--all

Resets all replication policies

{--verbose | -v}

Displays more detailed information.

isi sync policies resolve

Resolves a conflicted replication policy after the policy encounters an error and the cause of the error is fixed. If the cause of the error cannot be fixed, run the `isi sync policies reset` command instead.

Syntax

```
isi sync policies resolve <policy>
[--force]
```

Options

<policy>

Resolves the specified replication policy.

Specify as a replication policy name or ID.

{--force | -f}

Suppresses command-line prompts and messages.

isi sync policies view

Displays information about a replication policy.

Syntax

```
isi sync policies view <policy>
```

Options

<policy>

Displays information about the specified replication policy.

Specify as a replication policy name or ID.

isi sync recovery allow-write

Allows modifications to data in a target directory of a replication policy without breaking the association between the local cluster and the policy. The `isi sync target allow_write` command is most commonly used in failover and failback operations.

Syntax

```
isi sync recovery allow-write <policy-name>
[--revert]
[--log-level <level>]
[--workers-per-node <integer>]
[--verbose]
```

Options

<policy-name>

Allows writes for the target directory of the specified replication policy.

Specify as a replication policy name, a replication policy ID, or the path of a target directory.

--revert

Reverts an allow-writes operation on the local cluster only. This action does not affect the source cluster of the replication policy.

--log-level <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug
- trace

The default value is `info`.

{--workers-per-node | -w}<integer>

Specifies the number of workers per node that are generated by SyncIQ to perform the allow-writes job.

The default value is 3.

{--verbose | -v}

Displays more detailed information.

isi sync recovery resync-prep

Disables the specified policy, reverts the source directory of the policy to the last recovery point, and creates a mirror policy on the target cluster. The `isi sync resync prep` command is most commonly used in failback operations.

Syntax

```
isi sync recovery resync-prep <policy-name>
[--verbose]
```

Options

<policy-name>

Targets the following replication policy.

Specify as a replication policy name or ID. The replication policy must be a synchronization policy.

--verbose

Displays more detailed information.

isi sync reports list

Displays information about completed replication jobs targeting a remote cluster.

Syntax

```
isi sync reports list
[--policy-name <policy>]
[--state <state>]
[--reports-per-policy <integer>]
[--limit <integer>]
[--sort <attribute>]
[--descending]
```

```
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--policy-name <policy>

Displays only replication reports that were created for the specified policy.

--state <state>

Displays only replication reports whose jobs are in the specified state.

--reports-per-policy <integer>

Displays no more than the specified number of reports per policy. The default value is 10.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

start_time

Sorts output by when the replication job started.

end_time

Sorts output by when the replication job ended.

action

Sorts output by the action that the replication job performed.

state

Sorts output by the progress of the replication job.

id

Sorts output by the ID of the replication subreport.

policy_id

Sorts output by the ID of the replication policy

policy_name

Sorts output by the name of the replication policy.

job_id

Sorts output by the ID of the replication job.

total_files

Sorts output by the total number of files that were modified by the replication job.

files_transferred

Sorts output by the total number of files that were transferred to the target cluster.

bytes_transferred

Sorts output by the total number of files that were transferred to the target cluster.

duration

Sorts output by how long the replication job ran.

errors

Sorts output by errors that the replication job encountered.

warnings

Sorts output by warnings that the replication job triggered.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync reports rotate

If the number of replication reports has exceeded the maximum, deletes replication reports. The system intermittently deletes excess reports automatically. However, this command causes excess reports to be deleted immediately.

Syntax

```
isi sync reports rotate
[--verbose]
```

Options

{--verbose | -v}

Displays more detailed information.

isi sync reports subreports list

Displays subreports about completed replication jobs targeting remote clusters.

Syntax

```
isi sync reports subreports list <policy> <job-id>
[--limit]
[--sort <attribute>]
[--descending]
[--format {table | json | csv | list}]
[--no-header]
[--no-footer]
[--verbose]
```

Options

<policy>

Displays subreports about the specified policy.

<job-id>

Displays subreports about the job of the specified ID.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

start_time

Sorts output by when the replication job started.

end_time

Sorts output by when the replication job ended.

action

Sorts output by the action that the replication job performed.

state

Sorts output by the progress of the replication job.

id

Sorts output by the ID of the replication report.

policy_id

Sorts output by the ID of the replication policy

policy_name

Sorts output by the name of the replication policy.

job_id

Sorts output by the ID of the replication job.

total_files

Sorts output by the total number of files that were modified by the replication job.

files_transferred

Sorts output by the total number of files that were transferred to the target cluster.

bytes_transferred

Sorts output by the total number of files that were transferred to the target cluster.

duration

Sorts output by how long the replication job ran.

errors

Sorts output by errors that the replication job encountered.

warnings

Sorts output by warnings that the replication job triggered.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync reports subreports view

Displays a subreport about a completed replication job that targeted a remote cluster.

Syntax

```
isi sync reports subreports view <policy> <job-id> <subreport-id>
```

Options

<policy>

Displays a sub report about the specified replication policy. Specify as a replication policy name.

<job-id>

Displays a sub report about the specified replication job. Specify as a replication job ID.

<subreport-id>

Displays the subreport of the specified ID.

isi sync reports view

Displays information about a completed replication job that targeted a remote cluster.

Syntax

```
isi sync reports view <policy> <job-id>
```

Options

<policy>

Displays a replication report about the specified replication policy.

<job-id>

Displays a replication report about the job with the specified ID.

isi sync rules create

Creates a replication performance rule.

Syntax

```
isi sync rules create <type> <interval> <days> <limit>  
  [--description <string>]  
  [--verbose]
```

Options

<type>

Specifies the type of performance rule. The following values are valid:

file_count

Creates a performance rule that limits the number of files that can be sent by replication jobs per second.

bandwidth

Creates a performance rule that limits the amount of bandwidth that replication jobs are allowed to consume.

<interval>

Enforces the performance rule on the specified hours of the day. Specify in the following format:

```
<hh>:<mm>-<hh>:<mm>
```

<days>

Enforces the performance rule on the specified days of the week.

The following values are valid:

X	Specifies Sunday
M	Specifies Monday
T	Specifies Tuesday
W	Specifies Wednesday
R	Specifies Thursday
F	Specifies Friday
S	Specifies Saturday

You can include multiple days by specifying multiple values separated by commas. You can also include a range of days by specifying two values separated by a dash.

<limit>

Specifies the maximum number of files that can be sent or KBs that can be consumed per second by replication jobs.

--description <string>

Specifies a description of this performance rule.

--verbose

Displays more detailed information.

isi sync rules delete

Deletes a replication performance rule.

Syntax

```
isi sync rules delete {<id> | --all | --type <type>}  
  [--force]  
  [--verbose]
```

Options

<id>

Deletes the performance rule of the specified ID.

--all

Deletes all performance rules.

--type <type>

Deletes all performance rules of the specified type. The following values are valid:

file_count

Deletes all performance rules that limit the number of files that can be sent by replication jobs per second.

bandwidth

Deletes all performance rules that limit the amount of bandwidth that replication jobs are allowed to consume.

--force

Does not prompt you to confirm that you want to delete the performance rule.

--verbose

Displays more detailed information.

isi sync rules list

Displays a list of replication performance rules.

Syntax

```
isi sync rules list
  [--type <type>]
  [--limit]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--type <type>

Displays only performance rules of the specified type. The following values are valid:

file_count

Displays only performance rules that limit the number of files that can be sent by replication jobs per second.

bandwidth

Displays only performance rules that limit the amount of bandwidth that replication jobs are allowed to consume.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync rules modify

Modifies a replication performance rule.

Syntax

```
isi sync rules modify <id>
  [--interval <interval>]
  [--days <days>]
  [--limit <integer>]
  [--enabled {true | false}]
  [--description <string>]
  [--verbose]
```

Options

<id>

Modifies the replication performance rule of the specified ID.

{--interval | -i} <interval>

Specifies which hours of the day to enforce the performance rule. Specify in the following format:

```
<hh>: <mm>--<hh>: <mm>
```

{--days | -d} <days>

Specifies which days of the week to enforce the performance rule.

The following values are valid:

X	Specifies Sunday
M	Specifies Monday
T	Specifies Tuesday
W	Specifies Wednesday
R	Specifies Thursday
F	Specifies Friday
S	Specifies Saturday

You can include multiple days by specifying multiple values separated by commas. You can also include a range of days by specifying two values separated by a dash.

--limit <limit>

Specifies the maximum number of files that can be sent or KBs that can be consumed per second by replication jobs.

--enabled {true | false}

Determines whether the policy is enabled or disabled.

--description <string>

Specifies a description of this performance rule.

{--verbose | -v}

Displays more detailed information.

isi sync rules reports list

List SyncIQ bandwidth reports for running jobs.

Syntax

```
isi sync rules reports list
  [--policy-id <string>]
  [--start <integer>]
  [--end <integer>]
  [--sort (name | speed)]
  [--descending]
```

Options

--policy-id <string>

A SyncIQ policy identifier.

{--start | -s} <integer>

Beginning time stamp for the bandwidth report.

{--end | -e} <integer>

Ending time stamp for the bandwidth report.

--sort (*name* | *speed*)

Sort data by the specified field.

{--descending | -d}

Sort data in descending order.

isi sync rules reports view

View bandwidth for a running SynclQ job.

Syntax

```
isi sync rules reports view <policy-id>
  [--start <integer>]
  [--end <integer>]
```

Options

<policy-id>

A SynclQ policy identifier.

{--start | -s} <integer>

Beginning time stamp for the bandwidth report.

{--end | -e} <integer>

Ending time stamp for the bandwidth report.

isi sync rules view

Displays information about a replication performance rule.

Syntax

```
isi sync rules view <id>
```

Options

<id>

Displays information about the replication performance rule with the specified ID.

isi sync service policies create

Create a SynclQ service policy.

Syntax

```
isi sync service policies create <name> <target-host>
  [--enabled (yes | no)]
  [--description <string>]
  [--check-integrity (yes | no)]
  [--source-subnet <subnet> | --source-pool <pool>]
  [--target-snapshot-pattern <naming-pattern>]
  [--target-snapshot-expiration <duration>]
  [--target-snapshot-alias <naming-pattern>]
  [--source-snapshot-pattern <naming-pattern>]
```

```

[--source-snapshot-expiration <duration>]
[--snapshot-sync-pattern <pattern>]
[--snapshot-sync-existing (yes | no)]
[--schedule (<schedule> | when-source-modified
| when-snapshot-taken)]
[--rpo-alert <duration>]
[--log-level <level>]
[--workers-per-node <integer>]
[--report-max-age <duration>]
[--report-max-count <integer>]
[--restrict-target-network (on | off)]
[--target-compare-initial-sync (on | off)]
[--accelerated-failback (yes | no)]
[--priority (0 | 1 | normal | high)]
[--bandwidth-reservation <integer>]
[--target-certificate-id <string>]
[--ocsp-issuer-certificate-id <string>]
[--ocsp-address <string>]
[--encryption-cipher-list <string>]
[--linked-data-policies <string>]
[--replicated-services <string>]
[--service-history-max-age <duration>]
[--service-history-max-count <integer>]
[--password (<password>)]
[--set-password]
[--verbose]

```

Options

<name>

Specifies a name for the replication service policy.

Specify as any string.


<target-host>

Specifies the cluster that the policy replicates data to.

Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**

This will replicate data to another directory on the local cluster.

 NOTE: SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

--enabled (yes | no)

Determines whether the service policy is enabled or disabled.

The default value is yes.

--description <string>

Specifies a description of the replication service policy.

--check-integrity (yes | no)

Specifies whether to perform a checksum on each file data packet that is affected by the replication policy. If this option is set to yes, and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is yes.

--source-subnet <subnet>

Restricts replication policies to running only on nodes in the specified subnet on the local cluster. If you specify this option, you must also specify --source-pool.

--source-pool <pool>

Restricts replication policies to running only on nodes in the specified pool on the local cluster. If you specify this option, you must also specify `--source-subnet`.

--target-snapshot-pattern <naming-pattern>

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster. The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

--target-snapshot-expiration <duration>

Specifies an expiration period for archival snapshots on the target cluster.

If this option is not specified, archival snapshots will remain indefinitely on the target cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--target-snapshot-alias <naming-pattern>

Specifies a naming pattern for the most recent archival snapshot generated on the target cluster.

The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

--source-snapshot-pattern <naming-pattern>

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster.

For example, the following pattern is valid:

```
SIQ-source-#{PolicyName}-%Y-%m-%d_%H-%M
```

--source-snapshot-expiration <duration>

Specifies an expiration period for archival snapshots retained on the source cluster.

If this option is not specified, archival snapshots will exist indefinitely on the source cluster.

Specify in the following format:

```
<integer><units>
```

The following *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--snapshot-sync-pattern <string>

The naming pattern that a snapshot must match to trigger a replication job, when the schedule is set to **when-snapshot-taken**. The default value is asterisk (*).

`--snapshot-sync-existing` (yes | no)

If set to Yes, snapshot-triggered replication jobs will include replications taken before the policy creation time. The default is No. If set to yes, set `--schedule when-snapshot-taken`.

`{--schedule | -S}` (<schedule> | when-source-modified | when-snapshot-taken)

Specifies how often replication policy is executed. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified. Specifying `when-snapshot-taken` causes OneFS to replicate data every time that a snapshot is taken of the source directory.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify <interval> in one of the following formats:

- Every [{other | <integer>}] {weekday | day}
- Every [{other | <integer>}] week [on <day>]
- Every [{other | <integer>}] month [on the <integer>]
- Every [<day>[, ...] [of every [{other | <integer>}] week]]
- The last {day | weekday | <day>} of every [{other | <integer>}] month
- The <integer> {weekday | <day>} of every [{other | <integer>}] month
- Yearly on <month> <integer>
- Yearly on the {last | <integer>} [weekday | <day>] of <month>

Specify <frequency> in one of the following formats:

- at <hh>[:<mm>] [{AM | PM}]
- every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]
- every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]

You can optionally append "st", "th", or "rd" to <integer>. For example, you can specify "Every 1st month"

Specify <day> as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--rpo-alert` <duration>

Creates a OneFS event if the specified Recovery Point Objective (RPO) is exceeded. For example, assume you set an RPO of 5 hours; a job starts at 1:00 PM and completes at 3:00 PM; a second job starts at 3:30 PM; if the second job does not complete by 6:00 PM, SyncIQ will create a OneFS event.

The default value is 0, which will not generate events. This option is valid only if `--schedule` is set to <schedule>.



NOTE: This option is valid only if RPO alerts have been globally enabled through SyncIQ settings. The events have an event ID of 400040020.

`--log-level` <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice

- info
- copy
- debug
- trace

The default value is `info`.

`--workers-per-node | -w` <integer>

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is 3.

 **NOTE: This option has been deprecated and will not be recognized if configured.**

`--report-max-age` <duration>

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

`--report-max-count` <integer>

Specifies the maximum number of reports to retain for the replication policy.

`--restrict-target-network` (on | off)

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

`--accelerated-failback` (enable | disable)

If enabled, SyncIQ will perform failback configuration tasks the next time that a job is run, rather than waiting to perform those tasks during the failback process. Performing these tasks ahead of time will increase the speed of failback operations.

`--priority` (0 | 1)

Determines whether the policy has priority.

The default value is 0, which means that the policy does not have priority. If set to 1, the policy is high-priority.

`--bandwidth-reservation` <integer>

The desired bandwidth reservation for this policy, in kb/s. This feature does not activate unless a SyncIQ bandwidth rule is in effect.

`--target-certificate-id` <string>

The identifier of the target cluster certificate being used for encryption.

`--ocsp-issuer-certificate-id` <string>

The identifier of the certificate authority that issued the certificate whose revocation status is being checked.

`--ocsp-address` <string>

The address of the OCSP responder to which you want to connect.

`--encryption-cipher-list` <string>

The cipher list being used with cluster encryption. For SyncIQ targets, this is a list of supported ciphers. For SyncIQ sources, the list of ciphers is used in order.

`--linked-data-policies` <string>...

A list of SyncIQ policy identifiers whose source root directories will be used to filter service replication. Specify this option again for each additional service policy identifier.

--replicated-services <string>

A list of services to replicate. Specify again for each additional service.

--service-history-max-age <duration>

The maximum age of service information to maintain.

--service-history-max-count <integer>

The maximum number of service historical information records to maintain.

--password <password>

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--set-password`.

--set-password

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--password`.

{--verbose | -v}

Displays a message confirming that the snapshot schedule was created.

isi sync service policies delete

Delete one or all SyncIQ service policies.

Syntax

```
isi sync service policies delete <policy> | --all  
  [--local-only]  
  [--force]  
  [--verbose]
```

Options

<policy>

Deletes the specified replication service policy.

--all

Deletes all replication service policies.


--local-only

Does not break the target association on the target cluster. Not deleting a policy association on the target cluster will cause the target directory to remain in a read-only state.

 **NOTE: If SyncIQ is unable to communicate with the target cluster, you must specify this option to successfully delete the service policy.**

{--force | -f}

Deletes the service policy, even if an associated job is currently running. Also, does not prompt you to confirm the deletion.

 **CAUTION: Forcing a service policy to delete might cause errors if an associated replication job is currently running.**

{--verbose | -v}

Displays a confirmation message.

isi sync service policies disable

Disable one or all replication service policies.

Syntax

```
isi sync service policies disable <policy> | --all  
[--verbose]
```

Options

<policy>

Disables the specified replication policy. Specify as a replication policy name or a replication policy ID.

--all

Disables all replication policies on the cluster.

{--verbose | -v}

Displays more detailed information.

isi sync service policies enable

Enable one or all replication service policies.

Syntax

```
isi sync service policies enable <policy> | --all  
[--verbose]
```

Options

<policy>

Enables the specified replication service policy. Specify as a replication service policy name or ID.

--all

Enables all replication service policies on the cluster.

{--verbose | -v}

Displays more detailed information.

isi sync service policies list

List replication service policies.

Syntax

```
isi sync service policies list  
  [--limit <integer>]  
  [--sort (id | name | source_root_path | enabled | target_host | target_path |  
target_snapshot_pattern | target_snapshot_expiration | target_snapshot_alias |  
source_snapshot_pattern | source_snapshot_expiration | snapshot_sync_pattern |  
snapshot_sync_existing | schedule | rpo_alert | log_level | workers_per_node | report_max_age  
| report_max_count | force_interface | restrict_target_network | expected_data_loss |  
disable_fofb | disable_file_split | accelerated_failback | database_mirrored |  
source_domain_marked | priority | bandwidth_reservation | last_job_state | last_started |  
last_success | password_set | conflicted | has_sync_date | source_certificate_id |
```

```
ocsp_address | encryption_cipher_list | encrypted | linked_data_policies |
replicated_services | service_history_max_age | service_history_max_count) ]
[--descending]
[--format (table | json | csv | list)]
[--no-header]
[--no-footer]
[--verbose]
```

Options

If no options are specified, displays a table of all replication service policies.

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

{--descending | -d}

Displays output in reverse order.

--format (table | json | csv | list)

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync service policies modify

Modify a replication service policy.

Syntax

```
isi sync service policies modify <policy>
  [--name <service_policy_name>]
  [--enabled (yes | no)]
  [--target-host <target-cluster>]
  [--description <service-policy-description>]
  [--check-integrity (yes | no)]
  [--source-subnet <subnet> --source-pool <pool>]
  [--target-snapshot-pattern <naming-pattern>]
  [--target-snapshot-expiration <duration>]
  [--target-snapshot-alias <naming-pattern>]
  [--source-snapshot-pattern <naming-pattern>]
  [--source-snapshot-expiration <duration>]
  [--snapshot-sync-pattern <pattern>]
  [--snapshot-sync-existing (yes | no)]
  [--schedule (<schedule> | when-source-modified | when-snapshot-taken)]
  [--rpo-alert <duration>]
  [--clear-rpo-alert]
  [--log-level <level>]
  [--workers-per-node <integer> | --clear-rpo]
  [--report-max-age <duration>]
  [--report-max-count <integer>]
  [--restrict-target-network (on | off)]
  [--accelerated-failback (yes | no)]
  [--priority (0 | 1 | normal | high)]
  [--bandwidth-reservation <integer>]
  [--clear-bandwidth-reservation]
```



```

[--target-certificate-id <string>]
[--ocsp-issuer-certificate-id <string>]
[--ocsp-address <string>]
[--encryption-cipher-list <string>]
[--linked-data-policies <string>]
[--clear-linked-data-policies]
[--add-linked-data-policies <string>...]
[--remove-linked-data-policies <string>...]
[--replicated-services <string>]
[--clear-replicated-services]
[--add-replicated-services <string>...]
[--remove-replicated-services <string>...]
[--service-history-max-age <duration>]
[--service-history-max-count <integer>]
[--password (<password>]
[--set-password]
[--clear-source-network]
[--verbose]
[--force]

```

Options

<policy>

A replication service policy name.

{--name | -n} <service_policy_name>

Specifies a name for the replication service policy.

Specify as any string.

--enabled (yes | no)

Determines whether the service policy is enabled or disabled.

The default value is *yes*.

{--target-host | -C} <target-cluster>

Specifies the cluster that the service policy replicates data to.

Specify as one of the following:

- The fully qualified domain name of any node in the target cluster.
- The host name of any node in the target cluster.
- The name of a SmartConnect zone in the target cluster.
- The IPv4 or IPv6 address of any node in the target cluster.
- **localhost**

This will replicate data to another directory on the local cluster.



NOTE: SyncIQ does not support dynamically allocated IP address pools. If a replication job connects to a dynamically allocated IP address, SmartConnect might reassign the address while a replication job is running, which would disconnect the job and cause it to fail.

--description <service-policy-description>

Specifies a description of the replication service policy.

--check-integrity (yes | no)

Specifies whether to perform a checksum on each file data packet that is affected by the replication policy. If this option is set to *yes*, and the checksum values do not match, SyncIQ retransmits the file data packet.

The default value is *yes*.

--source-subnet <subnet>

Restricts replication service policies to running only on nodes in the specified subnet on the local cluster. If you specify this option, you must also specify `--source-pool`.

--source-pool <pool>

Restricts replication service policies to running only on nodes in the specified pool on the local cluster. If you specify this option, you must also specify `--source-subnet`.

--target-snapshot-pattern < naming-pattern >

Specifies the snapshot naming pattern for snapshots that are generated by replication jobs on the target cluster. The default naming pattern is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-%Y-%m-%d_%H-%M
```

--target-snapshot-expiration < duration >

Specifies an expiration period for archival snapshots on the target cluster. If this option is not specified, archival snapshots will remain indefinitely on the target cluster. Specify in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--target-snapshot-alias < naming-pattern >

Specifies a naming pattern for the most recent archival snapshot generated on the target cluster. The default alias is the following string:

```
SIQ-#{SrcCluster}-#{PolicyName}-latest
```

--source-snapshot-pattern < naming-pattern >

Specifies a naming pattern for the most recent archival snapshot generated on the source cluster. For example, the following pattern is valid:

```
SIQ-source-#{PolicyName}-%Y-%m-%d_%H-%M
```

--source-snapshot-expiration < duration >

Specifies an expiration period for archival snapshots retained on the source cluster. If this option is not specified, archival snapshots will exist indefinitely on the source cluster. Specify in the following format:

```
<integer><units>
```

The following <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

--snapshot-sync-pattern < string >

The naming pattern that a snapshot must match to trigger a replication job, when the schedule is set to **when-snapshot-taken**. The default value is asterisk (*).

--snapshot-sync-existing (yes | no)

If set to Yes, snapshot-triggered replication jobs will include replications taken before the policy creation time. The default is No. If set to yes, set **--schedule when-snapshot-taken**.

`--schedule | -S` (`<schedule>` | `when-source-modified` | `when-snapshot-taken`)

Specifies how often replication policy is executed. Specifying `when-source-modified` causes OneFS to replicate data every time that the source directory of the policy is modified. Specifying `when-snapshot-taken` causes OneFS to replicate data every time that a snapshot is taken of the source directory.

Specify in the following format:

```
"<interval> [<frequency>]"
```

Specify `<interval>` in one of the following formats:

- `Every [{other | <integer>}] {weekday | day}`
- `Every [{other | <integer>}] week [on <day>]`
- `Every [{other | <integer>}] month [on the <integer>]`
- `Every [<day>[, ...] [of every [{other | <integer>}] week]]`
- `The last {day | weekday | <day>} of every [{other | <integer>}] month`
- `The <integer> {weekday | <day>} of every [{other | <integer>}] month`
- `Yearly on <month> <integer>`
- `Yearly on the {last | <integer>} [weekday | <day>] of <month>`

Specify `<frequency>` in one of the following formats:

- `at <hh>[:<mm>] [{AM | PM}]`
- `every [<integer>] {hours | minutes} [between <hh>[:<mm>] [{AM | PM}] and <hh>[:<mm>] [{AM | PM}]]`
- `every [<integer>] {hours | minutes} [from <hh>[:<mm>] [{AM | PM}] to <hh>[:<mm>] [{AM | PM}]]`

You can optionally append "st", "th", or "rd" to `<integer>`. For example, you can specify "Every 1st month"

Specify `<day>` as any day of the week or a three-letter abbreviation for the day. For example, both "saturday" and "sat" are valid.

`--rpo-alert <duration>`

Creates a OneFS event if the specified Recovery Point Objective (RPO) is exceeded. For example, assume you set an RPO of 5 hours; a job starts at 1:00 PM and completes at 3:00 PM; a second job starts at 3:30 PM; if the second job does not complete by 6:00 PM, SyncIQ will create a OneFS event.

The default value is 0, which will not generate events. This option is valid only if `--schedule` is set to `<schedule>`.

 **NOTE: This option is valid only if RPO alerts have been globally enabled through SyncIQ settings. The events have an event ID of 400040020.**

`--clear-rpo-alert`

Clears an RPO alert.

`--log-level <level>`

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info

- copy
- debug
- trace

The default value is `info`.

`{--workers-per-node | -w} <integer>`

Specifies the number of workers per node that are generated by SyncIQ to perform each replication job for the policy.

The default value is 3.

 **NOTE: This option has been deprecated and will not be recognized if configured.**

`--report-max-age <duration>`

Specifies how long replication reports are retained before they are automatically deleted by SyncIQ.

Specify in the following format:

```
<integer><units>
```

The following `<units>` are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours

`--report-max-count <integer>`

Specifies the maximum number of reports to retain for the replication policy.

`--restrict-target-network (on | off)`

If you specify `on`, and you specify the target cluster as a SmartConnect zone, replication jobs connect only to nodes in the specified zone. If `off` is specified, does not restrict replication jobs to specific nodes on the target cluster.

`--accelerated-failback (enable | disable)`

If enabled, SyncIQ will perform failback configuration tasks the next time that a job is run, rather than waiting to perform those tasks during the failback process. Performing these tasks ahead of time will increase the speed of failback operations.

`--priority (0 | 1)`

Determines whether the policy has priority.

The default value is 0, which means that the policy does not have priority. If set to 1, the policy is high-priority.

`--bandwidth-reservation <integer>`

The desired bandwidth reservation for this policy, in kb/s. This feature does not activate unless a SyncIQ bandwidth rule is in effect.

`--clear-bandwidth-reservation`

Clears a bandwidth reservation for this service policy.

`--target-certificate-id <string>`

The identifier of the target cluster certificate being used for encryption.

`--ocsp-issuer-certificate-id <string>`

The identifier of the certificate authority that issued the certificate whose revocation status is being checked.

`--ocsp-address <string>`

The address of the OCSP responder to which you want to connect.

`--encryption-cipher-list <string>`

The cipher list being used with cluster encryption. For SyncIQ targets, this is a list of supported ciphers. For SyncIQ sources, the list of ciphers is used in order.

--linked-data-policies <string>...

A list of SyncIQ policy identifiers whose source root directories will be used to filter service replication. Specify this option again for each additional service policy identifier.

--clear-linked-data-policies

Clear the list of SyncIQ policy identifiers.

--add-linked-data-policies <string>...

Add items to the list of SyncIQ policy identifiers. Repeat for multiple policies.

--remove-linked-data-policies <string>...

Remove items from the list of SyncIQ policy identifiers. Repeat for multiple policies.

--replicated-services <string>...

A list of services to replicate. Specify again for each additional service.

--clear-replicated-services

Clear the list of services to replicate.

--add-replicated-services <string>...

Add items to the list of services to replicate. Repeat for multiple replicated services.

--remove-replicated-services <string>...

Remove items from the list of services to replicate. Repeat for multiple replicated services.

--service-history-max-age <duration>

The maximum age of service information to maintain.

--service-history-max-count <integer>

The maximum number of service historical information records to maintain.

--password <password>

Specifies a password to access the target cluster. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--set-password`.

--set-password

Prompts you to specify a password for the target cluster after the command is run. This can be useful if you do not want other users on the cluster to see the password you specify. If the target cluster requires a password for authentication purposes, you must specify this parameter or `--password`.

--clear-source-network

Clears the source subnet and pool.

{--verbose | -v}

Displays a message confirming that the snapshot schedule was created.

{--force | -f}

Does not prompt you to confirm modifications.

isi sync service policies reset

Resets a replication service policy after the policy encounters an error and the cause of the error cannot be identified or fixed. If you fix the cause of the error, run `isi sync service policies resolve` instead.

Resetting a replication service policy causes either a full replication or a differential replication to be performed the next time the policy is run.

Syntax

```
isi sync service policies reset <policy> | --all  
[--verbose]
```

Options

<policy>

Resets the specified replication service policy. Specify as a replication service policy name or ID

--all

Resets all replication policies

{--verbose | -v}

Displays more detailed information.

isi sync service policies resolve

Resolves a conflicted replication service policy after the policy encounters an error and the cause of the error is fixed. If the cause of the error cannot be fixed, contact Technical Support.

Syntax

```
isi sync service policies resolve <policy>
[--force]
```

Options

<policy>

Resolves the specified replication policy.

Specify as a replication policy name or ID.

{--force | -f}

Suppresses command-line prompts and messages.

isi sync service policies view

Displays information about a replication service policy.

Syntax

```
isi sync service policies view <policy>
```

Options

<policy>

Displays information about the specified replication service policy.

Specify as a replication service policy name or ID.

isi sync service recovery allow-write

Allow writes to a policy target path.

Syntax

```
isi sync service recovery allow-write <policy-name> <timestamp> <tgt-path>
[--log-level (fatal | error | notice | info | copy | debug | trace)]
[--skip-copy (yes | no)]
[--skip-map (yes | no)]
[--skip-failover (yes | no)]
[--verbose]
```

Options

<policy-name>

The policy name for the job to fail over/fail back.

<timestamp>

The time stamp for a service replication policy backup from which you are restoring.

<tgt-path>

The directory to output the service replication files for restoration.

--log-level <level>

Specifies the amount of data recorded in logs.

The following values are valid, organized from least to most information:

- fatal
- error
- notice
- info
- copy
- debug
- trace

The default value is `info`.

--skip-copy (yes | no)

Skips the copy phase of the service replication allow-write operation.

--skip-map (yes | no)

Skips the mapping phase of the service replication allow-write operation.

--skip-failover (yes | no)

Skips the data failover phase of the service replication allow-write operation.

{--verbose | -v}

Displays more detailed information.

isi sync service recovery resync-prep

Prepare a service policy for re-synchronization.

Syntax

```
isi sync service recovery resync-prep <policy-name>
[--verbose]
```

Options

<policy-name>

The policy name for the job you are preparing for re-synchronization.

{--verbose | -v}

Displays more detailed information.

isi sync service target break

Removes a service policy replication target association.

Syntax

```
isi sync service target break <policy> | --target-path <path>
  [--force]
  [--verbose]
```

Options

<policy>

A target replication policy name.

<target-path>

A target path for a replication policy.

{--force | -f}

Break the target association even if a replication job is running. Do not ask for confirmation.

{--verbose | -v}

Displays more detailed information.

isi sync service target cancel

Cancels a currently running service replication job targeted to this cluster.

Syntax

```
isi sync service target cancel <policy> | --target-path <path> | --all
  [--verbose]
```

Options

<policy>

A target replication policy name.

<target-path>

A target path for a replication policy.

--all

Cancel all running jobs targeted to this cluster.

{--verbose | -v}

Displays more detailed information.

isi sync service target list

List target policies.

Syntax

```
isi sync service target list
  [--target-path <path>]
  [--limit <integer>]
```



```
[--sort (name | source_host | target_path | last_job_state | failover_failback_state)]
[--descending]
[--format (table | json | csv | list)]
[--no-header]
[--no-footer]
[--verbose]
```

Options

--target-path <path>

Show target policies that have the specified target path.

--sort (name | source_host | target_path | last_job_state | failover_failback_state)

Sorts output displayed by the specified attribute.

{--descending | -d}

Displays output in reverse order.

--format (table | json | csv | list)

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync service target view

View properties of a target service replication policy.

Syntax

```
isi sync service target view <policy> | --target-path <path>
```

Options

<policy>

A target replication policy name.

<target-path>

A target path for a replication policy.

isi sync settings modify

Manages global replication settings.

Syntax

```
isi sync settings modify
  [--service(on | off | paused)]
  [--source-subnet <string>]
  [--source-pool <string>]
  [--restrict-target-network <boolean>]
  [--report-max-age <duration>]
```

```

[--report-max-count <integer>]
[--rpo-alerts <boolean>]
[--bandwidth-reservation-reserve-percentage <integer>]
[--bandwidth-reservation-reserve-absolute <integer> |
  --clear-bandwidth-reservation-reserve-absolute]
[--encryption-required <boolean>]
[--cluster-certificate-id <string>]
[--ocsp-issuer-certificate-id <string>]
[--ocsp-address <string>]
[--encryption-cipher-list <string>]
[--renegotiation-period <duration>]
[--service-history-max-age <duration>]
[--service-history-max-count <integer>]
[--use-workers-per-node <boolean>]
[{-v}]
[{-h}]

```

Options

If no options are specified, displays current default replication report settings.

--service {on | off | paused}

Determines the state of the SynclQ application.

--source-subnet <subnet>

Restricts replication jobs to running only on nodes in the specified subnet on the local cluster.

--source-pool <pool>

Restricts replication jobs to running only on nodes in the specified pool on the local cluster.

--restrict-target-network <boolean>

Restricts target to its nodes in the target zone name.

--report-max-age <duration>

Specifies the default maximum age of reports to retain for a replication policy.

--report-max-count <integer>

Specifies the default maximum number of reports to retain for a replication policy.

--rpo-alerts <boolean>

If disabled, no RPO alerts are generated.

--bandwidth-reservation-reserve-percentage <integer>

Specifies the percentage of SynclQ bandwidth to reserve for policies that did not specify a bandwidth reservation.

--bandwidth-reservation-reserve-absolute <integer>

Specifies the amount of SynclQ bandwidth to reserve in kb/s for policies that did not specify a bandwidth reservation. This field takes precedence over `bandwidth_reservation_reserve_percentage`.

--clear-bandwidth-reservation-reserve-absolute

Specifies the clear absolute bandwidth reservation.

--encryption-required <boolean>

Requires all SynclQ policies to utilize encrypted communications, if true.

--cluster-certificate-id <string>

Specifies the ID of the cluster's certificate to be used for encryption.

--ocsp-issuer-certificate-id <string>

Specifies the ID of the certificate authority that issued the certificate whose revocation status is being checked.

--ocsp-address <string>

Specifies the address of the OCSP responder to which it is to connect.

--encryption-cipher-list <string>

Specifies the cipher list that is being used with encryption.



NOTE: For SynclQ targets, this list serves as a list of supported ciphers. For SynclQ sources, the list of ciphers will be attempted to be used in order.

--renegotiation-period <duration>

Specifies the duration to persist encrypted connections before forcing a renegotiation.

--service-history-max-age <duration>

Specifies the maximum age of service information to maintain.

--service-history-max-count <integer>

Specifies the maximum number of historical service information records to maintain.

--use-workers-per-node <boolean>

If enabled, SyncIQ uses the deprecated `workers_per_node` field with worker pools functionality and limits workers accordingly.

{--verbose | -v}

Displays more detailed information.

{--help | -h}

Display help for this command.

isi sync settings view

Displays global replication settings.

Syntax

```
isi sync settings view
```

Options

There are no options for this command.

isi sync target break

Breaks the association between a local cluster and a target cluster for a replication policy.

i NOTE:

Breaking a source and target association requires you to reset the replication policy before you can run the policy again. Depending on the amount of data being replicated, a full or differential replication can take a very long time to complete.

Syntax

```
isi sync target break {<policy> | --target-path <path>}  
  [--force]  
  [--verbose]
```

Options

<policy>

Removes the association of the specified replication policy targeting this cluster.

Specify as a replication policy name, a replication policy ID, or the path of a target directory.

--target-path <path>

Removes the association of the replication policy targeting the specified directory path.

{--force | -f}

Forces the replication policy association to be removed, even if an associated job is currently running.

 **CAUTION:** Forcing a target break might cause errors if an associated replication job is currently running.

{--verbose | -v}

Displays more detailed information.

isi sync target cancel

Cancels running replication jobs targeting the local cluster.

Syntax

```
isi sync target cancel {<policy> | --target-path <path> | --all}
[--verbose]
```

Options

<policy>

Cancels a replication job created according to the specified replication policy.

Specify as a replication policy name or ID.

--target-path <path>

Cancels a replication job targeting the specified directory.

--all

Cancels all running replication jobs targeting the local cluster.

--verbose

Displays more detailed information.

isi sync target list

Displays a list of replication policies targeting the local cluster.

Syntax

```
isi sync target list
  [--target-path <path>]
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, displays a table of all replication policies currently targeting the local cluster.

--target-path <path>

Displays information about the replication policy targeting the specified directory.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

name	Sorts output by the name of the replication policy.
source_host	Sorts output by the name of the source cluster.
target_path	Sorts output by the path of the target directory.
last_job_status	Sorts output by the status of the last replication job created according to the policy.
failover_failback_state	Sorts output by whether the target directory is read only.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync target reports list

Displays information about completed replication jobs targeting the local cluster.

Syntax

```
isi sync target reports list
  [--state <state>]
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

If no options are specified, displays basic information about all completed replication jobs.

--state <state>

Displays information about only replication jobs in the specified state. The following states are valid:

- scheduled
- running
- paused
- finished
- failed
- canceled
- needs_attention
- unknown

{--limit | -1} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

start_time

Sorts output by when the replication job started.

end_time

Sorts output by when the replication job ended.

action

Sorts output by the action that the replication job performed.

state

Sorts output by the progress of the replication job.

id

Sorts output by the ID of the replication subreport.

policy_id

Sorts output by the ID of the replication policy

policy_name

Sorts output by the name of the replication policy.

job_id

Sorts output by the ID of the replication job.

total_files

Sorts output by the total number of files that were modified by the replication job.

files_transferred

Sorts output by the total number of files that were transferred to the target cluster.

bytes_transferred

Sorts output by the total number of files that were transferred to the target cluster.

duration

Sorts output by how long the replication job ran.

errors

Sorts output by errors that the replication job encountered.

warnings

Sorts output by warnings that the replication job triggered.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync target reports subreports list

Displays subreports about completed replication jobs targeting the local cluster.

Syntax

```
isi sync target reports subreports list <policy> <job-id>
  [--limit]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

<policy>

Displays subreports about the specified policy.

<job-id>

Displays subreports about the job of the specified ID.

{--limit | -1} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

start_time

Sorts output by when the replication job started.

end_time

Sorts output by when the replication job ended.

action

Sorts output by the action that the replication job performed.

state

Sorts output by the progress of the replication job.

id

Sorts output by the ID of the replication report.

policy_id

Sorts output by the ID of the replication policy

policy_name

Sorts output by the name of the replication policy.

job_id

Sorts output by the ID of the replication job.

total_files

Sorts output by the total number of files that were modified by the replication job.

files_transferred

Sorts output by the total number of files that were transferred to the target cluster.

bytes_transferred

Sorts output by the total number of files that were transferred to the target cluster.

duration

Sorts output by how long the replication job ran.

errors

Sorts output by errors that the replication job encountered.

warnings

Sorts output by warnings that the replication job triggered.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

isi sync target reports subreports view

Displays a subreport about a completed replication job targeting the local cluster.

Syntax

```
isi sync target reports subreports view <policy> <job-id> <subreport-id>
```

Options

<policy>

Displays a sub report about the specified replication policy. Specify as a replication policy name.

<job-id>

Displays a sub report about the specified replication job. Specify as a replication job ID.

<subreport-id>

Displays the subreport with the specified ID.

isi sync target reports view

Displays information about a completed replication job that targeted the local cluster.

Syntax

```
isi sync target reports view <policy> <job-id>
```

Options

<policy>

Displays a replication report about the specified replication policy.

<job-id>

Displays a replication report about the job with the specified ID.

isi sync target view

Displays information about a replication policy that is targeting the local cluster.

Syntax

```
isi sync target view {<policy-name> | --target-path <path>}
```

Options

<policy-name>

Displays information about the specified policy.

--target-path <path>

Displays information about the policy targeting the specified directory.

isi tape delete

Disconnects the cluster from an NDMP tape or media change device that is currently connected to a Backup Accelerator node on the cluster.

Syntax

```
isi tape delete  
  [--name <string>]  
  [--all]  
  [--force]  
  [--verbose]
```

Options

--name <string>

The name of the NDMP tape or media change device.

--all

Disconnects the cluster from all devices.

{--force | -f}

Skips the confirmation prompt.

{--verbose | -v}

Displays more detailed information.

Example

The following command disconnects `tape001` from the cluster:

```
isi tape delete tape001
```

isi tape list

Displays a list of NDMP devices that are currently connected to the cluster.

Syntax

```
isi tape list
  [--node <lnn>]
  [--tape]
  [--activepath]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--node <lnn>

Displays only devices that are attached to the node of the specified logical node number (LNN).

--tape

Displays only tape devices.

--activepath

Displays only the active paths of a device.

--format {table | json | csv | list}

Displays devices in table, JSON, CSV, or list format.

{--no-header | -a}

Does not display headers in table or CSV format.

{--no-footer | -z}

Does not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

Examples

To view a list of all NDMP devices, run the following command:

```
isi tape list
```

isi tape modify

Modifies the name or state of a tape or media changer device.

Syntax

```
isi tape modify --name <name>
  [--new-name <string>]
  [--close-device]
```

Options

--name <name>

The current device name.

--new-name <string>

The new device name.

--close-device

Forces the device state to closed if the device is currently open. If an NDMP session unexpectedly stops, a tape or media changer device may be left in an open state, which prevents the device from being opened again.

isi tape rename

Renames an NDMP device that is currently connected to a Backup Accelerator node on the cluster.

Syntax

```
isi tape rename <devname> <rename>
```

Options

<devname>

Modifies the name of the specified NDMP device.

<rename>

Specifies a new name for the given NDMP device.

Examples

The following example renames tape003 to tape005:

```
isi tape rename tape003 tape005
```

isi tape rescan

Scans Fibre Channel ports for undetected NDMP backup devices that are attached to Backup Accelerator nodes. If the scan reveals new devices, the cluster creates entries for the new devices.

Syntax

```
isi tape rescan  
  [--node <lnn>]  
  [--port <integer>]  
  [--reconcile]
```

Options

If no options are specified, scans all nodes and ports.

--node <lnn>

Scans only the node of the specified logical node number (LNN).

--port *<integer>*

Scans only the specified port. If you specify `--node`, scans only the specified port on the specified node. If you do not specify `--node`, scans the specified port on all nodes.

--reconcile

Removes entries for devices or paths that have become inaccessible.

Example

To scan the entire cluster for NDMP devices, and remove entries for devices and paths that have become inaccessible, run the following command:

```
isi tape rescan --reconcile
```

isi tape view

Displays information about a tape or media changer device.

Syntax

```
isi tape view --name <name>  
  [--activepath]  
  [--format {list | json}]
```

Options

<name>

The name of the tape or media changer device.

--activepath

Displays only the active paths of the device.

--format {list | json}

Displays devices in list or JSON format.

isi upgrade cluster add-nodes

Add new nodes to a running upgrade process.

Syntax

```
isi upgrade cluster add-nodes <nodes>  
  [--yes]
```

Options

<nodes>

List of comma-separated (1,3,7) or dash-separated (1-7) specified logical node numbers (LNNs) to mark for upgrade.

--yes

Automatically answer yes at the prompt.

isi upgrade cluster add-remaining-nodes

Let the system include any remaining or new nodes inside an existing upgrade.

Syntax

```
isi upgrade cluster add-remaining-nodes  
[--yes]
```

Options

--yes

Automatically answer yes at the prompt.

isi upgrade cluster archive

Start an archive of the upgrade framework.

Syntax

```
isi upgrade cluster archive  
[--clear]
```

Options

--clear

Clear the upgrade after an archive is complete.

isi upgrade cluster assess

Runs upgrade checks without starting an upgrade.

Syntax

```
isi upgrade cluster assess <install-image-path>
```

Options

<install-image-path>

The path of the upgrade install image. Must be within an `/ifs` or `http://` source.

isi upgrade cluster commit

Commits the upgrade to the new version. Rollback is not possible after you run this command.

Syntax

```
isi upgrade cluster commit  
[--yes]
```

Options

--yes

Automatically answers yes at the upgrade commitment prompt.

isi upgrade cluster firmware

This is the command-line interface for firmware upgrades.

Syntax


```
isi upgrade cluster firmware <action>
[--timeout <integer>]
```

Options

<action>

Specifies actions you can take against the firmware upgrade.

package	Lists all the nodes on the cluster and shows detailed firmware package information on the given node.
devices	Lists all the nodes on the cluster and shows detailed status of the current firmware for each node.
assess	Runs upgrade checks without starting a firmware upgrade.
view	Shows overview status of the current firmware upgrade activity.
start	Starts upgrade processes.

 **NOTE: All upgrade processes take a long time to run. The return status of a command only relates to the issuing of the command itself, not the successful completion of it.**

--timeout <integer>

Number of seconds for a command timeout.

Example

The following command runs upgrade checks without starting the firmware upgrade.

```
isi upgrade cluster firmware assess
```

isi upgrade cluster from-version

Displays the version of the cluster you are upgrading from.

Syntax

```
isi upgrade cluster from-version
```

Example

To view information about the cluster version you are upgrading from, run the following command:

```
isi upgrade cluster from-version
```

The system displays output similar to the following example:

```
Upgrading Current OS Version: 7.2.1.1
                               Major: 7
                               Maintenance: 0
                               Minor: 0
                               Bugfix: 0
```

isi upgrade cluster nodes firmware

This is the command-line interface for the non-disruptive upgrade firmware upgrade framework.

Syntax

```
isi upgrade cluster nodes firmware <action>
[--timeout <integer>]
```

Options

<action>

Specifies reporting actions you can take regarding node firmware updates.

devices	Reports devices on the nodes which are supported in the installed firmware package.
package	Reports the contents of the installed firmware package.
progress	Reports, in list or view format, status information regarding the firmware upgrade.

--timeout <integer>

Number of seconds for a command timeout.

Example

The following command displays the contents of the installed firmware package:

```
isi upgrade cluster nodes firmware package
```

isi upgrade cluster nodes list

List all nodes on the cluster and show detailed status of their upgrade activity.

Syntax

```
isi upgrade cluster nodes list
```

Example

To list upgrade status for all nodes on the cluster, run the following command:

```
isi upgrade cluster nodes list
```

The system displays output similar to the following example:

```
Node LNN: 1
Node Upgrade State: committed
Error Details: None
Last Upgrade Action: -
Last Action Result: -
Node Upgrade Progress: None
Node OS Version: 8.0.0.0

Node LNN: 2
Node Upgrade State: non-responsive
Error Details: None
Last Upgrade Action: -
Last Action Result: -
Node Upgrade Progress: unknown
Node OS Version: N/A

Node LNN: 3
Node Upgrade State: committed
Error Details: None
Last Upgrade Action: -
Last Action Result: -
Node Upgrade Progress: None
Node OS Version: 8.0.0.0
```

isi upgrade cluster nodes view

Show detailed status of the current upgrade activity on a specified node.

Syntax

```
isi upgrade cluster nodes view <lnn>
```

Options

<lnn>

The logical node number (LNN) of the node for which you want to view upgrade status.

Example

To view the upgrade status for a node with the LNN 1, run the following command:

```
isi upgrade cluster nodes view 1
```

The system displays output similar to the following example:

```
Node LNN: 1
Node Upgrade State: committed
Error Details: None
Last Upgrade Action: -
Last Action Result: -
Node Upgrade Progress: None
Node OS Version: 8.0.0.0
```


isi upgrade cluster retry-last-action

Retry the last upgrade action on a node, in case the previous action failed.

Syntax

```
isi upgrade cluster retry-last-action <nodes>
```

Options

<nodes>

A list of comma-separated (1,3,7) or dash-separated (1-7) logical node numbers to select. You can also use `a11` to select all the cluster's nodes at any given time.

isi upgrade cluster rollback

Stop upgrading a cluster, and return to the previous version. This causes a disruptive rollback of the upgrade.

Syntax

```
isi upgrade cluster rollback  
[--yes]
```

Options

--yes

Automatically answer yes to the confirmation prompt.

isi upgrade cluster rolling-reboot

Perform a rolling reboot of a cluster.

Syntax

```
isi upgrade cluster rolling-reboot  
[--nodes <integer-range_list>]  
[--force]
```

Options

--nodes <integer-range_list>

List of comma-specified (1,3,7...) or dash-specified (1-7) node logical node numbers (LNNs) to select. At any given time during the reboot, you can also enter `a11` to select all the cluster nodes.

{--force | -f}

Do not prompt for confirmation of the reboot.

isi upgrade cluster settings

Show the settings of the currently running upgrade.

Syntax

```
isi upgrade cluster settings
```

Options

There are no options for this command.

isi upgrade cluster start

Start an upgrade process.

Syntax

```
isi upgrade cluster start <install-image-path>  
  [--skip-optional]  
  [--yes]  
  [--simultaneous]  
  [--nodes <integer_range_list>]
```

Options

<install-image-path>

The file path of the location of the upgrade install image. The file path must be accessible in a `/ifs` directory or by an `https://` URL.

--skip-optional

Skips the optional pre-upgrade checks.

--yes

Automatically answer yes to the confirmation prompt.

--simultaneous

Start a simultaneous upgrade.

--nodes <integer_range_list>

List of comma-separated (1,3,7) or dash-separated (1-7) logical node numbers (LNNs) to select for upgrade.

isi upgrade cluster to-version

Show the version of the cluster to which you are upgrading.

Syntax

```
isi upgrade cluster to-version
```

Options

There are no options for this command.

isi upgrade cluster view

Show status of the current upgrade activity on the cluster.

Syntax

```
isi upgrade cluster view
```

Options

There are no options for this command.

isi upgrade patches abort

Repairs the patch system by attempting to discontinue the most recent failed action.

Syntax

```
isi upgrade patches abort  
  [--force]
```

Options

{--force | -f}

Skips the confirmation prompt for this command.

isi upgrade patches install

Installs a system patch.

Syntax

```
isi upgrade patches install <patch>  
  [--rolling {yes | no}]  
  [--override]
```

Options

<patch>

The file path location of the patch to install. This must be an absolute path within the `/ifs` file system.

{--rolling | -r} {yes | no}

Performs a rolling patch install. A value of `no` will install simultaneously on all nodes.

{--override | -o}

Overrides the patch system validation, and forces the patch installation.

isi upgrade patches list

Lists all system patches.

Syntax

```
isi upgrade patches list
  [--local]
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

--local

Lists patch information only on the local node.

{--limit | -l} <integer>

The number of upgrade patches to display.

--format {table | json | csv | list}

Displays upgrade patches in table, JSON, CSV or list format.

{--no-header | -a}

Does not display headers in CSV or table formats.

{--no-footer | -z}

Does not display table summary footer information.

{--verbose | -v}

Displays more detailed information.

isi upgrade patches uninstall

Uninstalls a system patch.

Syntax

```
isi upgrade patches uninstall <patch>
  [--rolling {yes | no}]
  [--override]
  [--force]
```

Options

<patch>

The name or ID of the patch to uninstall.

{--rolling | -r} {yes | no}

Performs a rolling patch uninstall. A value of `no` will uninstall simultaneously on all nodes.

{--override | -o}

Overrides the patch system validation, and forces the patch uninstallation.

{--force | -f}

Skips the confirmation prompt.

isi upgrade patches view

Shows details of a system patch.

Syntax

```
isi upgrade patches view <patch>
[--local]
```

Options

<patch>

The name or ID of the patch to view.

--local

Shows patch information only for the local node.

isi version

Displays cluster version information.

Syntax

```
isi version
[--format {list | json}]
[--verbose
```

Options

--format {list | json}

Displays the cluster version information in list or JSON format.

{--verbose | -v}

Displays more detailed cluster version information.

isi worm cdate set

Sets the SmartLock compliance clock to the current time on the system clock.

 **CAUTION:**

You can set the compliance clock only once. After the compliance clock has been set, you cannot modify the compliance clock time.

Syntax

```
isi worm cdate set
```

Options

There are no options for this command.

isi worm cdate view

Displays whether or not the SmartLock compliance clock is set. If the compliance clock is set, displays the current time on the compliance clock.

Syntax

```
isi worm cdate view
```

Options

There are no options for this command.

isi worm create

Designates an existing directory as a WORM root directory. The **isi smartlock create** command is an alias of this command.

Syntax

```
isi worm create <path>
```

Options

<path>

Designates the specified directory as a SmartLock directory. The specified directory must be empty.

Specify as a directory path.

isi worm domains create

Creates a SmartLock directory.

Syntax

```
isi worm domains create <path>
  [--compliance]
  [--autocommit-offset <duration>]
  [--override-date <timestamp>]
  [--privileged-delete {true | false}]
  [--disable-privileged-delete]
  [--default-retention {<duration> | forever | use_min
  | use_max}]
  [--min-retention {<duration> | forever}]
  [--max-retention <duration>]
  [--mkdir]
  [--force]
  [--verbose]
```

Options

<path>

Creates a SmartLock directory at the specified path.

Specify as a directory path.

{--compliance | -C}

Specifies the SmartLock directory as a SmartLock compliance directory. This option is valid only on clusters running in SmartLock compliance mode.

{--autocommit-offset | -a} <duration>

Specifies an autocommit time period. After a file exists in a SmartLock directory without being modified for the specified length of time, the file automatically committed to a WORM state.

Specify <duration> in the following format:

```
<integer><units>
```

Specify <units> are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To specify no autocommit time period, specify `none`. The default value is `none`.

{--override-date | -o} <timestamp>

Specifies an override retention date for the directory. Files committed to a WORM state are not released from a WORM state until after the specified date, regardless of the maximum retention period for the directory or whether a user specifies an earlier date to release a file from a WORM state.

Specify <timestamp> in the following format:

```
<YYYY>-<MM>-<DD>[T<hh>:<mm>[:<ss>]]
```

{--privileged-delete | -p} {true | false}

Determines whether files in the directory can be deleted through the `isi worm files delete` command. This option is available only for SmartLock enterprise directories.

The default value is `false`.

--disable-privileged-delete

Permanently prevents WORM committed files from being deleted from the SmartLock directory.

NOTE:

If you specify this option, you can never enable the privileged delete functionality for the directory. If a file is then committed to a WORM state in the directory, you will not be able to delete the file until the retention period has passed.

{--default-retention | -d} {<duration> | forever | use_min | use_max}

Specifies a default retention period. If a user does not explicitly assign a retention period expiration date, the default retention period is assigned to the file when it is committed to a WORM state.

Specify <duration> in the following format:

```
<integer><units>
```

Specify <units> are valid:

Y	Specifies years
M	Specifies months

W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To permanently retain WORM committed files by default, specify `forever`. To assign the minimum retention period as the default retention period, specify `use_min`. To assign the maximum retention period as the default retention period, specify `use_max`.

{--min-retention | -m} {<duration> | forever}

Specifies a minimum retention period. Files are retained in a WORM state for at least the specified amount of time.

Specify <duration> in the following format:

```
<integer><units>
```

Specify <units> as one of the following values:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To permanently retain all WORM committed files, specify `forever`.

{--max-retention | -x} {<duration> | forever}

Specifies a maximum retention period. Files cannot be retained in a WORM state for more than the specified amount of time, even if a user specifies an expiration date that results in a longer retention period.

Specify <duration> in the following format:

```
<integer><units>
```

Specify <units> as one of the following values:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To specify no maximum retention period, specify `forever`.

{--mkdir | -M}

Creates the specified directory if it does not already exist.

{--force | -f}

Does not prompt you to confirm the creation of the SmartLock directory.

{--verbose | -v}

Displays more detailed information.

isi worm domains list

Displays a list of WORM directories.

Syntax

```
isi worm domains list
  [--limit <integer>]
  [--sort <attribute>]
  [--descending]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--sort <attribute>

Sorts output displayed by the specified attribute.

The following values are valid:

id	Sorts output by the SmartLock directory ID.
path	Sorts output by the path of the SmartLock directory.
type	Sorts output based on whether the SmartLock directory is a compliance directory.
lin	Sorts output by the inode number of the SmartLock directory.
autocommit_off set	Sorts output by the autocommit time period of the SmartLock directory.
override_date	Sorts output by the override retention date of the SmartLock directory.
privileged_del ete	Sorts output based on whether the privileged delete functionality is enabled for the SmartLock directory.
default_retent ion	Sorts output by the default retention period of the SmartLock directory.
min_retention	Sorts output by the minimum retention period of the SmartLock directory.
max_retention	Sorts output by the maximum retention period of the SmartLock directory.
total_modifies	Sorts output by the total number of times that the SmartLock directory has been modified.

{--descending | -d}

Displays output in reverse order.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table output without headers.

{--no-footer | -z}

Displays table output without footers. Footers display snapshot totals, such as the total amount of storage space consumed by snapshots.

{--verbose | -v}

Displays more detailed information.

isi worm domains modify

Modifies SmartLock settings of a SmartLock directory.

Syntax

```
isi worm domains modify <domain>
  [--compliance]
  [{--autocommit-offset <duration> | --clear-autocommit-offset}]
  [{--override-date <timestamp> | --clear-override-date}]
  [--privileged-delete {true | false}]
  [--disable-privileged-delete]
  [--default-retention {<duration> | forever | use_min
  | use_max} | --clear-default-retention}]
  [--min-retention {<duration> | forever} | --clear-min-retention}]
  [--max-retention | -x] (<duration> | forever) | --clear-max-retention]
  [{--exclude | -x} <string>]
  [--set-pending-delete]
  [{--force | -f}]
  [{--verbose | -v}]
  [{--help | -h}]
```

Options

<domain>

Modifies the specified SmartLock directory.

Specify as a directory path, ID, or LIN of a SmartLock directory.

{--compliance | -C}

Specifies the SmartLock directory as a SmartLock compliance directory. This option is valid only on clusters running in SmartLock compliance mode.

{--autocommit-offset | -a} <duration>

Specifies an autocommit time period. After a file exists in a SmartLock directory without being modified for the specified length of time, the file automatically committed to a WORM state.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To specify no autocommit time period, specify `none`. The default value is `none`.

`--clear-autocommit-offset`

Removes the autocommit time period for the given SmartLock directory.

`{--override-date | -o} <timestamp>`

Specifies an override retention date for the directory. Files committed to a WORM state are not released from a WORM state until after the specified date, regardless of the maximum retention period for the directory or whether a user specifies an earlier date to release a file from a WORM state.

Specify `<timestamp>` in the following format:

```
<YYYY>--<MM>--<DD>[T<hh>:<mm>[:<ss>]]
```

`--clear-override-date`

Removes the override retention date for the given SmartLock directory.

`{--privileged-delete | -p} {true | false}`

Determines whether files in the directory can be deleted through the `isi worm files delete` command. This option is available only for SmartLock enterprise directories.

The default value is `false`.

`--disable-privileged-delete`

Permanently prevents WORM committed files from being deleted from the SmartLock directory.

NOTE:

If you specify this option, you can never enable the privileged delete functionality for the SmartLock directory. If a file is then committed to a WORM state in the directory, you will not be able to delete the file until the retention period expiration date has passed.

`{--default-retention | -d} {<duration> | forever | use_min | use_max}`

Specifies a default retention period. If a user does not explicitly assign a retention period expiration date, the default retention period is assigned to the file when it is committed to a WORM state.

Specify `<duration>` in the following format:

```
<integer><units>
```

Specify `<units>` are valid:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To permanently retain WORM committed files by default, specify `forever`. To assign the minimum retention period as the default retention period, specify `use_min`. To assign the maximum retention period as the default retention period, specify `use_max`.

`--clear-default-retention`

Removes the default retention period for the given SmartLock directory.

`{--min-retention | -m} {<duration> | forever}`

Specifies a minimum retention period. Files are retained in a WORM state for at least the specified amount of time.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* as one of the following values:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To permanently retain all WORM committed files, specify *forever*.

--clear-min-retention

Removes the minimum retention period for the given SmartLock directory.

{--max-retention | -x} {<duration> | forever}

Specifies a maximum retention period. Files cannot be retained in a WORM state for more than the specified amount of time, even if a user specifies an expiration date that results in a longer retention period.

Specify *<duration>* in the following format:

```
<integer><units>
```

Specify *<units>* as one of the following values:

Y	Specifies years
M	Specifies months
W	Specifies weeks
D	Specifies days
H	Specifies hours
m	Specifies minutes
s	Specifies seconds

To specify no maximum retention period, specify *forever*.

--clear-max-retention

Removes the maximum retention period for the given SmartLock directory.

{--exclude | -x} <string>

Path excluded from the WORM domain. Specify *--exclude* for each additional path.

--set-pending-delete

Mark a compliance domain for deletion. This action is irreversible.

{--force | -f}

Does not prompt you to confirm the creation of the SmartLock directory.

{--verbose | -v}

Displays more detailed information.

isi worm domains view

Displays WORM information about a specific directory or file.

Syntax

```
isi worm domains view <domain>
```

Options

<domain>

Displays information about the specified SmartLock directory.

Specify as a directory path, ID, or LIN of a SmartLock directory.

isi worm files delete

Deletes a file committed to a WORM state. This command can be run only by the root user or compliance administrator.

Syntax

```
isi worm files delete <path>
  [--force]
  [--verbose]
```

Options

<path>

Deletes the specified file. The file must exist in a SmartLock enterprise directory with the privileged delete functionality enabled.

Specify as a file path.

--force

Does not prompt you to confirm that you want to delete the file.

--verbose

Displays more detailed information.

isi worm files view

Displays information about a file committed to a WORM state.

Syntax

```
isi worm files view <path>
  [--no-symlinks]
```

Options

<path>

Displays information about the specified file. The file must be committed to a WORM state.

Specify as a file path.

--no-symlinks

If <path> refers to a file, and the given file is a symbolic link, displays WORM information about the symbolic link. If this option is not specified, and the file is a symbolic link, displays WORM information about the file that the symbolic link refers to.

isi zone restrictions create

Prohibits user or group access to the /ifs directory. Attempts to read or write files by restricted users or groups return ACCESS DENIED errors.

Syntax

```
isi zone restrictions create <zone> {<user> | --uid <integer>
| --group <string> | --gid <integer> | --sid <string>
| --wellknown <string>}
[--verbose]
```

Options

<zone>

Specifies an access zone by name.

<user>

Specifies a user by name.

--uid <integer>

Specifies a user by UID.

--group <string>

Specifies a group by name.

--gid <integer>

Specifies a group by GID.

--sid <string>

Specifies an object by user or group SID.

--wellknown <name>

Specifies a well-known user, group, machine, or account name.

{--verbose | -v}

Returns a success or fail message after running the command.

isi zone restrictions delete

Removes a restriction that prohibits user or group access to the /ifs directory.

Syntax

```
isi zone restrictions delete <zone> {<user> | --uid <integer>
| --group <string> | --gid <integer> | --sid <string>
| --wellknown <string>}
[--force]
[--verbose]
```

Options

- <zone>**
Specifies an access zone by name.
- <user>**
Specifies a user by name.
- uid <integer>**
Specifies a user by UID.
- group <string>**
Specifies a group by name.
- gid <integer>**
Specifies a group by GID.
- sid <string>**
Specifies an object by user or group SID.
- wellknown <string>**
Specifies an object by well-known SID.
- {--force | -f}**
Suppresses command-line prompts and messages.
- {--verbose | -v}**
Returns a success or fail message after running the command.

isi zone restrictions list

Displays a list of users or groups that are prohibited from accessing the `/ifs` directory.

Syntax

```
isi zone restrictions list <zone>  
  [--limit <integer>]  
  [--format {table | json | csv | list}]  
  [--no-header]  
  [--no-footer]  
  [--verbose]
```

Options

- <zone>**
Specifies an access zone by name.
- {--limit | -l} <integer>**
Displays no more than the specified number of items.
- format {table | json | csv | list}**
Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.
- {--no-header | -a}**
Displays table and CSV output without headers.
- {--no-footer | -z}**
Displays table output without footers.
- {--verbose | -v}**
Displays more detailed information.

Examples

To display a list of restricted users for the built-in System zone, run the following command:

```
isi zone restrictions list system
```

isi zone zones create

Creates an access zone.

Syntax

```
isi zone zones create <name> <path>
  [--map-untrusted <workgroup>]
  [--auth-providers <provider-type>:<provider-name>]
  [--netbios-name <string>]
  [--user-mapping-rules <string>]
  [--home-directory-umask <integer>]
  [--skeleton-directory <string>]
  [--cache-entry-expiry <duration>]
  [--create-path]
  [--force-overlap]
  [--groupnet <groupnet>]
  [--verbose]
```

Options

<name>

Specifies the name of the access zone.

<path>

Specifies the base directory path for the zone.

--map-untrusted <workgroup>

Maps untrusted domains to the specified NetBIOS workgroup during authentication.

--auth-providers <provider-type>:<provider-name>

Specifies one or more authentication providers, separated by commas, for authentication to the access zone. Authentication providers are checked in the order specified. You must specify the name of the authentication provider in the following format: <provider-type>:<provider-name>.

--netbios-name <string>

Specifies the NetBIOS name.

--user-mapping-rules <string>

Specifies one or more user mapping rules, separated by commas, for the access zone.

--home-directory-umask <integer>

Specifies the permissions to set on auto-created user home directories.

--skeleton-directory <string>

Sets the skeleton directory for user home directories.

--cache-entry-expiry <duration>

Specifies duration of time to cache a user/group.

--create-path

Specifies that the value entered as the access zone path is to be created if it does not already exist.

--force-overlap

Allows the base directory to overlap with the base directory of another access zone.

--groupnet <string>

Specifies the groupnet referenced by the access zone.

{--verbose | -v}

Displays the results of running the command.

isi zone zones delete

Deletes an access zone. All authentication providers that are associated with the access zone remain available to other zones, but IP addresses are not reassigned. You cannot delete the built-in System zone.

Syntax

```
isi zone zones delete <zone>
  [--force]
  [--verbose]
```

Options

<zone>

Specifies the name of the access zone to delete.

{--force | -f}

Suppresses command-line prompts and messages.

{--verbose | -v}

Displays the results of running the command.

isi zone zones list

Displays a list of access zones in the cluster.

Syntax

```
isi zone zones list
  [--limit <integer>]
  [--format {table | json | csv | list}]
  [--no-header]
  [--no-footer]
  [--verbose]
```

Options

{--limit | -l} <integer>

Displays no more than the specified number of items.

--format {table | json | csv | list}

Displays output in table (default), JavaScript Object Notation (JSON), comma-separated value (CSV), or list format.

{--no-header | -a}

Displays table and CSV output without headers.

{--no-footer | -z}

Displays table output without footers.

{--verbose | -v}

Displays more detailed information.

Examples

To view a list of all access zones in the cluster, run the following command:

```
isi zone zones list
```

isi zone zones modify

Modifies an access zone.

Syntax

```
isi zone zones modify <zone>
  [--name <string>]
  [--path <path>]
  [--map-untrusted <string>]
  [--auth-providers <provider-type>:<provider-name>]
  [--clear-auth-providers]
  [--add-auth-providers <provider-type>:<provider-name>]
  [--remove-auth-providers <provider-type>:<provider-name>]
  [--netbios-name <string>]
  [--user-mapping-rules <string>]
  [--clear-user-mapping-rules]
  [--add-user-mapping-rules <string>]
  [--remove-user-mapping-rules <string>]
  [--home-directory-umask <integer>]
  [--skeleton-directory <string>]
  [--cache-entry-expiry <duration>]
  [--revert-cache-entry-expiry]
  [--create-path]
  [--force-overlap]
  [--verbose]
```

Options

<zone>

Specifies the name of the access zone to modify.

--name <string>

Specifies a new name for the access zone. You cannot change the name of the built-in System access zone.

--path <path>

Specifies the base directory path for the zone.

--map-untrusted <string>

Specifies the NetBIOS workgroup to map untrusted domains to during authentication.

--auth-providers <provider-type>:<provider-name>

Specifies one or more authentication providers, separated by commas, for authentication to the access zone. This option overwrites any existing entries in the authentication providers list. To add or remove providers without affecting the current entries, configure settings for `--add-auth-providers` or `--remove-auth-providers`.

--clear-auth-providers

Removes all authentication providers from the access zone.

--add-auth-providers <provider-type>:<provider-name>

Adds one or more authentication providers, separated by commas, to the access zone.

--remove-auth-providers <provider-type>:<provider-name>

Removes one or more authentication providers, separated by commas, from the access zone.

--netbios-name <string>

Specifies the NetBIOS name.

--user-mapping-rules <string>

Specifies one or more user mapping rules, separated by commas, for the access zone. This option overwrites all entries in the user mapping rules list. To add or remove mapping rules without overwriting the current entries, configure settings with `--add-user-mapping-rules` or `--remove-user-mapping-rules`.

--clear-user-mapping-rules

Removes all user mapping rules from the access zone.

--add-user-mapping-rules <string>

Adds one or more user mapping rules, separated by commas, to the access zone.

--remove-user-mapping-rules <string>

Removes one or more user mapping rules, separated by commas, from the access zone.

--home-directory-umask <integer>

Specifies the permissions to set on auto-created user home directories.

--skeleton-directory <string>

Sets the skeleton directory for user home directories.

--cache-entry-expiry <duration>

Specifies duration of time to cache a user/group.

--cache-entry-expiry

Sets the value of `--cache-entry-expiry` to the system default.

--create-path

Specifies that the zone path is to be created if it doesn't already exist.

--force-overlap

Allows the base directory to overlap with the base directory of another access zone.

{--verbose | -v}

Displays the results of running the command.

isi zone zones view

Displays the properties of an access zone.

Syntax

```
isi zone zones view <zone>
```

Options

<zone>

Specifies the name of the access zone to view.