

Dell EMC Data Protection Advisor

Version 18.1

Installation and Administration Guide

302-004-935

Rev 11

February 2020

Copyright © 2005-2020 Dell Inc. or its subsidiaries. All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Figures		7
Tables		9
Preface		11
Chapter 1	Preparing to install DPA	15
	Overview.....	16
	System requirements.....	16
	DPA Server platforms.....	16
	Datastore storage.....	17
	Permissions.....	17
	NTP time synchronization.....	17
	Installation considerations.....	17
	Configuring virtual infrastructure memory and CPU.....	18
	OS resource optimization.....	18
	Communications settings in DPA.....	18
	DPA port settings	20
	Installation and configuration overview.....	23
Chapter 2	Installing DPA	29
	DPA server installation.....	30
	Installing the Datastore Service.....	30
	Installing the Application Service.....	32
	Application clustering.....	35
	Datastore Replication.....	45
	DPA Agent installation.....	51
	Installing the DPA Agent.....	51
	Setting DPA Agent registration password.....	53
	Configure DPA version 18.1 Agent to go back and collect backup application data.....	53
	Configure DPA to show all VM Image Backups in Avamar	55
	Installing by using command line installation.....	55
	DPA postinstallation steps.....	60
	Encryption of the DPA Application server.....	63
	Encrypting Application Server cluster.....	63
	Configuring antivirus software with DPA.....	64
	Upgrades.....	64
	Upgrade prerequisites.....	65
	Upgrading DPA.....	66
	Upgrading DPA Agents	67
	Upgrading DPA Agents previous to version 6.5 alongside DPA version 6.5 Agents and version 6.5 Server.....	67
	Upgrading DPA with a LINUX version running glibc earlier than 2.12 ...	67
	Upgrading existing clusters.....	68

	Upgrading with Datastore Replication enabled with DPA 6.3 and later....	68
	Upgrading with Datastore Replication enabled with DPA versions earlier than 6.3.....	69
	Upgrading with Datastore Replication and existing clusters.....	70
Chapter 3	Administering DPA	71
	License management.....	72
	Evaluation license bundled with DPA.....	72
	Licensing types in DPA.....	72
	CLP and WLS license coexistence in DPA.....	72
	Expired licenses.....	72
	License removal.....	73
	Adding new licenses.....	73
	Disabling automatic temporary licence expiration pop-up.....	73
	Users and security.....	73
	User accounts.....	73
	User roles and privileges.....	76
	External authentication, LDAP integration, and binding.....	79
	Automated user provisioning.....	81
	System settings.....	84
	Configuring backup and restore resolution fields.....	84
	Viewing and editing settings.....	85
	System Settings.....	86
	Agentless Discovery.....	89
	Server data deletion.....	89
	Configuring Data Deletion Schedule.....	90
	Root Cause Analysis Settings.....	90
	Gathering historical backup data using DPA web console.....	91
	Generate Support Bundle.....	91
	Digital certificate.....	92
	Time periods.....	92
	Time zones in DPA.....	92
	Automatic report prioritization.....	93
	Schedules.....	94
	Manage Data Collection Defaults.....	94
	Manage Sites.....	114
	Application service administration.....	115
	Running Linux DPA Application as non-root user.....	115
	Setting TLS protocol version 1.2 only after installation or upgrade.....	115
	Customization of service information.....	115
	Clustering administration.....	119
	Datastore service administration.....	121
	Backup of the Datastore.....	121
	Datastore Replication administration.....	122
	DPA Database superuser password.....	126
	DPA command line operations.....	127
	Sourcing the DPA config file for UNIX users.....	127
	dpa CLI command.....	127
	dpa agent commands.....	128
	dpa application commands.....	130
	dpa datastore commands.....	138
	dpa service commands.....	145
	Loading historical backup job data.....	147

Chapter 4	Environment discovery in DPA	151
	Configuring the environment for discovery.....	152
	Discovery overview.....	152
	Defining objects to be monitored.....	153
	Before you run the Discovery Wizard.....	154
	Monitoring of backup applications.....	156
	Monitoring of Databases.....	169
	Monitoring of applications using cloud-based solutions.....	180
	Monitoring of hosts.....	181
	Monitoring of primary storage.....	186
	Monitoring of protection storage.....	188
	Monitoring of switches and I/O devices.....	192
	Virtualization management.....	194
	Monitoring of clusters.....	195
	Discovering a host or object manually.....	196
	About job data gathering after discovery.....	197
	Monitored objects and groups.....	198
	Objects overview.....	198
	Groups.....	199
	Object attributes.....	200
	Smart Groups.....	200
	Gathering historical backup data using DPA web console.....	203
	Configuring policies, rules, and alerts.....	203
	Policies and alerts overview.....	203
	Policies.....	203
	Policies and generating events.....	230
	Parameters for generating alerts from scripts.....	231
	Rule Template.....	232
	Policy application.....	233
Chapter 5	Uninstalling DPA	235
	Uninstalling the software.....	236
	Uninstalling by using silent command line.....	236
	Uninstalling through user interface on Windows.....	236
	Agent-only uninstallation.....	236
Chapter 6	Troubleshooting	237
	Installation troubleshooting.....	238
	DPA Agent does not restart or register after DPA Server password change.....	238
	DPA Datastore on Linux failure to start after installation.....	238
	DPA web console launch failure on Windows Server 2012.....	238
	Postinstallation memory adjustment.....	238
	Error messages during upgrades.....	239
	Log files.....	239
	Changing default log detail level.....	239
	Viewing install log file.....	239
	Viewing server log files.....	239
	Server log files.....	240
	Viewing agent log files.....	240
	Managing log files.....	240
	Enabling alternative log rotation on VMs running Windows.....	240
	Erroneous memory data in installer log file.....	240

Running a DPA Agent request in debug mode using DPA web console....	240
Default modtest deletion schedule.....	241
Generate Support Bundle.....	241
Data collection troubleshooting.....	242
Troubleshooting data collection: first actions.....	242
Troubleshooting data collection: second actions.....	242
Preparing a log file for submission to EMC Support.....	242
Client/Storage discovery for replication analysis troubleshooting.....	243
Client/storage discovery using remote execution.....	243
Client/storage discovery with agent.....	245
General client/storage discovery.....	245
Incorrect Recovery Point times Synchronization.....	247
Troubleshooting report output failure.....	248
Troubleshooting report generation or publishing problems.....	248
System clock synchronization.....	249

FIGURES

1	DPA ports and protocols.....	19
2	DPA installation workflow.....	23
3	Relationship between DPA Application nodes and DPA Agents monitoring applications..	152
4	Object library Multilevel Smart Group configuration example.....	202

TABLES

1	Revision history.....	11
2	Style conventions.....	13
3	DPA Application ports settings.....	20
4	DPA Datastore port settings.....	21
5	DPA Agent port settings.....	21
6	DPA cluster port settings.....	21
7	Installation and configuration overview	23
8	Installer command line options.....	56
9	Datastore installer variables.....	57
10	Datastore Advanced options Replication variables.....	57
11	Datastore Agent variables.....	58
12	Application installer variables.....	58
13	Application server Agent variables.....	59
14	Application server Cluster Advanced option variables.....	59
15	Standalone Agent Installer variables.....	60
16	Password Policy.....	74
17	Password History Policy.....	75
18	Login Limit.....	75
19	Password Expiration.....	76
20	User roles.....	76
21	LDAP Authentication configuration in DPA.....	79
22	Open LDAP server settings.....	82
23	Data Collection Agent settings	86
24	Server settings.....	86
25	SharePoint settings.....	88
26	Replication Analysis settings.....	88
27	Agentless Discovery settings.....	89
28	Default collected data retention periods.....	90
29	Default system-generated data retention periods.....	90
30	Data collection request options by module.....	95
31	VTL templates.....	116
32	Command and option abbreviations	128
33	Data monitoring setup summary	153
34	Connectivity details for configuring data collection through the Discovery Wizard	154
35	HP Data Protector 6.1 patch IDs.....	161
36	System monitoring modules.....	181
37	Multilevel Smart Group example.....	201
38	Capacity planning.....	214
39	Change management.....	215
40	Configuration.....	216
41	Data protection.....	217
42	Licensing.....	218
43	Performance.....	218
44	Provisioning.....	219
45	Recoverability.....	219
46	Resource utilization.....	223
47	Service Level Agreement.....	225
48	Status.....	225
49	Troubleshooting.....	227
50	Recoverability checks	229
51	Script field parameters.....	231
52	Script alert arguments.....	232
53	Client/storage discovery problems and solutions	243


Tables

54	Client/storage discovery with agent problems and solutions.....	245
55	General client/storage discovery problems and solutions	245

Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

 **Note:** This document was accurate at publication time. Go to EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

Purpose

This document provides information on how to install DPA and set up DPA to monitor a data protection environment. This document also describes administrative functions such as creating users and roles, updating system settings, creating policies, and troubleshooting data collection.

ISO 9001 certification

The management system governing the design and development of this product is ISO 9001:2015 certified.

Audience

This document is intended for system administrators. Readers of this document must be familiar with the following tasks:

- Identifying the different hardware and software components that make up the backup and replication environment.
- Following procedures to configure backup and replication operations.
- Following guidelines to locate problems and implement solutions.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	July 6, 2018	First release of this document for DPA 18.1
02	July 13, 2018	Updates to Upgrading DPA Agents on page 67
03	August 10, 2018	Updates to Data collection request options by module on page 95
04	September 7, 2018	Updates to Data collection request options by module on page 95
05	September 21, 2018	Updates to the following sections: <ul style="list-style-type: none">• Before starting the Discovery Wizard for monitoring Oracle on page 174• Before starting the Discovery Wizard for monitoring RMAN on page 175

Table 1 Revision history (continued)

Revision	Date	Description
06	October 1, 2018	Updates to the following sections: <ul style="list-style-type: none"> • Data collection request options by module on page 95 • Enable support of TLS 1.2 only on page 172
07	December 21, 2018	Addition of Configure DPA to show all VM Image Backups in Avamar on page 55
08	February 21, 2019	Updates to the following sections: <ul style="list-style-type: none"> • Upgrading DPA on page 66 • Running Linux DPA Application as non-root user on page 115
09	October 16, 2019	Updates to the following section: <ul style="list-style-type: none"> • System requirements on page 16
10	January 31, 2020	Updates to the following section: <ul style="list-style-type: none"> • System requirements on page 16
11	February 28, 2020	Updates to the "Installing DPA" chapter.

Related documentation

The DPA documentation set includes the following publications:

- *Data Protection Advisor Custom Reporting Guide*
- *Data Protection Advisor Data Collection Reference Guide*
- *Data Protection Advisor Installation and Administration Guide*
- *Data Protection Advisor Migrator Technical Notes*
- *Data Protection Advisor online help system*
- *Data Protection Advisor Product Guide*
- *Data Protection Advisor Release Notes*
- *Data Protection Advisor Report Reference Guide*
- *Programmers' Guide to Using Data Protection Advisor REST API*
- *Data Protection Advisor Security Configuration Guide*
- *Data Protection Advisor Software Compatibility Guide*
- *Other Technical Notes/White Papers*

Special notice conventions used in this document

EMC uses the following conventions for special notices:

 **NOTICE** Addresses practices not related to personal injury.


 **Note:** Presents information that is important, but not hazard-related.

Table 2 Style conventions

Bold	Used for names of interface elements, such as names of buttons, fields, tab names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications that are referenced in text
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate non-essential information that is omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at <https://support.emc.com>.

Technical support

Go to EMC Online Support at <https://support.emc.com>, and click **Service Center**. Several options for contacting EMC Technical Support appear on the site. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Online communities

Go to the EMC Community Network at <https://community.emc.com> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all EMC products.

Your comments

Your suggestions help to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to DPAD.Doc.Feedback@emc.com.

CHAPTER 1

Preparing to install DPA

This chapter includes the following sections:

- [Overview](#)16
- [System requirements](#)..... 16
- [Installation considerations](#)..... 17
- [Communications settings in DPA](#)..... 18
- [DPA port settings](#)20
- [Installation and configuration overview](#).....23

Overview

All of DPA deployments include the following installations:

- DPA Datastore server and a DPA agent on one host
- DPA Application server and a DPA agent on another host

When you install DPA, the installation wizard takes you step by step through placement of these components.

Installing the Application and Datastore servers on a single host is not supported. You can connect multiple Application servers to the same Datastore server, where each additional Application server is on its own host and the Application servers are installed as a DPA cluster. You can install additional DPAAgents for system monitoring and remote data collection. DPA supports Datastore Replication to enable continuous, safe, and reliable replication so that DPA can maintain a replica copy, or Slave, of the primary Datastore, or Master, for resilience against a single point of failure.

System requirements

DPA has the following basic minimum system requirements. The *Data Protection Advisor Software Compatibility Guide* provides a comprehensive list of system requirements.

DPA Server platforms

DPA servers support for 64-bit operating systems only. Work with your Account Representative to determine appropriate sizing for your environment.

Memory requirements

- 16 GB RAM/4 cores for the DPA Datastore server
- 16 GB RAM/4 cores for the DPA Application Server

Hard Disk Drive requirements:

- 18 GB of locally attached disk storage for the Application server
- 20 GB of locally attached disk storage for the Datastore Server
- 5 GB of free space is required for database upgrade in the DPA installation directory
 - ① **Note:** Co-located Application and Datastore systems are not supported in production systems. Although the installer provides a co-located system option, when it is selected, a dialogue stating that it is not supported in a production systems displays.
- 5 GB of free space is required in the system temp directory for Server or DataStore installation
- 5 GB of free space is required in the system temp directory for Agent installation
 - ① **Note:** DPA installer uses system temp directory (OS temp directory) to unpack the installation package. Conventional system temp folder location might vary depending on the OS used. Most OS allows to change temp directory with system environment variables. To find out where the temp folder is located in your environment and how to change it, refer to the OS documentation.
- The DPA Application server and DPA Datastore servers must not be used to run other applications. The DPA Application server host and DPA Datastore server host resources must be dedicated to DPA.
- If you are running DPA in a virtualized environment the allocated CPU and memory must be reserved for the DPA servers

- The DPA installer has a soft threshold of 7892 MB and a hard threshold of 5844 MB. The soft threshold allows the installation to continue, but the hard threshold does not.
- Automatic sizing and tuning of internal DPA resource usage takes place during installation. If resources (CPU, Memory) are taken away from the installation by other applications performance of DPA could be adversely affected.
- Storage size requirement for Datastore is highly dependent on the expected data collection and retention rates.
- DPA installations require Transport Layer Security (TLS) v 1.2 or later.
- Operating systems:
The *Data Protection Advisor Software Compatibility Guide* provides information on the supported operating systems.

Datastore storage

For performance reasons, the installation of the DPA Datastore server on NAS-based file systems, such as CIFS or NFS shares is not recommended because these file systems might not have the bandwidth to manage the required I/O.

Although the standard datastore file system layout is adequate for most deployments, you can distribute different file systems across different file systems to optimize performance during installation under Advanced installation options.

Permissions

Ensure that you have the following permissions before you install the software to avoid installation failure:

- Windows:
 - Administrator privileges (domain or local with full access)
 - If User Account Control (UAC) is enabled, use Run As Administrator
- UNIX / Linux:
 - Root user
 - If using security software to manage access to the root account, ensure the permissions allow the creation of new users after you become root. This must include the ability to create default home directories for the account to be created.

NTP time synchronization

It is a best practice to have Network Time Protocol (NTP) available to synchronize the DPA Server and the DPA Agent hosts. This ensures accurate and consistent data collection.

The DPA User Authentication process requires that the times on the system clock on the client machine and on the server be synchronized within one minute of one another.

Installation considerations

The DPA installation wizard presents advanced options for configuring Datastore Replication with Master and Slave Datastores, and for configuring clustered Application objects. If using either or both of these options, ensure that you:

- Plan the final deployment topology before beginning installation.
- Have all hosts and IP addresses predetermined and available.

If you are planning an advanced installation, contact your Account Representative for help with advanced architecture solution design.

Configuring virtual infrastructure memory and CPU

If you plan to deploy DPA in a virtualized infrastructure, perform the following steps:

Procedure

- Ensure that the memory allocated is reserved exclusively for each VM.
- Place the DPA Application and Datastore VMs in a resource pool where the resource allocation shares are set to High. Alternatively, select High Share Allocation for each individual VM.
- Select Thick Provision Eager Zeroed for Datastore disks. Thick Provision Eager Zeroed disk allocation causes all space to be allocated upfront, and the full disk file is zeroed before the system is made available for usage.

OS resource optimization

General tuning

During installation, the installer tunes the DPA Datastore Service for the host environment on which it is being deployed. This tuning assumes that the host is dedicated to DPA and takes into account resources such as Disk Space, Total Memory, and CPU cores. If during the lifetime of the DPA Datastore Service any of these physical resources are increased or decreased, execute the `dpa datastore tune` command on the Datastore host. [dpa datastore tune](#) on page 145 provides more information.

Hardware issues with tuning

For deployments where optimal performance is a concern, the type and quality of the hardware you use for your Datastore host server drastically impacts the performance of the Datastore Service.

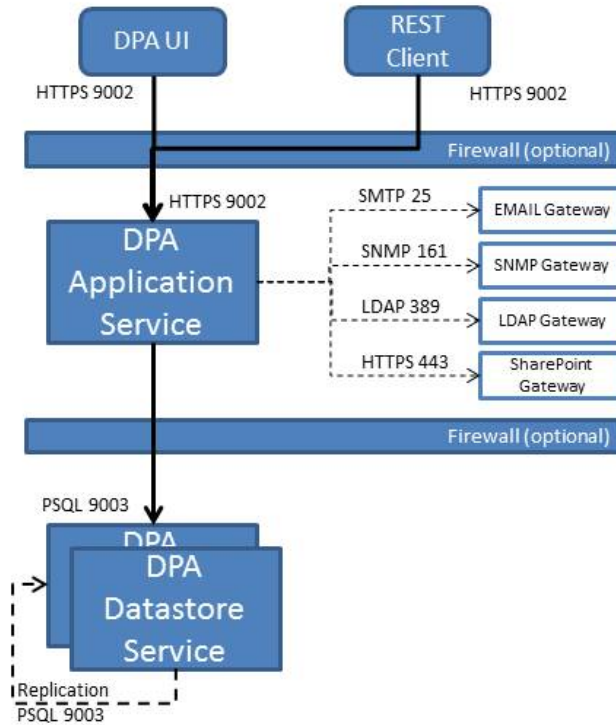
Usually, the performance is better when you have more RAM and disk spindles in your system. This is because with the extra RAM you will access your disks less. And the extra spindles help spread the reads and writes over multiple disks to increase throughput and to reduce drive head congestion.

For production purposes the DPA Application Service and the DPA Datastore Service should be placed onto different hardware. Not only does this provide more hardware dedicated to the Datastore Service, but the operating system's disk cache will contain more Datastore data and not any other application or system data.

Communications settings in DPA

To ensure communication between the DPA Server and DPA Agents, configure the firewalls in the network to allow communication on these ports, as shown in the following figure. Additional firewall configuration can be required for other ports depending on what you plan to monitor. For example, if you monitor Avamar, open port 5555 between the Avamar server and the DPA Agent. "Environment discovery in DPA" provides more information.

Figure 1 DPA ports and protocols



Note: *Application servers and Collectors can be one or many.

In the graphic above, the arrows show the initiation direction. The DPA Agent initiates connection to DPA Application Server on 9002. For firewalls, it is based upon who initiates the connection and on what port, and who is listening on the other side. DPA Agent to DPA Application Server communication is on 9002 and 3741 TCP. The communications are secure, encrypted, and compressed between the Agent and DPA server.

The following tables detail the additional ports required on deployment hosts to allow DPA to function correctly. The ports listed must be able to accept connections and allow responses back on any established connection. Some network vendors describe such handshaking communication as Bi-Directional; and such network security devices should reflect this accordingly.

DPA port settings

The following tables provide the ports needed by DPA to function correctly. Additional ports can be required for the DPA Agents depending on the systems being monitored. The *Data Protection Advisor Installation and Administration Guide* provides information on installation requirements.

Table 3 DPA Application ports settings

Port	Description	Traffic direction
25	TCP port used for the SMTP service	Outbound connection to SMTP server.
80	TCP port used for the SharePoint service	Outbound connection to SharePoint server.
161	UDP port used for SNMP service	Outbound connection to SNMP devices.
389/636 (over SSL)	TCP port used for LDAP integration	Outbound connection to LDAP server.
3741	TCP port used for DPA Agents communications.	Outbound connection to DPA agents
4447	TCP port used for intra-service communication	Inbound connection
4712	TCP port used for intra-service communication	Localhost connection
4713	TCP port used for intra-service communication	Localhost connection
5445	TCP port used for intra-service communication	Localhost connection
5455	TCP port used for intra-service communication	Localhost connection
8090	TCP port used for intra-service communication	Localhost connection
9002	TCP port used for the HTTPS service.	Inbound connection over SSL from UI, CLI and REST API clients.
9003	TCP port used for DPA Datastore communications.	Outbound connection to DPA Datastore.
9005	TCP port used for Jboss Management	Localhost connection
9999	TCP port used for Jboss Management	Localhost connection

Table 4 DPA Datastore port settings

Port	Description	Traffic direction
3741	TCP port used for DPA Agents communications.	Inbound connection from DPA Application server.
9002	TCP port used for the HTTPS service.	Outbound connection over SSL to DPA Application server.
9003	TCP port used for DPA Datastore communications.	Inbound connection from DPA Application server.

Table 5 DPA Agent port settings

Port	Description	Traffic direction
3741	TCP port used for DPA Agents communications.	Inbound connection from DPA Application server.
9002	TCP port used for the HTTPS service.	Outbound connection over SSL to DPA Application server.

Table 6 DPA cluster port settings

Port	Description	Traffic direction
25	TCP port used for the SMTP service	Outbound connection to SMTP server.
80	TCP port used for the SharePoint service	Outbound connection to SharePoint server.
161	UDP port used for SNMP service	Outbound connection to SNMP devices.
389/636 (over SSL)	TCP port used for LDAP integration	Outbound connection to LDAP server.
3741	TCP port used for DPA Agents communications.	Outbound connection to DPA agents
4447	TCP port used for intra-service communication	Inbound connection
4712	TCP port used for intra-service communication	Localhost connection
4713	TCP port used for intra-service communication	Localhost connection
5445	TCP port used for intra-service communication	Bidirectional connection for Cluster
5455	TCP port used for intra-service communication	Bidirectional connection for Cluster

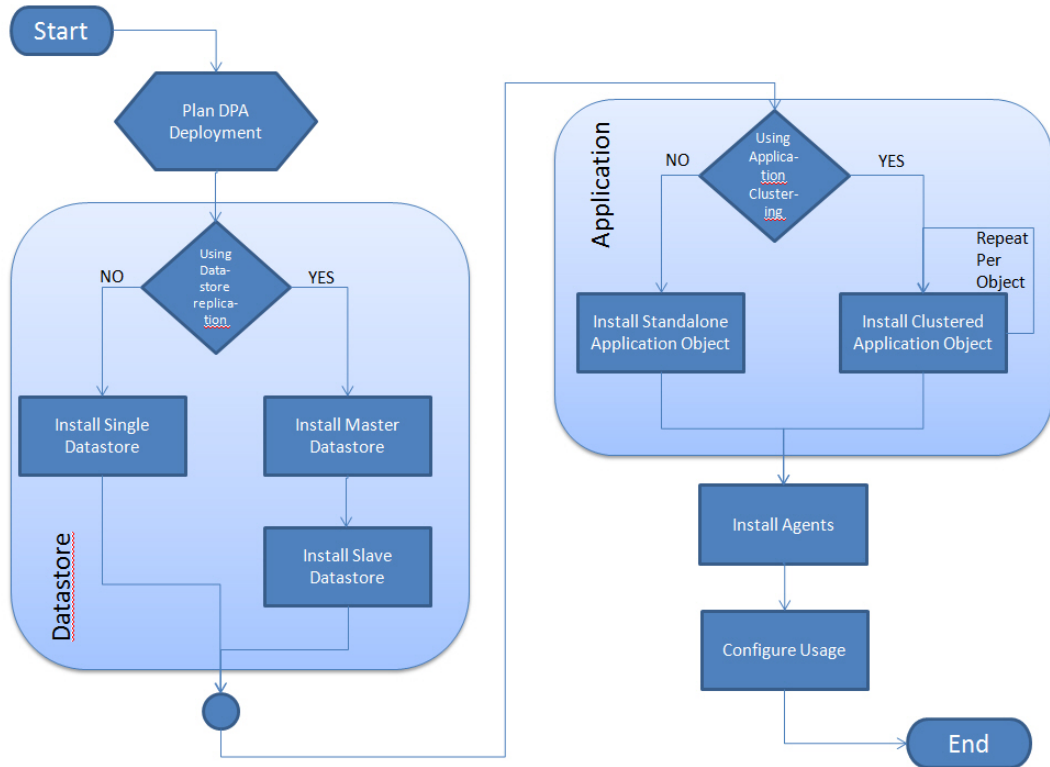
Table 6 DPA cluster port settings (continued)

Port	Description	Traffic direction
7500	Multicast over UDP	Bidirectional connection for Cluster
7600	Multicast over TCP	Inbound connection for Cluster
8090	TCP port used for intra-service communication	Localhost connection
9002	TCP port used for the HTTPS service.	Inbound connection over SSL from UI, CLI and REST API clients.
9003	TCP port used for DPA Datastore communications.	Outbound connection to DPA Datastore.
9005	TCP port used for Jboss Management	Localhost connection
9876	Multicast over TCP	Bidirectional connection for Cluster
9999	TCP port used for Jboss Management	Localhost connection
23364	Multicast over TCP	Bidirectional connection for Cluster
45688	Multicast over TCP	Bidirectional connection for Cluster
45689	Multicast over TCP	Bidirectional connection for Cluster
45700	Multicast over UDP	Bidirectional connection for Cluster
54200	Multicast over UDP	Bidirectional connection for Cluster
54201	Multicast over UDP	Bidirectional connection for Cluster
55200	Multicast over UDP	Bidirectional connection for Cluster
55201	Multicast over UDP	Bidirectional connection for Cluster
57600	Multicast over TCP	Bidirectional connection for Cluster

Installation and configuration overview

The DPA installation workflow provides a high-level workflow of tasks for installing DPA with various configurations.

Figure 2 DPA installation workflow



The Installation and configuration overview lists the tasks you need to perform for installing DPA and configuring data monitoring.

Table 7 Installation and configuration overview

Action	Comments
Set up host computer	
Provide at least two hosts for DPA server installation: One for the initial DPA Application server, and one for the Datastore. A separate host is required for the Datastore and Application server so that the operating system on each server can successfully and properly manage the IO performance needs of one service and the RAM and caching requirements of the other service, without the two services competing with each other for resources.	DPA must not be installed on servers already running other applications. For installation in a production environment, you need one host for the Application Service and a separate host for the Datastore Service. recommends that you use a dedicated server with at least 2GB of temporary space. The Compatibility Guide provides more information.

Table 7 Installation and configuration overview (continued)

Action	Comments
<p>Provide a host for DPA Agent installation (optional).</p>	<p>If the DPA server is running on Windows and the discovered host(s) are also Windows, you need not install an Agent on the discovered host. However, we recommend that you use the Agent installed on the DPA Server hosts for DPA Server monitoring only.</p> <p>If the DPA server resides on a Linux host and you are performing client discovery of Windows hosts, at least one DPA agent must be installed on a Windows Agent.</p>
<p>Ensure that DPA and all its components are configured as exceptions in any antivirus software.</p>	<p>Occasionally DPA components are shut down or associated files are quarantined by antivirus software if not defined as exceptions.</p>
<p>Provision networking infrastructure and a shared directory if installing multiple Application servers (DPA clustering).</p>	<ul style="list-style-type: none"> • Allocate a dedicated VLAN for use by the DPA Application servers. If a dedicated VLAN is not available, ask your network administrator for a UDP Multicast group address that can be used for the DPA cluster. • To increase resiliency and quality of service, provision a hardware load-balancing switch as a gateway to the DPA Application servers. • Configure a shared directory that will be accessible by all Application Servers. DPA will use this shared directory for writing scheduled reports and other temporary files that all Application Servers need to access.
<p>Check VMware or Hyper-V requirements.</p>	<p>DPA has been certified to work on a Linux or Windows virtual machine in a VMware or Hyper-V environment. The Software Compatibility Guide provides more information.</p>
<p>Configure virtual infrastructure memory and CPU</p>	<p>Configuring virtual infrastructure memory and CPU on page 18 provides more information.</p>
<p>Open or disable firewalls for communication between the DPA servers.</p>	<p>If you want to use secure communication for connecting to the Application server on port 9002, ensure that TLS (Transport Layer Security) settings are enabled for secure communication in your browser settings.</p> <p>When installing on DPA Servers, the operating system/software-based firewalls can be disabled or have ports opened for communication between the DPA Application server, the DPA Datastore server, and the</p>

Table 7 Installation and configuration overview (continued)

Action	Comments
	<p>DPA Agents prior to installing the DPA components.</p> <p>Typically, the network in which the DPA servers and DPA Agents reside are secure and behind a network firewall. This means that you could choose to disable operating system/software based firewalls. If you choose to leave the operating system/software based in effect, you must open/unblock the required ports. Communications settings in DPA on page 18 provides information.</p> <p>If on Linux and you choose to disable the firewall, run the following commands to disable and ensure that the firewall remains disabled after startup or reboot:</p> <ul style="list-style-type: none"> • Run <code>iptables stop</code>. • Set the <code>chkconfig</code> utility to <code>iptables off</code>.
Install the host operating system on the DPA Server(s) and Agent host and install all required patches.	The Software Compatibility Guide lists the required architectures and patches.
Install all required software on the agent host after the DPA latest release Application Server is ready.	When monitoring applications or devices remotely, you may need to install additional software on the Agent host. For example, the NetWorker client must be installed on the Agent host if the Agent will be used to monitor NetWorker remotely. For more information see Environment discovery in DPA on page 151
If DNS is not enabled in the environment, add the IP address and FQDN of the SharePoint server on the DPA Application server's hosts file.	DPA and SharePoint integration requires the IP address and FQDN to enable you to publish reports to SharePoint and to configure the SharePoint port. The SharePoint port is configurable. The default port, if no port is specified, is 80. You can set the port by using a standard URL in the existing URL field in the SharePoint settings dialog. System Settings on page 86, SharePoint settings table, provides information.
If you are going to use LDAP User Authentication on your DPA server, gather the information needed for configuration	<p>You need the following information for LDAP User Authentication configuration:</p> <ul style="list-style-type: none"> • LDAP Server Name/IP • Use SSL? • LDAP Server Port

Table 7 Installation and configuration overview (continued)

Action	Comments
	<ul style="list-style-type: none"> • LDAP Version • Distinguished Name of Base Directory • Identification Attribute
<p>Download and save the DPA binaries</p>	<p>To download the DPA Server and Agent binaries, go to the DPA downloads section of http://support.emc.com.</p> <p>Save the DPA Server and Agent binaries locally.</p>
<p>Obtain and save DPA Licenses</p>	
<p>Save the required license files on your local machine for easy access during installation. The DPA installation wizard prompts you to browse for the license file at license installation.</p>	<p>You must know the IP address of the primary Datastore server.</p> <p>For more information on obtaining DPA licenses or types of DPA licenses available and required, contact your Account Representative.</p>
<ul style="list-style-type: none"> • For new non-migrated installations - Obtain DPA licenses for all components that will be monitored. • For migrated 5.x installations - Existing licenses will be migrated. • The CLP license is required for new DPA functionality and increased capacity on a DPA instance. If you are not adding capacity or changing to new DPA latest release functionality, import of CLP licenses is not required. If you are migrating from DPA version 5.x to DPA, the existing licenses are migrated with your configuration and data. When not increasing capacity or changing functionality on existing WLS licenses, WLS licenses can only coexist with CLP license types if they are imported before CLP licenses. CLP and WLS license coexistence in DPA on page 72 provides more information. 	<p>A DPA license is required to administer DPA after installation.</p> <p>DPA is bundled with a 90-day evaluation license. The evaluation license is created from the time of DPA installation, is valid for up to 90 days, and allows access to all features. If you import a license during 90-day evaluation license period, the evaluation license is removed and you have access to DPA features according to license you imported.</p> <p>For information on required DPA licenses or on purchasing licenses for your DPA installation, contact your Sales Representative.</p>
<p>Provide the Solutions Enabler (SE) licenses.</p>	<ul style="list-style-type: none"> • A minimum of one gatekeeper per HBA per Symmetrix is required. • One Solutions Enabler host can discover all VNX/CLARiiON arrays through IP address. For VNX/CLARiiON discovery, we recommend installing Solutions Enabler on the DPA server.

Table 7 Installation and configuration overview (continued)

Action	Comments
	<ul style="list-style-type: none"> The <i>Software Compatibility Guide</i> describes the versions of Solutions Enabler required for storage array discovery.
Install DPA	
Install the DPA software.	Install the DPA server and agent according to the installation instructions. Installing the Datastore Service on page 30, Installing the Application Service on page 32, and Installing the DPA Agent on page 51 provide more information.
Configure host array discovery and Solutions Enabler hosts	
Configure Symmetrix and VNX/CLARiiON array for discovery	Configuration of storage arrays for replication analysis on page 186 provides more information. The steps in this section apply only if you are monitoring a storage array, database, or Microsoft Exchange Server for replication analysis.
Provide the Solutions Enabler host used to discover Symmetrix or VNX/CLARiiON storage arrays.	The <i>Software Compatibility Guide</i> describes the versions of Solutions Enabler required for storage array discovery, and the software that must be installed on the Solutions Enabler host. The host must be able to connect to the Symmetrix array by a SAN connection. The host must have the TCP port 443 or 2163 enabled for the VNX/CLARiiON connection.
Configure the environment for data protection monitoring	
Ensure that the required ports between the DPA Agent host and the monitored server or devices are open and communication is possible over the protocol.	Communications settings in DPA on page 18 lists the protocols and default DPA ports required for communication between the agent and the monitored device or server.
Ensure that the DPA credential used to connect to the monitored device or server is sufficient, or have the new credential details ready.	Permissions on page 17 lists the default settings for the DPA credentials that are installed with DPA.
Set up monitoring of RecoverPoint (if applicable).	RecoverPoint agent host and application host requirements are listed in Monitoring of RecoverPoint on page 185
Discover and configure Application Host import (if monitoring Microsoft Exchange or a database).	<ul style="list-style-type: none"> If a remote agent is being used to import hosts, the DPA server must be able to resolve the agent host. If application discovery is being performed without an agent, Configuring

Table 7 Installation and configuration overview (continued)

Action	Comments
	<p>for Replication Analysis on page 184 provides more information.</p>
<p>Define the data protection policies</p>	
<p>Prepare the details of the policies that DPA will monitor for compliance.</p>	<p>For replication analysis, the Data protection policy details consist of:</p> <ul style="list-style-type: none"> • The type of replication, (SRDF/S, SRDF/A, MirrorView, RecoverPoint, and so forth). • Whether the replication is Point-in-Time or continuous. • The replication target destination. For data protection reporting, the policies are: • Chargeback Policies - For financial cost analysis of data protection operations. • Protection Policies - To analyze compliance with recovery time objective (RTO) and recovery point objective (RPO) data protection targets. <p>Policies on page 203 provides more information.</p>

CHAPTER 2

Installing DPA

This chapter includes the following sections:

- [DPA server installation](#)..... 30
- [DPA Agent installation](#)..... 51
- [Installing by using command line installation](#).....55
- [DPA postinstallation steps](#).....60
- [Upgrades](#).....64

DPA server installation

The DPA server installation involves two stages:

1. Installing the Datastore service
2. Installing the Application service

[Application clustering](#) on page 35 provides information on installing with clustering. [Datastore Replication](#) on page 45 provides information on installing with Datastore Replication.

Installation of the Application service before the Datastore service results in failure of Application service installation. If you encounter issues during the installation, [Troubleshooting](#) on page 237 provides information.

The procedures in this section are applicable to new installations. For upgrades from previously supported DPA versions to DPA 18.1, and to install version 18.1, see [Upgrades](#). The DPA Release Notes provides information on supported upgrades.

The DPA installer runs on Windows and Linux, provided that your Linux installation supports running a UI. The following procedures describe installation in a Windows 64-bit environment.

Installing the Datastore Service

This procedure includes implementation for a normal Datastore installation without clustering and Datastore Replication.

Before you begin

- Ensure that you log in as a local administrator or a Domain administrator with full local access.
- If UAC is enabled on a Windows host, start the installer by Run as Administrator.
- Copy the installation binary to the server or to your local machine.
- If installing on UNIX/Linux, ensure that you are logged in as root. You could experience problems with the Datastore server if you install after becoming root through certain SU-type security software; for example, using the `sesu` command.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that you have the IP Address of the Application server for the Agent to communicate with. If installing on Linux IPv6, ensure that you also have the IPv6 Interface ID of the Datastore server. You are prompted for this in the **Configure Agent** window of the Datastore installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Linux Agent machine and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface Id (in this example, `2`).

Procedure

1. Double-click the DPA server binary to start the installation.
For Linux, provide execute permission to the Linux binary and execute it as `./DPA-Server-Linux-x86_64****.bin`.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to activate the option to accept the terms of the License Agreement. Click **Next**.

4. In the Installation Options screen, select to install Datastore service, click **Next**.
5. If you do not perform an advanced installation, click **Next** and follow the installation wizard.

To perform an advanced installation, select the Show Advanced Installation Options checkbox in the Advanced Installation screen, click **Next**, and follow the installation wizard.

The Advanced Options are:

- **Do not register DPA services:** Prevent the registration of the Datastore service with the operating system manager. This will prevent the Datastore service from being started after host reboot. You must use the DPA Command Line Interface to install the service with the operating system.
 - **Do not start DPA services:** Prevent the starting of the Datastore services after installation. Use of the DPA Command Line Interface will be required to start the service.
 - **Install with advanced datastore layout:** Configure the Datastore service with the required filesystems distributed across different disks to optimize performance.
 - **Install services under specified account:** Run datastore and agent services under a specified account.
6. When prompted, choose the installation folder.
Choose the default location or browse to another folder location.
 7. Review the Pre-Installation Summary, the disk space information in particular, click **Install**.
The installation proceeds.

If there is not enough disk space, cancel the installation or choose a different drive on which to install DPA.
 8. When prompted, select the IP addresses that the Datastore should listen on for connections from the DPA Application Server(s).
 9. On Linux, when prompted for Datastore Replication Option, select either Y or N.

```
By default the DPA datastore service is installed not configured for
replication.
If replication is required please enter 'Y' and then the role of this
datastore installation.
Do you wish to configure for replication (Y/N):
```

10. When prompted, enter the IP address of the DPA Application Server that will use the Datastore from step 8 and then click **Add** and **Next**. On Linux, select option **Add an Application Client Address** and **Review and Complete**.

```
Please enter the IP addresses for all DPA application service hosts
that will
connect to and use this datastore.
```

```
At least one IP address must be provided.
```

```
Additional clients can be added to the datastore access using the DPA
command
line interface.
```

```
1- Add an Application Client Address
2- Remove an Application Client Address
```

```
3- Review and Complete
Select action:
```

11. When prompted, specify the Datastore password.

Note the following regarding the Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa datastore dspassword` command can be used to reset the DPA Datastore password . [dpa datastore dspassword](#) on page 139 provides more information.

12. When prompted, specify the DPA Agent password:

Note the following regarding the Agent password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa agent --set-credentials` command can be used to reset the DPA Agent password . [dpa agent --set-credentials](#) provides more information.

13. When the DPA Datastore Server installation is complete, click **Done**.

Installing the Application Service

This procedure includes implementation for a normal Application service installation without clustering and Datastore Replication.

Before you begin

About this task

- To ensure secure communication between the DPA Server and Agent, set the Agent registration password using the `dpa app agentpwdCLI` command on the DPA Application Server host. You must also set this password on all DPA Agent hosts.[dpa application agentpwd](#) provides information. Then restart the Application service. Ensure that you set this password for each Agent.
- Copy the Agent installation binary to the server or to your local machine.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that the Datastore service option is checked, and that the Datastore service is running.
- If installing with Advanced Options on Linux IPv6, and the Agent wants to talk to a different application server or a Load Balancer, for example, in case of a cluster, ensure that you have the IP Address of the Application server for the Agent to communicate with. You are prompted for this in the **Configure Agent** window of the Application server installation. To get the IPv6

Interface ID, run the `ip addr show` command on the Application server and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server or the load balancer to which the Agent wants to connect (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface ID of the current Application server (in this example, `2`).

- If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at <https://support.emc.com/downloads/37716 EMC-Secure-Remote-Services-Virtual-Edition> on EMC Online Support provides more information on ESRS-VE installations.

The Application service installation process is similar to installing the Datastore service.

Procedure

1. Double-click the DPA server binary to start the installation.
For Linux, provide execute permission to the Linux binary and execute it as `./DPA-Server-Linux-x86_64****.bin`.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to enable the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Application service, click **Next**.
5. If you do not perform an advanced installation, click **Next** and follow the installation wizard.

The Advanced Options are:

- **Do not register DPA services:** Prevents the registration of the service with the operating system service manager. This option prevents the DPA services from being started after a host reboot.
 - **Do not start DPA services:** Prevents the DPA services from being started after installation. Use of the DPA command line interface is required to start the service.
 - **Install the DPA services as clusterable:** Configures the DPA service to discover and join any present DPA cluster.
 - **Install services under specified account:** Run application and agent services under a specified account. Not applicable for clusters.
The rest of the installation is similar to the Datastore installation.
6. In Linux, specify the user account.

```
Specify user
-----

Please enter a user account to run services.

Please enter the user name (Default: root):
```

7. For Linux, when prompted for security warning, select the appropriate option.

```
Are you sure you want to install services under root account?
```

```
->1- OK
    2- Cancel
```

```
ENTER THE NUMBER OF THE DESIRED CHOICE, OR PRESS <ENTER> TO ACCEPT THE
DEFAULT:
```

8. Review the Pre-Installation Summary, the disk space information in particular, click **Install**. The installation proceeds.

If there is not enough disk space, cancel the installation or choose a different drive to install DPA on.

Note: A datastore connection failure error might occur if the relevant firewalls required to communicate between Application Server and the Datastore are not open. [Communications settings in DPA](#) on page 18 provides information.

9. In the **Connect to Remote DPA Datastore** step, enter the IP address for the DPA Datastore server previously installed.

The installation resumes.

10. When prompted, specify the name or IP address of the DPA Application server host with which the DPA Agent will communicate. By default the Agent communicates with the local Application server with IP address 127.0.0.1. In a clustered configuration provide the IP address of the load balancing switch placed in front of the Application servers. Click **Next**.

The DPA Application service installation is now complete.

11. When prompted, specify the Datastore password.

Note the following regarding Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa application dspassword` configures the DPA Datastore password. [dpa application dspassword](#) on page 133 provides more information.

12. When prompted, specify the Administrator password.

Note the following regarding the Administrator password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa app adminpassword` command can be used to reset the DPA Administrator's password and enable the DPA Administrator account when the DPA Datastore service is up and running. [dpa application adminpassword](#) on page 130 provides more information.

13. When prompted, specify the DPA Agent password.

Note the following regarding the Agent password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa agent --set-credentials` command can be used to reset the DPA Agent password. [dpaagent --set-credentials](#) provides more information.

14. Click **Done**.

After the installation is complete, start the DPA Server and license the Server. [DPA postinstallation steps](#) on page 60 provides more information.

Application clustering

You can set DPA up in a clustered configuration, with multiple DPA Application Servers working with a single DPA Datastore Server. Clustering allows the ability for Application servers to dynamically start, share workload with other Application servers, and be stopped as demand decreases.

Clustered Application servers provide many benefits:

- Increased resiliency
- Load balancing of workload when placed behind a load-balancing switch that you provide
- Ability to scale the DPA deployment rapidly
- Flexible, green resource management
- Reduction of single points of failure

Once multiple Application Servers have been configured as a cluster you can start and stop individual application servers based on load, such as powering-on additional servers for end-of-month reporting or other high-usage periods. You can add new servers to running clusters to improve performance due to load.

Ensure that all cluster nodes are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.

You can configure Application clustering:

- during a fresh installation; [Installing the Master Application Service with clustering](#) on page 38 and [Installing the Slave Application Service with clustering](#) on page 41 provide information.
- during an upgrade; [Upgrading existing clusters](#) on page 68 and [Upgrading with Datastore Replication and existing clusters](#) on page 70 provide more information.
- after installation and configuration; [Adding an Application server to a cluster after DPA deployment](#) on page 119 provides more information.

Restrictions and recommendations for clustering

Observe the following restrictions and recommendations when configuring Clusters:

- DPA supports a maximum of four nodes in a cluster:

- One Master
- Three Slaves
- Each cluster of Application servers must be on its own LAN/VLAN.
 - Spanning LANs is not possible.
 - Clustering is UDP-broadcast based.
- Clusters can communicate cross-LAN to Datastore.
- A physical load-balancing switch should be placed in front of the Application server cluster to manage the load between DPA Application server objects. The use of software load-balancing switches is not recommended.
- Any configuration accessible via the DPA web console is stored in the Datastore and is accessible cluster-wide. Any configuration operation that requires the use of the `dpa` executive utility, such as `"dpa application promote"`, is local to the object on which it was executed. [Adding an Application server to a cluster after DPA deployment](#) on page 119 and [dpa application commands](#) on page 130 provide information on the `dpa application promote` command.
- If you are implementing Application server clustering, ensure that you complete all cluster configuration before enabling encryption on Application servers.

Installing the Datastore Service with clustering

This procedure includes implementation of a cluster with a load balancer, Datastore, Master Application server, and one or more Slave Application Server.

Before you begin

- Ensure that you log in as a local administrator or a Domain administrator with full local access.
- If UAC is enabled on a Windows host, start the installer by Run as Administrator.
- Copy the installation binary to the server or to your local machine.
- If installing on UNIX/Linux, ensure that you are logged in as root. You could experience problems with the Datastore server if you install after becoming root through certain SU-type security software; for example, using the `sesu` command.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that you create a common shared directory for reports that is accessible from both the Application nodes. For example, on Windows Cluster Datastore1 \ `\WinClusterDS1\cluster_share`. The shared directory must have read/write permissions for the users in ClusterApp1 and ClusterApp2 who own the DPA service.
- Ensure that you have the IP Address of the Application server for the Agent to communicate with. If installing on Linux IPv6, ensure that you also have the IPv6 Interface ID of the Datastore server. You are prompted for this in the **Configure Agent** window of the Datastore installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Linux Agent machine and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface Id (in this example, `2`).

- Verify that all the machines are on the same network adapter in vCenter.
- If installing Datastore Replication:
 - Plan the final deployment topology before beginning installation. Additional resources are available on the EMC Community Network (ECN) that provide guidance and best practice for planning your deployment.
 - Have all hosts and IP addresses predetermined and available.
 - Ensure that all Datastore server or Application server, including clustered nodes, are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.

Procedure

1. Double-click the DPA server binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to activate the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Datastore service, click **Next**.
5. If you do not perform an advanced installation, click **Next** and follow the installation wizard.
To perform an advanced installation, select the Show Advanced Installation Options checkbox in the Advanced Installation screen, click Next, and follow the installation wizard.

The Advanced Options are:

- **Do not register DPA services:** Prevent the registration of the Datastore service with the operating system manager. This will prevent the Datastore service from being started after host reboot. You must use of the DPA Command Line Interface to install the service with the operating system.
 - **Do not start DPA services:** Prevent the starting of the Datastore services after installation. Use of the DPA Command Line Interface will be required to start the service.
 - **Install with advanced datastore layout:** Configure the datastore service with the required filesystems distributed across different disks to optimize performance.
Selecting Advanced Installation Options also enables you to configure Datastore Replication and select a replication role for this server later on in the installer.
6. When prompted, choose the installation folder.
Choose the default location or browse to another folder location.
 7. Review the Pre-Installation Summary, the disk space information in particular, click **Install**.
The installation proceeds.
If there is not enough disk space, cancel the installation or choose a different drive on which to install DPA.
 8. In the **Datastore Listening Addresses** window, specify the IP addresses that the Datastore service should listen on for connections from the DPA Application services.
 9. In the **Configure Datastore Access** window, enter the IP addresses of the DPA Application Servers that will use the Datastore and then click **Add** and **Next**.
Enter IP addresses for each DPA Application Server in the clustered configuration.
 10. In the **Datastore Agent Address** window, specify the alternative address for the Datastore Agent to be the Load Balancer IP Address.
 11. If you are configuring Datastore Replication, select **Enable datastore replication** > **and select the replication role for this server** > **SLAVE**. Click **Next**.

- a. Provide the IP address or FQDN of the Master Datastore server.
 - b. When prompted in the **Configure Agent** window, enter the FQDN or IP address of the DPA Application service that the installed DPA Agent needs to communicate with.
By default, the Agent communicates with the Application server specified earlier in the wizard.
 - c. If you are working in a Linux IPv6 environment, provide the load balancer's FQDN/IP address in the following format: `IPV6Address%Interface_Id`
Click **Next**.
12. When prompted, specify the Datastore password.
Note the following regarding the Datastore password:
 - Blank passwords are not supported.
 - Minimum length is 9 characters.
 - The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
 - The `dpa datastore dspassword` command can be used to reset the DPA Datastore password . [dpa datastore dspassword](#) on page 139 provides more information.
 13. When the DPA Datastore Server installation is complete, click **Done**.
 14. On a command prompt, run the `dpa svc status` command to verify that the Datastore service is running.
 15. Set the database connection pool size in all Datastore nodes. Run:


```
# dpa ds tune --connections xxx <RAM>GB
```

 where `xxx` is approximately 250 per each Application server and `RAM` is the amount of RAM . For example, you would set a `xxx` figure of 500 for a two-node cluster.

If the cluster is enabled with Datastore Replication, run this command for all Datastore Slaves.

Installing the Master Application Service with clustering

Before you begin

- Copy the installation binary to the server or to your local machine.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that the Datastore service is running.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- If installing with Advanced Options on Linux IPv6, and the Agent wants to talk to a different application server or a Load Balancer, for example, in case of a cluster, ensure that you have the IP Address of the Application server for the Agent to communicate with. You are prompted for this in the **Configure Agent** window of the Application server installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Application server and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the % refer to the IPv6 of the Application server or the load balancer to which the Agent wants to connect (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface ID of the current Application server (in this example, `2`).

- Plan the final deployment topology before beginning installation. Additional resources are available on the EMC Community Network (ECN) that provide guidance and best practice for planning your deployment.
- Have all hosts and IP addresses predetermined and available, including the IP address configured for the load-balancing switch that will be placed in front of the Application servers.
- Ensure that all cluster nodes are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.
- Specify a common directory that is shared across all nodes. This is the location of the folder where the reports generated by the DPA Application node are stored.
- If installing Application server clustering on UNIX, ensure that you specify the common shared directory to a local directory mapped to a UNIX NFS or CIFS network share.
 - Ensure that you create a username in all Application nodes within the cluster with the same UID and GID. During installation, you are prompted to log on with a valid UNIX username and password. System users like `ftuser` and `bin` cannot be used.
 - Ensure that you have read and write access to the shared directory that you specify.
 - Ensure that you validate the path if it is tied to a network share.
- If installing Application server clustering on Windows, ensure that you specify the common shared directory as a UNC (Windows *Universal Naming Convention*) path.
 - Ensure that you validate the path specified.
 - Configure and grant read and write access to a user account (username and password) to the share that you specify above. This user account must have the **Log on as a service** Windows permissions enabled.
- If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at https://support.emc.com/downloads/37716_EMCSecure-Remote-Services-Virtual-Edition on EMC Online Support provides more information on ESRS-VE installations.

The Application service installation process is similar to installing the Datastore service.

Procedure

1. Double-click the DPA server binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to enable the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Application service, click **Next**.
5. Ensure that **Show Advanced Installation Options** is enabled and click **Next**.

The Advanced Options are:

- **Do not register DPA services:** Prevents the registration of the service with the operating system service manager. This option prevents the DPA services from being started after a host reboot.
- **Do not start DPA services:** Prevents the DPA services from being started after installation. Use of the DPA command line interface is required to start the service.
- **Install the DPA services as clusterable:** Configures the DPA service to discover and join any present DPA cluster.

If you would like to add an Application Object to a cluster, select **Install the DPA services as clusterable** and follow the steps in the wizard.

At the prompt for a common location for Application servers for reports, ensure that you specify a common directory that is shared across all nodes. The Shared Directory for reports is required when you run multiple Application nodes.

If installing on UNIX, the installer prompts you to specify the user account username of a valid user that has read and write access to the share specified in “Before you begin.”

If installing on Windows, ensure that you configure the required common and shared UNC folder and enter the Domain username and password with access to that specified directory. “Before you begin” provides more information.

The rest of the installation is similar to the Datastore installation.


6. In the **Application Advanced Options** window, ensure that **Install the DPA services as clusterable** is enabled and click **Next**.
7. In the **Identify the DPA Datastore to connect to** window, specify the Datastore IP address and click **Next**.
8. In the **Application Cluster Address** window, select the IP address that the Application Server wants to listen on and click **Next**.
9. In the **Application Cluster Options** window, select the Application Role as **Master** from the dropdown menu and click **Next**.
10. In the **Choose a Folder** window, specify the shared folder that you would like to be used for reporting and click **Next**.
11. In the **Username** window, specify the username and password for that user who will now own the DPA Service. Click **Next**

Ensure that the user has read and write permissions to the shared folder specified in step 11.

The username should be in the form of <Domain\User> if it is a domain. If it is not a domain, the username should be in the form of <HOSTNAME\User>.

12. In the **Enter Alternative Agent Address** window, specify the alternative Agent address to be the load balancer's IP Address and click **Next**
13. Review the Pre-Installation Summary, the disk space information in particular, click **Install**. The installation proceeds.

If there is not enough disk space, cancel the installation or choose a different drive to install DPA on.

 **Note:** A datastore connection failure error might occur if the relevant firewalls required to communicate between Application Server and the Datastore are not open. [Communications settings in DPA](#) on page 18 provides information.

14. In the **Connect to Remote DPA Datastore** step, enter the IP address for the DPA Datastore server previously installed.

The installation resumes.

15. When prompted, specify the Datastore password.

Note the following regarding the Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:

- A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
 - The `dpa datastore dspassword` command can be used to reset the DPA Datastore password . [dpa datastore dspassword](#) on page 139 provides more information.
16. When prompted, specify the Administrator password.
- Note the following regarding Administrator password:
- Blank passwords are not supported.
 - Minimum length is 9 characters.
 - The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
 - The `dpa app adminpassword` command can be used to reset the DPA Administrator's password and enable the DPA Administrator account when the DPA Datastore service is up and running. [dpa application adminpassword](#) on page 130 provides more information.
17. Click **Done**.
- After the installation is complete, start the DPA Server and license the Server. [DPA postinstallation steps](#) on page 60 provides more information.
18. On the command prompt, run the `dpa app con` command to check the Application Server configuration.
- You may notice after running the `dpa app con` command that the bind address is set to 0.0.0.0. DPA does this to allow for any connection address.
- The output should indicate that the operation mode is cluster and that the cluster role is Master.
19. If you are adding a multicast address to the cluster, demote the cluster to a standalone and then promote it to a cluster node:
- If you are not adding a multicast address to the cluster, proceed to step 19.
- a. On the command prompt, run the `dpa app stop` command to stop the Application Server.
 - b. Run the `dpa app demote` command to demote the node to a standalone node.
 - c. Run the `dpa app promote` command to promote the Application node to a cluster. Ensure that you include the bind address, the multicast address and the shared folder path. Ensure also that you specify the role.
20. On the command prompt, run the `dpa app start` command to start Application service.
21. Verify correct installation and configuration in the `server.log` file for the message `DPA master started successfully`.

Installing the Slave Application Service with clustering

Before you begin

- Copy the installation binary to the server or to your local machine.

- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that the Datastore service is running.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- If installing with Advanced Options on Linux IPv6, and the Agent wants to talk to a different application server or a Load Balancer, for example, in case of a cluster, ensure that you have the IP Address of the Application server for the Agent to communicate with. You are prompted for this in the **Configure Agent** window of the Application server installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Application server and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server or the load balancer to which the Agent wants to connect (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface ID of the current Application server (in this example, `2`).

- Plan the final deployment topology before beginning installation. Additional resources are available on the EMC Community Network (ECN) that provide guidance and best practice for planning your deployment.
- Have all hosts and IP addresses predetermined and available, including the IP address configured for the load-balancing switch that will be placed in front of the Application servers.
- Ensure that all cluster nodes are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.
- Specify a common directory that is shared across all nodes. This is the location of the folder where the reports generated by the DPA Application node are stored.
- If installing Application server clustering on UNIX, ensure that you specify the common shared directory to a local directory mapped to a UNIX NFS or CIFS network share.
 - Ensure that you create a username in all Application nodes within the cluster with the same UID and GID. During installation, you are prompted to log on with a valid UNIX username and password. System users like `ftpuser` and `bin` cannot be used.
 - Ensure that you have read and write access to the shared directory that you specify.
 - Ensure that you validate the path if it is tied to a network share.
- If installing Application server clustering on Windows, ensure that you specify the common shared directory as a UNC (Windows *Universal Naming Convention*) path.
 - Ensure that you validate the path specified.
 - Configure and grant read and write access to a user account (username and password) to the share that you specify above. This user account must have the **Log on as a service** Windows permissions enabled.
- If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at <https://support.emc.com/downloads/37716 EMC-Secure-Remote-Services-Virtual-Edition> on EMC Online Support provides more information on ESRS-VE installations.

The Application service installation process is similar to installing the Datastore service.

Procedure

1. Double-click the DPA server binary to start the installation.

2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to enable the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Application service, click **Next**.
5. Ensure that **Show Advanced Installation Options** is enabled and click **Next**.

The Advanced Options are:

- **Do not register DPA services:** Prevents the registration of the service with the operating system service manager. This option prevents the DPA services from being started after a host reboot.
- **Do not start DPA services:** Prevents the DPA services from being started after installation. Use of the DPA command line interface is required to start the service.
- **Install the DPA services as clusterable:** Configures the DPA service to discover and join any present DPA cluster.

If you would like to add an Application Object to a cluster, select **Install the DPA services as clusterable** and follow the steps in the wizard.

At the prompt for a common location for Application servers for reports, ensure that you specify a common directory that is shared across all nodes. The Shared Directory for reports is required when you run multiple Application nodes.

If installing on UNIX, the installer prompts you to specify the user account username of a valid user that has read and write access to the share specified in “Before you begin.”

If installing on Windows, ensure that you configure the required common and shared UNC folder and enter the Domain username and password with access to that specified directory. “Before you begin” provides more information.

The rest of the installation is similar to the Datastore installation.

6. In the **Application Advanced Options** window, ensure that **Install the DPA services as clusterable** is enabled and click **Next**.
7. In the Identify the DPA Datastore to connect to window, specify the Datastore IP address and click **Next**.
8. In the **Application Cluster Address** window, select the IP address that the Application Server wants to listen on and click **Next**.
9. In the **Application Cluster Options** window, select the Application Role as **Slave** from the dropdown menu and click **Next**.
10. In the **Application Cluster Option** window, specify the Master node IP Address or FQDN with which the Slave should communicate and click **Next**.
11. In the **Username** window, specify the username and password for that user who will now own the DPA Service. Click **Next**

Ensure that the user has read and write permissions to the shared folder specified in step 10.

The username should be in the form of <Domain\User> if it is a domain. If it is not a domain, the username should be in the form of <HOSTNAME\User>.

12. In the **Enter Alternative Agent Address** window, specify the alternative Agent address to be the load balancer's IP Address, and click **Next**
13. Review the Pre-Installation Summary, the disk space information in particular, click **Install**. The installation proceeds.

If there is not enough disk space, cancel the installation or choose a different drive to install DPA on.

Note: A datastore connection failure error might occur if the relevant firewalls required to communicate between Application Server and the Datastore are not open. [Communications settings in DPA](#) on page 18 provides information.

14. In the **Connect to Remote DPA Datastore** step, enter the IP address for the DPA Datastore server previously installed.

The installation resumes.

15. When prompted, specify the Datastore password.

Note the following regarding the Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa datastore dspassword` command can be used to reset the DPA Datastore password. [dpa datastore dspassword](#) on page 139 provides more information.

16. When prompted, specify the Administrator password.

Note the following regarding Administrator password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa app adminpassword` command can be used to reset the DPA Administrator's password and enable the DPA Administrator account when the DPA Datastore service is up and running. [dpa application adminpassword](#) on page 130 provides more information.

17. Click **Done**.

After the installation is complete, start the DPA Server and license the Server. [DPA postinstallation steps](#) on page 60 provides more information.

18. On the command prompt, run the `dpa app con` command to check the Application Server configuration.

The output should indicate that the operation mode is cluster and that the cluster role is Slave.

19. If you are adding a multicast address to the cluster, demote the cluster to a standalone and then promote it to a cluster node:

If you are not adding a multicast address to the cluster, proceed to step 19.

- a. On the command prompt, run the `dpa app stop` command to stop the Application Server.
- b. Run the `dpa app demote` command to demote the node to a standalone node.

- c. Run the `dpa app promote` command to promote the Application node to a cluster. Ensure that you include the bind address, the multicast address and the shared folder path. Ensure also that you specify the role as Slave and the Master node IP address. For example:

```
dpa app promote --bind 10.10.211.212 --multicast 210.1.2.33 --role SLAVE
10.10.211.213 --path \\WinClusterDS1\cluster_share
```

20. On the command prompt, run the `dpa app start` command to start Application service.
21. Verify correct installation and configuration in the `server.log` file for the message `DPA slave started successfully`.

Datstore Replication

DPA Datstore Replication enables continuous, safe, and reliable replication so that DPA can maintain a replica copy, or *Slave*, of the primary Datstore, or *Master*, for resilience against a single point of failure. You can add additional slaves in a cascading fashion to the standard Master Slave configuration if required.

In the event of failure of the Master Datstore, the Slave can be updated to the Master role using the manual failover command, and the Application servers are then configured to use this new Master. Reconfiguration should normally take the same amount of time to take effect as the DPA Application and Datstore services startup take. [Carrying out Datstore server failover](#) on page 124 provides more information.

There can be only one Master Datstore per deployment. All Datstores are Masters on installation. Replication is enabled once a Slave Datstore can communicate with the Master Datstore. Data starts being replicated when an Application server is started.

You can configure Datstore Replication:

- during a fresh installation; [Installing the Master Datstore Service with Datstore Replication](#) and [Installing the Slave Datstore Service with Datstore Replication](#) provide information.
- during an upgrade; [Upgrading with Datstore Replication enabled with DPA 6.3 and later](#) on page 68 and [Upgrading with Datstore Replication and existing clusters](#) on page 70 provide information.
- after installation and deployment; [Configuring Datstore Replication after deployment](#) on page 122 provides more information.

Ensure that all Datstore nodes are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.

Configuring Datstore Replication

Procedure

1. Configure the Slave Datstore, either during or after installation.
2. Configure the Master Datstore, either during or after installation.
3. Install or, if already installed, start the Application server.

Installing the Master Datstore Service with Datstore Replication

This procedure includes implementation for a Master Datstore installation implementing Datstore Replication.

Before you begin

- Ensure that you log in as a local administrator or a Domain administrator with full local access.
- If UAC is enabled on a Windows host, start the installer by Run as Administrator.

- Copy the installation binary to the server or to your local machine.
- If installing on UNIX/Linux, ensure that you are logged in as root. You could experience problems with the Datastore server if you install after becoming root through certain SU-type security software; for example, using the `sesu` command.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that you have the IP Address of the Application server for the Agent to communicate with. If installing on Linux IPv6, ensure that you also have the IPv6 Interface ID of the Datastore server. You are prompted for this in the **Configure Agent** window of the Datastore installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Linux Agent machine and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface Id (in this example, `2`).

- Plan the final Datastore Replication deployment topology before beginning installation. Additional resources are available on the EMC Community Network (ECN) that provide guidance and best practice for planning your deployment.
- Have all hosts and IP addresses predetermined and available.
- Ensure that all Datastore server or Application server, including clustered nodes, are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.
- Ensure that the Application server chosen is the same one that the Master Datastore is using.

Procedure

1. Double-click the DPA server binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to activate the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Datastore service, click **Next**.
5. Select the **Show Advanced Installation Options** checkbox in the **Advanced Installation** screen, click **Next**.
6. Select **Install with advanced datastore layout** and click **Next**.
7. When prompted, choose the installation folder.
Choose the default location or browse to another folder location.
8. Review the Pre-Installation Summary, the disk space information in particular, click **Install**.
The installation proceeds.
If there is not enough disk space, cancel the installation or choose a different drive on which to install DPA.
9. In the **Datastore Listening Addresses** window, specify the IP addresses that the Datastore service should listen on for connections from the DPA Application services.
10. In the **Configure Datastore Access** window, enter the IP addresses of the DPA Application Servers that will use the Datastore and then click **Add** and **Next**.
Enter IP addresses for each DPA Application Server in the clustered configuration.

11. In the **Datastore Agent Address** window, specify the alternative address for the Datastore Agent to be the Load Balancer IP Address.
12. Select **Enable datastore replication** > **and select the replication role for this server** > **SLAVE**. Click **Next**.

- a. Provide the IP address or FQDN of the Master Datastore server.

- b. When prompted in the **Configure Agent** window, enter the FQDN or IP address of the DPA Application service that the installed DPA Agent needs to communicate with.

By default, the Agent communicates with the Application server specified earlier in the wizard.

- c. If you are using clustered DPA Application servers, provide the load balancer's FQDN/IP address. Provide the Application server/Load Balancer's IPV6 Address in the following format: `IPV6Address%Interface_Id`

The FQDN/IPAddress default value is left blank in case of a cluster and in case you are using clustered DPA Application servers Linux IPv6 application server because you must manually enter the `IPV6%Interface_Id`. In all other cases, the FQDN/IP Address is automatically populated with the default value of the Application server's IP Address.

Click **Next**.

13. When prompted, specify the Datastore password.

Note the following regarding the Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
- The `dpa datastore dspassword` command can be used to reset the DPA Datastore password. [dpa datastore dspassword](#) on page 139 provides more information.

14. When the DPA Datastore Server installation is complete, click **Done**.

15. On a command prompt, run the `dpa svc status` command to verify that the Datastore service is running.

16. Set the database connection pool size in all Datastore nodes. Run:

```
# dpa ds tune --connections xxx <RAM>GB where xxx is approximately 250 per each
Application server and RAM is the amount of RAM. For example, you would set a xxx figure
of 500 for a two-node cluster.
```

If the cluster is enabled with Datastore Replication, run this command for all Datastore Slaves.

Installing the Slave Datastore Service with Datastore Replication

This procedure includes implementation for a Slave Datastore installation implementing Datastore Replication.

Before you begin

- Ensure that you log in as a local administrator or a Domain administrator with full local access.

- If UAC is enabled on a Windows host, start the installer by Run as Administrator.
- Copy the installation binary to the server or to your local machine.
- If installing on UNIX/Linux, ensure that you are logged in as root. You could experience problems with the Datastore server if you install after becoming root through certain SU-type security software; for example, using the `sesu` command.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that you have the IP Address of the Application server for the Agent to communicate with. If installing on Linux IPv6, ensure that you also have the IPv6 Interface ID of the Datastore server. You are prompted for this in the **Configure Agent** window of the Datastore installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Linux Agent machine and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface Id (in this example, `2`).

- Plan the final Datastore Replication deployment topology before beginning installation. Additional resources are available on the EMC Community Network (ECN) that provide guidance and best practice for planning your deployment.
- Have all hosts and IP addresses predetermined and available.
- Ensure that all Datastore server or Application server, including clustered nodes, are using the same IP type of IP addressing, either IPv4 addresses or IPv6 addresses.
- Ensure that the Application server chosen is the same one that the Master Datastore is using.

Procedure

1. Double-click the DPA server binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to activate the option to accept the terms of the License Agreement. Click **Next**.
4. In the Installation Options screen, select to install Datastore service, click **Next**.
5. Select the **Show Advanced Installation Options** checkbox in the **Advanced Installation** screen, click **Next**.
6. Select **Install with advanced datastore layout** and click **Next**.
7. When prompted, choose the installation folder.
Choose the default location or browse to another folder location.
8. Review the Pre-Installation Summary, the disk space information in particular, click **Install**.
The installation proceeds.
If there is not enough disk space, cancel the installation or choose a different drive on which to install DPA.
9. In the **Datastore Listening Addresses** window, specify the IP addresses that the Datastore service should listen on for connections from the DPA Application services.
10. In the **Configure Datastore Access** window, enter the IP addresses of the DPA Application Servers that will use the Datastore and then click **Add** and **Next**.

Enter IP addresses for each DPA Application Server in the clustered configuration.

11. In the **Datastore Agent Address** window, specify the alternative address for the Datastore Agent to be the Load Balancer IP Address.
12. Select **Enable datastore replication** > **and select the replication role for this server** > **SLAVE**. Click **Next**.
 - a. Provide the IP address or FQDN of the Master Datastore server.
 - b. When prompted in the **Configure Agent** window, enter the FQDN or IP address of the DPA Application service that the installed DPA Agent needs to communicate with.

By default, the Agent communicates with the Application server specified earlier in the wizard.
 - c. If you are using clustered DPA Application servers, provide the load balancer's FQDN/IP address. Provide the Application server/Load Balancer's IPV6 Address in the following format: `IPV6Address%Interface_Id`

The FQDN/IPAddress default value is left blank in case of a cluster and in case you are using clustered DPA Application servers Linux IPv6 application server because you must manually enter the `IPV6%Interface_Id`. In all other cases, the FQDN/IP Address is automatically populated with the default value of the Application server's IP Address.

Click **Next**.

13. When the DPA Datastore Server installation is complete, click **Done**.
14. On a command prompt, run the `dpa svc status` command to verify that the Datastore service is running.
15. Set the database connection pool size in all Datastore nodes. Run:


```
# dpa ds tune --connections xxx <RAM>GB
```

 where `xxx` is approximately 250 per each Application server and `RAM` is the amount of RAM. For example, you would set a `xxx` figure of 500 for a two-node cluster.

If the cluster is enabled with Datastore Replication, run this command for all Datastore Slaves.

Installing the Application Service with Datastore Replication

This procedure for installing for the Application service installation is included for completeness. There is no special Application service implementation for Datastore Replication.

Before you begin

About this task

- Copy the installation binary to the server or to your local machine.
- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.
- Ensure that the Datastore service option is checked, and that the Datastore service is running.
- If installing on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- If installing with Advanced Options on Linux IPv6, and the Agent wants to talk to a different application server or a Load Balancer, for example, in case of a cluster, ensure that you have the IP Address of the Application server for the Agent to communicate with. You are prompted for this in the **Configure Agent** window of the Application server installation. To get the IPv6

Interface ID, run the `ip addr show` command on the Application server and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the Application server or the load balancer to which the Agent wants to connect (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface ID of the current Application server (in this example, `2`).


- If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at <https://support.emc.com/downloads/37716 EMC-Secure-Remote-Services-Virtual-Edition> on EMC Online Support provides more information on ESRS-VE installations.

The Application service installation process is similar to installing the Datastore service.

Procedure

1. Double-click the DPA server binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Scroll to the end of the agreement to enable the option to accept the terms of the License Agreement. Click **Next**.
4. In the **Installation Options** screen, select to install Application service and click **Next**.
5. If you do not perform an advanced installation, click **Next** and follow the installation wizard.
6. Review the Pre-Installation Summary, the disk space information in particular, click **Install**. The installation proceeds.

If there is not enough disk space, cancel the installation or choose a different drive to install DPA on.

 **Note:** A datastore connection failure error might occur if the relevant firewalls required to communicate between Application Server and the Datastore are not open. [Communications settings in DPA](#) on page 18 provides information.

7. In the **Connect to Remote DPA Datastore** step, enter the IP address for the DPA Master Datastore server previously installed.

The installation resumes.

8. When prompted, specify the name or IP address of the DPA Application server host with which the DPA Agent will communicate. By default the Agent communicates with the local Application server with IP address 127.0.0.1. In a clustered configuration provide the IP address of the load balancing switch placed in front of the Application servers. Click **Next**.

The DPA Application service installation is now complete.

9. When prompted, specify the Datastore password.

Note the following regarding the Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol

- A minimum of 1 special character
 - The `dpa datastore dspassword` command can be used to reset the DPA Datastore password. [dpa datastore dspassword](#) on page 139 provides more information.
10. Set the Administrator password.
- Note the following regarding the Administrator password:
- Blank passwords are not supported.
 - Minimum length is 9 characters.
 - The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
 - The `dpa app adminpassword` command can be used to reset the DPA Administrator's password and enable the DPA Administrator account when the DPA Datastore service is up and running. [dpa application adminpassword](#) on page 130 provides more information.
11. Click **Done**.
- After the installation is complete, start the DPA Server and license the Server. [DPA postinstallation steps](#) on page 60 provides more information.

Datastore Replication best practices

Observe the following best practices for Datastore Replication:

- You must restart the Datastore service any time the role between Master Datastore and Slave Datastore is changed.
- Use the replication configuration command `dpa ds rep` to check the status of replication. Running the `dpa ds rep` command on the Master Datastore displays if replication is streaming and what the Slave Datastore is. Running on the Slave Datastore tells you what the Master Datastore is.
- Before exporting a Datastore, ensure that you create an empty directory on the Datastore to which to export the Datastore file set. For example, `/tmp/export`.
- Master and Slave Datastores should have the same performance specifications and be installed on the same version of DPA.

DPA Agent installation

This section describes how to install the DPA Agent using the agent-only installation package. It is applicable to new installations.

An Agent is automatically installed on the DPA Application and Datastore servers. Therefore do not run this procedure on the DPA servers. For upgrades from previous DPA service packs to DPA 18.1, and to install DPA 18.1 latest versions, see [Upgrades](#).

Installing the DPA Agent

The following procedure explains installing the DPA Agent in a Windows environment.

Before you begin

- Ensure that ports are opened or disabled for communication between the DPA servers. [Installation and configuration overview](#) on page 23 provides information.

- Ensure that you have the IP Address of the DPA Application server for the Agent to communicate with. If installing on Linux IPv6, ensure that you also have the IPv6 Interface ID of the Agent. You are prompted for this in the **Configure Agent** window of the Agent installation. To get the IPv6 Interface ID, run the `ip addr show` command on the Linux Agent machine and use the output to find the IPv6 Interface ID. For example:

```
fe80::9c9b:36f:2ab:d7a2%2
```

Where the values before the `%` refer to the IPv6 of the DPA Application server (in this example, `fe80::9c9b:36f:2ab:d7a2`) and those after refer to the interface ID of the Agent (in this example, `2`).

Procedure

1. Double-click the DPA Agent binary to start the installation.
2. Click **Next**.
3. Read and accept End User License Agreement. Click **Next**.
4. Choose an installation folder and click **Next**.
5. Verify the Pre-Installation Summary and click **Install**.
6. Choose the Agent installation options:
 - **Do not start DPA Agent service** - this option prevents starting of the DPA Agent service after installation.
If you select this option, you must manually start the DPA Agent from the command line.
If you select Do not start DPA Agent service, click **Next**.
Type the fully qualified domain name or the IP address of the DPA Server that communicates with the DPA Agent.
 - **Agent will be used to monitor Oracle Database:** Select this option to monitor an Oracle database with the DPA Agent.
If you select this option, browse to the directory where the DPA Agent can find the Oracle Database device driver files.
7. Click **Next**.
8. In the **Configure Agent** window, enter the fully qualified domain name or the IP address of the DPA Application Server that communicates with the DPA Agent.
If you are installing on Linux IPv6 and are installing Linux Agents, enter the IPv6 Interface ID of the Linux Agent.
Click **Next**.
9. Set the same Agent password that you set during the DPA Datastore installation:
Note the following regarding the Agent password:
 - Blank passwords are not supported.
 - Minimum length is 9 characters.
 - The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character

10. Click **Done** to complete the installation.
11. Restart the Agent service.

After you finish

Follow the steps in [Setting DPA Agent registration password](#) on page 53.

Setting DPA Agent registration password

After installation of the DPA Agent, set the Agent password.

Procedure

1. Run `dpaagent --set-credentials` to set the DPA agent password.
[dpaagent --set-credentials](#) provides full command information.

Configure DPA version 18.1 Agent to go back and collect backup application data

By default, the newly installed DPA Agent starts collecting data from backup applications starting from the current date and time. If you would like to see alerts for failed backups within the previous days for auditing or other reasons, or if for any other reason you wish to collect days of backup application data, you can configure the newly installed DPA agent to collect data for user-defined number of hours.

Before you begin

You must have DPA 18.1 or later installed for this procedure.

For Linux

Procedure

1. Install the DPA Agent. Do not start the DPA Agent.
2. Add the following two lines to the `dpa.config` file:

```
VARIABLE_NAME=NUMBER_OF_BACKUP_HOURS
export VARIABLE_NAME
```

Where *VARIABLE_NAME* is the following for these backup applications:

NetWorker: AGENT_NSR_JOB_STARTTIME

Avamar: AGENT_AXION_JOB_STARTTIME

TSM: AGENT_TSM_JOB_STARTTIME

HPDP: AGENT_DP_JOB_STARTTIME

CommVault: AGENT_CV_JOB_STARTTIME

NetBackup: AGENT_NB_JOB_STARTTIME

ArcServe: AGENT_AS_JOB_STARTTIME

DB2: AGENT_DB2_JOB_STARTTIME

SAP HANA: AGENT_SAP_HANA_JOB_STARTTIME

RMAN: AGENT_RMAN_JOB_STARTTIME

MSSQL: AGENT_MSSQLDB_JOB_STARTTIME

The `NUMBER_OF_BACKUP_HOURS` is the number of backup hours before the current time.

For example the following two lines in `dpa.config` should make the DPA Agent start collecting data from the 14 days previous:

```
AGENT_AXION_JOB_STARTTIME=336
export AGENT_AXION_JOB_STARTTIME
```

3. Start the DPA Agent.

For Windows

Procedure

1. Export the key system registry to the registry path `HKEY_LOCAL_MACHINE\SOFTWARE\emc\DPA\AGENT` with the following information:

```
VARIABLE_NAME=NUMBER_OF_BACKUP_HOURS
```

Where `VARIABLE_NAME` is the following for these backup applications:

NetWorker : `NSR_JOB_STARTTIME`

Avamar : `AXION_JOB_STARTTIME`

TSM : `TSM_JOB_STARTTIME`

HPDP: `DP_JOB_STARTTIME`

CommVault: `CV_JOB_STARTTIME`

NetBackup: `NB_JOB_STARTTIME`

ArcServe: `AS_JOB_STARTTIME`

DB2: `DB2_JOB_STARTTIME`

SAP HANA: `SAP_HANA_JOB_STARTTIME`

RMAN: `RMAN_JOB_STARTTIME`

MSSQL: `MSSQLDB_JOB_STARTTIME`

The `NUMBER_OF_BACKUP_HOURS` is the number of backup hours before the current time.

For example, add the following 3 lines as a contents of the `avamar.reg` file and start it from cmd to export to registry so that the DPA Agent collects data from NetWorker starting from 14 days previous:

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\emc\DPA\AGENT]
NSR_JOB_STARTTIME="336"
```

2. Install and start the DPA Agent.

Configure DPA to show all VM Image Backups in Avamar

By default, the Backup Job Config and Backup Server Mapping data sources show the VM Image Backup clients that are being actively backed up in the last 30 days only in DPA for Avamar. Complete the procedure below to show all the VM clients configured in Avamar server.

Procedure

1. Set the option `Show potentially disabled VM clients and HLE conatiners` to `True` on the Avamar Configuration request.
2. On the agent monitoring the Avamar server, set the environment variable `AGENT_AXION_DATASET_BACKUP_DAYS` to 18000.

This value can be used to override the default cut-off threshold of 30 days.
3. Depending on your OS on which you are running your DPA system, follow the steps outlined in the subtasks below.

For Linux

Procedure

1. Edit the `<install_path>/dpa/agent/etc/dpa.config` file to add the following two lines:


```
AGENT_AXION_DATASET_BACKUP_DAYS=18000
export AGENT_AXION_DATASET_BACKUP_DAYS
```
2. Restart the DPA Agent service.

For Windows

Procedure

1. Open regedit to create a registry key named `AXION_DATASET_BACKUP_DAYS` under the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\EMC\DPA\AGENT`.
2. Route to `HKEY_LOCAL_MACHINE\SOFTWARE\EMC\DPA\AGENT` .
3. Right-click to create **New > String Value** with name `AXION_DATASET_BACKUP_DAYS`.
4. Modify the registry key `AXION_DATASET_BACKUP_DAYS` with value 18000.

Installing by using command line installation

Use the appropriate command line.

Before you begin

If you are installing DPA on any of the UNIX OSes, run the `chmod 755` command to change the binary execute permission.

About this task

- Linux

```
DPA-<component>-Linux-<architecture>-<version>.xxx.install.bin [option]
```

where *option* is one of the options listed for a silent or an interactive installation in Table 7.

For example: `DPA-Agent-Linux-x86_64-6.5.0.1.bin -i silent -DUSER_INSTALL_DIR="/opt/custom/emc/dpa"`

- AIX

```
./DPA-<component>-AIX-<architecture>-<version>.bin
```

For example: `./DPA-Agent-AIX-PPC64-6.5.0.1.bin`

- Windows

```
DPA-<component>-Windows-<architecture>-<version>.xxx.install.exe [option]
```

where *option* is one of the options listed for a silent or an interactive installation in Table 7.

For example: `DPA-Agent-Windows-x86_64-6.5.0.1.exe -i silent -DUSER_INSTALL_DIR="C:\custom\emc\dpa"`

Ensure that you carry out the steps provided in [DPA postinstallation steps](#) on page 60.

Table 8 Installer command line options

Option	Description
-?	Displays help text
-i [swing console silent]	Specify the user interface mode for the installer: swing - Graphical interface console - console only silent - no user interaction
-D <name>="<value>"	Shows the installer name-value pairs that might be set on the command line (using the -D option) to override default installer values, or placed in a response file and used with the -f option. Quotes must be used around the value. Example: -D<variable name>="<value>" Where: For example: <code>DPA-Agent-Windows-x86_64-6.5.0.1.exe -i silent -DUSER_INSTALL_DIR="C:\custom\emc\dpa"</code> <variable name> and <value> descriptions are included in the following tables.

Table 9 Datastore installer variables

Variable Name	Description	Possible Values	Default Values
USER_INSTALL_DIR	Installation location	Valid Path	Windows: C:\Program Files\EMC\DPA Linux: /opt/emc/dpa
CHOSEN_INSTALL_SET	Installation set	DS	N/A
VAR_INSTALL_SERVICE	Advanced option to install the Datastore Service	TRUE/FALSE	TRUE
VAR_START_SERVICE	Advanced option to start/stop the Datastore service	TRUE/FALSE	TRUE
VAR_DATASTORE_DATA_LOCATION	Advanced Datastore layout option to specify Datastore server data directory for optimizing performance	Valid Path	\$(USER_INSTALL_DIR)\services\datastore\
VAR_DATASTORE_XLOG_LOCATION	Advanced Datastore layout option to specify Datastore server Xlog directory for optimizing performance	Valid Path	\$(USER_INSTALL_DIR)\services\datastore\data\
VAR_USERNAME (LINUX only)	Advanced option to specify an existing UNIX user account to install the Datastore service	Existing username	N/A
VAR_DATASTORE_BIND_ADDRESSES	IP Address for Postgres to listen on	Valid IP Address	N/A
VAR_DATASTORE_CLIENTS_ADDRESSES IP Address of	IP Address of Application server(s) which will connect to the Datastore service	Valid IP Addresses separated by ", "	N/A
VAR_APOLLO_USER_PASSWORD	DPA Datastore password	[Set at installation or reset using DPA CLI.]	N/A

Table 10 Datastore Advanced options Replication variables

Variable Name	Description	Possible Values	Default Values
VAR_DATASTORE_REPLICATION	Role for Datastore replication	MASTER/SLAVE	N/A
VAR_DATASTORE_REPLICATION_	The IP Address of Master or Slave. If VAR_DATASTORE_REPLICATION_ROLE is set as "MASTER", then the Slave's IP Address needs to be entered and vice versa when VAR_DATASTORE_REPLI	Valid IP Address of Master or Slave	N/A

Table 10 Datastore Advanced options Replication variables (continued)

Variable Name	Description	Possible Values	Default Values
	CATION_ROLE is set as "SLAVE "		

Table 11 Datastore Agent variables

Variable Name	Description	Possible Values	Default Values
VAR_AGENT_APPLICATION_ADDRESS DPA Server FQDN or IP Address to manage the Datastore Agent	DPA Server FQDN or IP Address to manage the Datastore Agent In case of linux IPv6, <IPv6Address> %<Interface_Id_Of_Datastore_Agent>	Valid IP Address or hostname	For multiple application servers and for cases where the datastore service is communicating with linux IPv6 application server(s), this value will be empty. Otherwise the default value is the same as VAR_DATASTORE_CLIENTS_ADDRESSES
VAR_AGENT_START_SERVICE	Advanced option to start/stop Datastore Agent after install	TRUE/FALSE	TRUE
VAR_AGENT_ORACLE_DIRECTORY	Advanced option used for monitoring Oracle by the Datastore Agent. Path where the Oracle Database device driver files can be found	Valid Path	N/A

Table 12 Application installer variables

Variable Name	Description	Possible Values	Default Values
USER_INSTALL_DIR	Installation location	Valid Path	Windows: C:\Program Files\EMC\DPA Linux: /opt/emc/dpa
CHOSEN_INSTALL_SET	Installation set	APP	N/A
VAR_INSTALL_SERVICE	Advanced option to Install the Application Service	TRUE/FALSE	TRUE
VAR_START_SERVICE	Advanced option to start/stop the Application service after installation	TRUE/FALSE	TRUE
VAR_APPLICATION_DATASTORE_ADDRESS	IPAddress of the Datastore server	Valid IP Address where Datastore service is installed and running	N/A
VAR_ADMIN_PASSWORD	DPA Application administrator password	[Set at installation or reset using DPA CLI.]	N/A

Table 12 Application installer variables (continued)

Variable Name	Description	Possible Values	Default Values
VAR_APOLLO_USER_PASSWORD	DPA Datastore password	[Set at installation or reset using DPA CLI.]	N/A

Table 13 Application server Agent variables

Variable Name	Description	Possible Values	Default Values
VAR_AGENT_APPLICATION_ADDRESS	DPA Server FQDN or IP Address to manage the Application server's Agent	Valid IP Address or hostname	127.0.0.1
VAR_AGENT_START_SERVER	Advanced option to start/stop the Application server's Agent after install	TRUE/FALSE	TRUE
AVAR_AGENT_ORACLE_DIRECTORY	Advanced option used for monitoring Oracle by the Application server's Agent. Path where the Oracle Database device driver files can be found	Valid Path	N/A

Table 14 Application server Cluster Advanced option variables

Variable Name	Description	Possible Values	Default Values
VAR_APPLICATION_ADDRESS	The IP Address used by the Application server to announce itself to other DPA application nodes.	Valid IP Address	N/A
VAR_APPLICATION_CLUSTER_ROLE	Role of the application node in a cluster	MASTER/SLAVE	N/A
VAR_APPLICATION_MASTER_ADDRESS	If VAR_APPLICATION_CLUSTER_ROLE="SLAVE", this value needs to be entered.	Valid IP Address	N/A
VAR_APPLICATION_REPORT_DIRECTORY	Path to the network shared report folder	Valid path	N/A
VAR_APPLICATION_REPORT_USERNAME	The user who will be owning the Application service and has permissions to the shared report folder	Existing DOMAIN\ \Username for windows existing username for UNIX	N/A
VAR_APPLICATION_REPORT_PASSWORD (Windows only)	The password of the above user		N/A

Table 15 Standalone Agent Installer variables

Variable Name	Description	Possible Values	Default Values
USER_INSTALL_DIR	Installation location	Valid Path	Windows: C:\Program Files\EMC\DPA Linux: /opt/emc/dpa
VAR_AGENT_APPLICATION_ADDRESS	DPA Server FQDN or IP Address to manage this Agent Valid IP Address or hostname.	In case of linux IPv6, <IPv6Address> %<Interface_Id_Of_Agent >	N/A
VAR_AGENT_START_SERVICE	Advanced Option to start/stop the Agent after install	TRUE/FALSE	TRUE
VAR_AGENT_ORACLE_DIRECTORY	Advanced option used for monitoring Oracle. Path where the Oracle Database device driver files can be found	Valid Path	N/A

DPA postinstallation steps

After you install or upgrade DPA and access the DPA web console, a message is displayed that indicates the DPA Server the status of the initialization process. The initialization process can take approximately 10 minutes to complete.

About this task

During the initialization time, DPA is creating the database schemas, tables, views, and the DPA Datastore. It also creates the various system reports and dashboards templates, the default system users, Analysis Engine Rulesets, and various other default and initial objects. Your network connection time affects the speed at which all these actions complete. Ensure that you perform the following steps after installing DPA.

Procedure

1. If you have upgraded or migrated to DPA 18.1, delete the browsing history/cache in your browser before using DPA 18.1.
2. (Optional) Carry out the following steps to verify whether initialization is still in progress or completed:
 - a. If you installed on Linux and the install is done to a non-default location, log out and back in to the session. Alternatively, run from a new login window.
A new shell is required for the executive command paths to be found before running `dpa.sh svc status`.
 - b. On the DPA Application server, go to `<install_dir>\services\applications`.
 - c. Check the `*.rar` ; `*.ear`, and `*.war` files for `*.deployed`, `*.isdeploying`, or `.failed` extensions.
 - If files have an extension of `*.isdeploying`, then server initialization is still in progress.
 - If files have an extension of `*.deployed`, then server initialization is complete and you can login to the DPA web console.
 - If files have an extension of `*.failed`, then server initialization failed; contact Technical Support.

3. Start the web console to verify successful DPA installation.

All DPA services must be running when you launch the web console. The Adobe Flash plugin in your web browser is required to launch the web console.

- a. Start a browser and connect to DPA Server over https on port 9002. Ensure that all pop-up blockers are disabled. For example:

```
https://<server_name>:9002
```

where *<server_name>* is the name or IP address of the server or localhost.

Alternatively, use

```
https://<server_name>:9002/flexui url
```

. if you choose to continue using the Flex-based DPA 18.1 web console.

- b. Type the username and password. Username and password fields are case-sensitive.
c. Click **Login**

4. Add licenses to the DPA server.

The DPA server is installed with a 90-day temporary license.

If you are upgrading and you are not adding capacity or changing to new DPA 18.1 functionality, no licensing changes are needed.

The CLP license is required for new DPA 18.1 functionality and increased capacity on a DPA instance. If you are migrating from DPA version 5.x to DPA 18.1, the existing licenses are migrated with your configuration and data. [CLP and WLS license coexistence in DPA](#) on page 72 provides more information.

If you are adding CLP licenses, ensure that you select license files with the .lic file extension.

If you are adding WLS licenses, select license files with the .wls file extension.

After you install the license file, the DPA web console prompts you to close so it can register the license file.

5. Log back in to the DPA web console.
6. (Recommended) If you added CLP licenses in step 4, register the DPA Application server with the ESRS-VE. This registration process enables Customer Support to service the DPA instance.

Observe the following:

- If you are upgrading a previously registered ESRS, it is possible that ESRS will show that it is already registered with the following error:
[ERROR] This node is already registered with an EMC Secure Remote Support Service.

Then, ESRS shows that host IP is not available anymore with the following errors:

```
[ERROR] This node failed to delete with EMC Secure Remote Support Service.
```

```
Offline: Validation error
```

Search EMC Knowledgebase article xxxxxx, available on <http://www.support.emc.com>, for more information. This is an environment issue not related to DPA.

- Registering ESRS after a fresh installation requires that an ESRS-VE be already installed and reachable from the DPA Application server. If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at https://support.emc.com/downloads/37716_EMV-Secure-Remote-Services-Virtual-Edition on EMC Online Support provides more information on ESRS-VE installations. The *Data Protection Advisor Software Compatibility Guide* provides supported ESRS-VE module and version information.
- Register a single Application service. The registration includes both DPA Datastore and Application servers.
- If you are working in a clustered environment, register the Master Application server with ESRS. Use the `dpa app con` command to check if your Application server is Master or Slave server. The CLI section provides more information.
- When prompted for EMC Secure Remote Support username and password, provide EMC online support credentials for registration. For example:

```
dpa app support --register 10.11.110.111
Dell EMC Data Protection Advisor
Enter Data Protection Advisor Administrator username :
Enter Data Protection Advisor Administrator password :
Enter EMC Secure Remote Support username :
Enter EMC Secure Remote Support password :
```

- Note the following: In a clustered environment, do not use the Application server registered with ESRS for scheduled reports. Any problems with the scheduled reports or data collection on the listener are propagated across the Application servers in the cluster.
 - a. Log in to the Application server using Remote Desktop Connection for Windows or PuTTY for Linux.
 - b. Type the `dpa app support --register ESRS_IP` command to register a DPA server.

Where *ESRS_IP* is the IP address of the ESRS Gateway. For example:

```
C:\Program Files\EMC\DPA\services\bin>dpa app support --register
10.11.110.111
```

- c. When prompted, type the EMC Secure Remote Support username and password.

Output appears that indicates that the request to register the DPA server with IP address that you typed is approved and the command is successful.

7. (Recommended) If you registered the DPA Application server with the ESRS-VE in step 6, enable Health Service on the DPA Application server. On the DPA Application server, type:
 - a. `$ dpa health install`
 - b. `$ dpa health start`
8. (Optional) If you want to configure alerting on Replication Monitoring, ensure that you create Recoverability rules to the Analysis Policy and assign the rules to the desired object. Go to **Policies > Analysis Policies**
9. (Optional) If you have upgraded from a previous 6.x version and you would like to display the Data Domain Overview dashboard and the Data Domain Details dashboard:
 - a. Go to **Dashboard > + icon > Open Existing Dashboard**.

The **Open Existing Dashboard** window appears.

- b. Select **Data Domain** and Click OK.
10. (Optional) If you are monitoring a Data Domain OS 5.7 and later and would like to ensure configuration of Physical Capacity Reporting data collection:
 - a. Manually assign the request to any Data Domain OS 5.7 boxes.
 - b. Run the request so that the statistics are gathered on the Data Domain and the schedule is created. Then, when you are ready to run the first report, data is returned.

Encryption of the DPA Application server

To encrypt the information flowing between the Application server and the DPA web console, you must install a certificate on the Application server.

Encrypting the DPA Application server

About this task

Out of the box, the information that flows between the DPA Application server and the DPA web console is encrypted using the self-signed certificate that is included with the DPA Application server. This certificate is generated during install as well as the key store password.

Before you begin:

- Ensure that you have requested and obtained a Trusted Certificate and private key for the Application server from a CA.
- Ensure that you have merged the Trusted Certificate and the private key inside a keystore file. Refer to CA vendor documentation for information.
- If you are implementing Application server clustering, ensure that you complete all cluster configuration before enabling encryption on the Datastore and Application servers.

Procedure

1. Use the `dpa app impcert -kf` command to import the self-signed certificate:

```
dpa app impcert -kf "C:\work\new.keystore" -al newkey -pw password
```

This is the password of the newly generated keystore file. This password is located at `C:\work\new.keystore`.

2. Restart the DPA Application service. The `dpa app --help` command provides additional information.
3. (Optional) Install the certificate in those browsers that you use to access DPA. Follow the instructions of your chosen browser.

It may take a few minutes on initial connection to open DPA when using a secure connection.

Encrypting Application Server cluster

To encrypt Application Server clusters, you must have one domain (wildcard) certificate from the trusted certificate vendor. Install this certificate on all the DPA Application cluster nodes.

About this task

You should not install an individual certificate for each Application node in the cluster.

Configuring antivirus software with DPA

Configure the following antivirus configuration. Refer to your particular antivirus software documentation for information on how to configure the software so that there is no real-time monitoring of these processes or monitoring of the files that they read.

About this task

It is not necessary to have all DPA file systems monitored by antivirus software, and scanning certain file systems and processes can potentially degrade overall performance due to the impact of increased disk IO activity.

Procedure


1. Exclude the following files and processes from antivirus monitoring.

If you are configuring antivirus software on Linux, the following file names will not have a `.exe` extension.

- DPA Application Server:
 - `<install_dir>services\executive\wrapper.exe`
 - `<install_dir>\agent\bin\dpaagent.exe`
 - `<install_dir>\services_jre\bin\java.exe`
- DPA Datastore Server:
 - `<install_dir>\services\datastore\engine\bin\postgres.exe`
 - `<install_dir>\agent\bin\dpaagent.exe`

2. Exclude the following specific directories from being monitored by your antivirus software.

- DPA Application Server:
 - `<install_dir>\services\standalone**`
 - `<install_dir>\services\tmp**`
 - `<install_dir>\services\shared**`
- File space on the DPA Datastore Server:

 **Note:** If you selected advanced file system layout during Datastore installation, then alternative directories may be used instead of the following defaults.

 - `<install_dir>\services\datastore\data**`
 - `<install_dir>\services\datastore\data\pg_log**`

Upgrades

You can upgrade from previous DPA releases to DPA 18.1 and minor releases. The *Data Protection Advisor Release Notes* provide information on supported upgrades.

Note that the DPA 18.1 upgrade installer does not provide the option to use TLS protocol version 1.2 only. Additionally, DPA retains your existing TLS protocol version settings after upgrade. You can change the TLS protocol version to 1.2 only after upgrade. [Setting TLS protocol version 1.2 only after installation or upgrade](#) provides information.

Upgrade prerequisites

There are a set of recommended best practices before you carry out an upgrade of the DPA server.

- Back up the DPA Datastore by using the `dpa ds export` command. [Backup of the Datastore](#) on page 121 provides information. The DPA Installer prompts you to do this.
- For Datastore and Application server upgrades, the DPA Agent on those servers is also upgraded as part of the server upgrade. You must carry out a separate upgrade for a DPA Agent in the case of standalone DPA Agents only.
- To ensure secure communication between the DPA Server and Agent, set the Agent registration password using the `dpa app agentpwd` CLI command on the DPA Application Server host. You must also set this password on all DPA Agent hosts. [dpa application agentpwd](#) provides information. Then restart the Application service. Ensure that you set this password for each Agent. The exception to this is if you are concurrently running DPA Agents previous to version 6.5 along with the upgrade to version 18.1 Agents. [Upgrading DPA Agents previous to version 6.5 alongside DPA version 6.5 Agents and version 6.5 Server](#) on page 67 provides information.
- Take note of the previous DPA 6.x build installed on your system by running `dpa app ver` and recording the output. This output is important when verifying package installation.
- Stop the DPA Application server. Good practice is to perform a complete backup of the host running DPA Application server.
- Stop the DPA Datastore. Good practice is to perform a complete backup of the host running DPA Datastore server.
- If your infrastructure is running on VM, stop the DPA Application and Datastore servers and take a snapshot of the DPA Application and Datastore servers to facilitate restoring them in case of upgrade problems.
- Clear the browser cache.
- Ensure that you have admin/root privileges.
- If upgrading on UNIX/Linux, ensure that the `unzip` command for InstallAnywhere is installed on your system.
- When upgrading or installing patches in clustered environments, stop the DPA Application service on all servers. Upgrade the Datastore first, and then upgrade the Application servers. You must stop the Application service because when the services are on separate machines, the installer cannot stop the services.
Start the upgraded DPA Application. Confirm initialization completed and that you can login to the DPA web console before upgrading the remaining clustered Application servers.
- In relation to the database upgrade:
 - Ensure that you have 3GB of free space for the database upgrade.
 - Ensure that you are running a LINUX version with a minimum glibc version of 2.12. If your LINUX version is running a glibc version earlier than 2.12, use the procedure provided in [Upgrading DPA with a LINUX version running glibc earlier than 2.12](#) on page 67
- If you are currently using DPA for RMAN reporting through an existing DPA backup license, contact your Account Representative for the DPA for Enterprise Applications license. The DPA for Enterprise Applications license allows you to expand the number of RMAN servers being reported in DPA when you upgrade to DPA 6.3 and minor releases. Enter the DDBEA license into DPA 6.3 and minor releases after installation. The *DPA 6.2 Release Notes* provides more information on the license is for DDBEA.

- If you are upgrading from DPA 6.1, ensure that you review and edit the retention period on collection requests to match organizational policies before upgrading. Data collection requests contain a different default retention period in DPA 6.1.

Upgrading DPA

Use this procedure to upgrade DPA if you do not have clusters or Datastore Replication configured, and if the LINUX version you are running has a minimum glibc version of 2.12, as applicable.

Before you begin

Add support for Upgrading installations where the database tablespaces have been configured to reside on different filesystems.

- Ensure that you carry out the prerequisites in [Upgrade prerequisites](#) on page 65.
- Ensure that you run the installer as admin/root user.

About this task

If you are running a LINUX version that has a glibc version earlier than 2.12, follow the procedure provided in [Upgrading DPA with a LINUX version running glibc earlier than 2.12](#) on page 67

Procedure

1. If you have not already done so, shut down the Application Service.
2. Upgrade the Datastore. Follow the installation steps as directed in the Installer. Ensure that the existing DPA installation directory is specified correctly.

You must install the DPA update package in the same installation directory as your existing DPA package.

3. Upgrade the Application server. Follow the installation steps as directed in the Installer. Ensure that the existing DPA installation directory is specified correctly on the installer.

You must install the DPA update package in the same installation directory as your existing DPA package.

4. If you are upgrading on Linux, and if you are running the Linux DPA Application as non-root user, carry out the following steps to restore the required FS permissions and owners:

- a. Transfer ownership of DPA services installation directory to the OS user, which will run the DPA Application. Type: `chown --dereference -LR`

`<user_to_run_dpa>:<group_of_user> <dpa_install_dir>/services`

- b. Modify the `<dpa_install_dir>/services/executive/applnsvc.sh` file.

Change the line `RUN_AS_USER=` to `RUN_AS_USER=<user_to_run_dpa>`

5. Restart the DPA web console.
6. Wait for the files to be deployed under the installation folder.

In Windows: `C:\Program Files\EMC\DPA\<install_dir>services\applications`

In Linux: `/opt/emc/dpa/services/applications`

The DPA web console UI splash page displays upgrade status.

7. Carry out the steps provided in [DPA postinstallation steps](#) on page 60.

Upgrading DPA Agents

Procedure

1. Shut down the DPA Agent service.
2. Upgrade the Agent using Agent Installer suitable for your OS. Follow the installation steps as directed in the Installer.

You must install the Agent update package in the same installation directory as your existing DPA package.

Consider that during the upgrade, the Agent is stopped and as such, requests that are holding during the upgrade can be failed. After upgrade finishes, the DPA Agent continues to work normally.

Upgrading DPA Agents previous to version 6.5 alongside DPA version 6.5 Agents and version 6.5 Server

You may need to run versions of the DPA Agent previous to 6.5, which do not support the Agent password. In these situations if you set the Agent registration password on the DPA Server then all the previous version DPA Agents which do not support the Agent password fail to connect. Follow the procedure below to avoid this situation.

About this task

You may need to run DPA Agent versions previous to version 6.5 for collecting for systems which are no longer supported, or you may have so many agents that you cannot upgrade them all at once.

Procedure

1. Upgrade the DPA server to version 6.5.

Do not set the Agent registration password. 3. Do not uninstall the old version of agent then install 6.5.

2. Upgrade the DPA Agents which require an upgrade to version 6.5 following the normal the upgrade process.

In the process of upgrading the DPA Agent it will not request that an Agent password be set. This is different from a fresh installation, which would request that an Agent password be set.

Upgrading DPA with a LINUX version running glibc earlier than 2.12

Before you begin

- Ensure that you carry out the prerequisites in [Upgrade prerequisites](#) on page 65.
- Ensure that you run the installer as admin/root user.

Procedure

1. Stop the Application Service.
2. Export the Datastore. [Backup of the Datastore](#) on page 121 provides information.
3. Install a new Datastore with the latest version of DPA and a version of LINUX that is running glibc version 2.12.
4. Import the existing Datastore to the newly installed Datastore with the latest version of DPA and the supported version of LINUX with glibc version 2.12

5. Point the DPA Application server to the newly installed and imported Datastore. Run: `dpa app configure --master <datastore_ip>`
6. Upgrade the Datastore. Follow the procedure provided in [Upgrading DPA](#) on page 66.

Upgrading existing clusters

Use this procedure to upgrade an already existing cluster.

Before you begin

- Ensure that you carry out all the steps provided in [Upgrade prerequisites](#) on page 65.
- If you are running UNIX machines, ensure that you are a root user.
- Stop the load balancer on the DPA Application and Datastore servers. The command to stop the load balancer varies by OS. Refer to your OS documentation for information.

Procedure

1. Stop the Application service on the cluster Application nodes:
 - a. Stop the Slave Application server.
 - b. Stop the Master Application server.

Run:

```
# dpa app stop
```
2. Upgrade the DPA Datastore server:
 - a. Launch the DPA installer and follow the prompts.
 - b. Ensure that the Datastore has installed and started successfully. [DPA postinstallation steps](#) on page 60 provides information.
3. Upgrade the Master Application node:
 - a. Launch the DPA installer and follow the prompts.
 - b. Wait for the Application service to start. Verify that the `server.log` file includes output such as `DPA master started successfully`.
4. Upgrade the Slave Application nodes:
 - a. Launch the DPA installer and follow the prompts.
 - b. Wait for the Application service to start. Verify that the `server.log` file includes output such as `DPA slave started successfully`.
5. Restart the load balancer application on the DPA Application and Datastore servers. The command to start the load balancer varies by OS. Refer to your OS documentation for information.

Upgrading with Datastore Replication enabled with DPA 6.3 and later

To upgrade with Datastore Replication enabled follow the following procedure:

Before you begin

- Ensure that you have carried out all the steps provided in [Upgrade prerequisites](#) on page 65.
- If you are running UNIX machines, ensure that you are a root user.
- Ensure that all processes in each step are complete before starting the process in the next step.

Procedure

1. If you have not already done so, on the Application servers stop the Application Service.
Run:

```
# dpa app stop
```
2. Upgrade the Slave Datastore.
Launch the DPA installer and follow the prompts.

If you are implementing Cascading Replication, upgrade the Datastore at the end of the chain first.
3. Upgrade the Master Datastore.
Launch the DPA installer and follow the prompts.
4. Upgrade the Application server(s).
Launch the DPA installer and follow the prompts.
5. Verify that Datastore Replication is running. Run:

```
# dpa ds rep
```


Output should show STREAMING.

Upgrading with Datastore Replication enabled with DPA versions earlier than 6.3

Upgrading with Datastore Replication is automated and does not require user interaction, except when upgrading the Replication Slave Datastore.

Before you begin

- Ensure that you have carried out all the steps provided in [Upgrade prerequisites](#) on page 65.
- If you are running UNIX machines, ensure that you are a root user.
- Ensure that all processes in each step are complete before starting the process in the next step.

Procedure

1. Stop all services:
 - a. Run `# dpa app stop` on the Application Server.
 - b. Run `# dpa ds stop` on the Master Datastore.
 - c. Run `# dpa ds stop` on the Slave Datastore.
2. Upgrade Master Datastore:
 - a. Launch the DPA installer and follow the prompts.
 - b. Verify that Datastore Replication is running. Run: `# dpa ds rep`
3. Create a copy of the Master Datastore. Type: `dpa ds rep -e <empty_dir>`
4. Uninstall the existing Slave Datastore.
5. Install a clean Datastore Server with the same install location as the Master Datastore, and configure the newly installed Datastore Server as a Slave Datastore. Type: `dpa.sh ds rep --role SLAVE <IP of master>`.

Do not start or stop services.

6. Initialize the Slave Datastore from Master copy. Type: `dpa ds rep -i <master_copy>`
7. Start the Slave Datastore.
8. Upgrade the Application Server.

Upgrading with Datastore Replication and existing clusters

Use this procedure to upgrade a system with Datastore Replication and existing clusters.

Before you begin

- Ensure that you carry out all the steps provided in [Upgrade prerequisites](#) on page 65.
- If you are running UNIX machines, ensure that you are a root user.
- Stop the load balancer on the DPA Application and Datastore servers. The command to stop the load balancer varies by OS. Refer to your OS documentation for information.

Procedure

1. If you haven't already done so, stop the Application service on the cluster Application nodes:
 - a. Stop the Slave Application server.
 - b. Stop the Master Application server.

Run:

```
# dpa app stop
```

2. Carry out the steps provided in [Upgrading with Datastore Replication enabled with DPA 6.3 and later](#) on page 68.

If you are upgrading with a DPA version previous to 6.3, carry out the steps provided in "Upgrading with Datastore Replication enabled with DPA versions earlier than 6.3."

3. Upgrade the Slave Application nodes:
 - a. Launch the DPA installer and follow the prompts.
 - b. Wait for the Application service to start. Verify that the `server.log` file includes output such as `DPA slave started successfully`.
4. Restart the load balancer application on the DPA Application and Datastore servers. The command to start the load balancer varies by OS. Refer to your OS documentation for information.

CHAPTER 3

Administering DPA

This chapter includes the following sections:

- [License management](#)..... 72
- [Users and security](#)..... 73
- [System settings](#).....84
- [Application service administration](#)..... 115
- [Datastore service administration](#)..... 121
- [DPA command line operations](#)..... 127

License management

This section describes license management in DPA.

Evaluation license bundled with DPA

DPA is bundled with a 90-day evaluation license.

The evaluation license is created from the time of DPA installation, is valid for up to 90 days, and allows access to all features. If you import a license during 90-day evaluation license period, the evaluation license is removed and you have access to DPA features according to license you imported.

Licensing types in DPA

DPA uses the *Common Licensing Platform (CLP)* license type.

The CLP license coexists with and, in certain circumstances, replaces the legacy *Wysdm Licensing System (WLS)* license type that was previously used with DPA before the product name was changed to DPA.

CLP and WLS license coexistence in DPA

The CLP license is required for DPA functionality.

If you are not adding capacity or changing to DPA functionality from a version of DPA later than DPA 6.2, import of CLP licenses is not required. However, if you are upgrading to the latest version of DPA from a version of DPA previous to DPA 6.2, contact licensing@emc.com after upgrade or migration to assist you with legacy license transition to CLP licenses of all your WLS licenses. If you are migrating from DPA version 5.x to the latest version of DPA, the existing licenses are migrated with your configuration and data. You need to add CLP licenses only for the latest version of DPA functionality or for increasing current license capacity.

CLP licenses work on a replacement model. When you import a CLP license, the CLP license replaces all the existing licenses of the same type. Additionally, the base and Enterprise license functionality is moved into each CLP license. You must be aware of the existing license count when you order CLP licenses of the same type, then add on the new capacity required and order for the total. For information on purchasing licenses for your DPA installation, contact your Account Representative.

A system that has been migrated or upgraded from a previous version of the former DPA will contain WLS licenses. WLS and CLP can coexist only where they aren't for the same functionality.

Expired licenses

If a license expires, a license violation warning appears in the report title for reports run from all objects enabled by the expired license. In addition, new objects cannot be added in the web console for module components enabled by an expired license.

License removal

Removing a license causes a license violation warning to appear when running reports against objects for that license. New objects of that type cannot be added in the web console until a replacement license is supplied.

If you are using temporary licenses that have an expiration date, the License Expiration dialog appears to notify you of the expiration of your temporary licenses. Permanent licenses do not display.

Adding new licenses

Go to **Admin > System** and then click **Manage Licenses**.

About this task

Disabling automatic temporary licence expiration pop-up

Go to **User Properties > Show License Expiration** and uncheck the box.

Users and security

User accounts

Four default users are supplied by default within DPA: Administrator, Application Owner, Engineer, and User.

The Administrator account is the only account active after DPA installation. The user sets the Administrator account password during the DPA installation process.

The Administrator must set passwords for the other default user accounts before they can be used to access DPA. If the Administrator does not set passwords for the other user accounts, they remain in a disabled state.

Managing users

The DPA Administrator can manage user accounts in the **Manage Users** section. Go to **Admin > Users & Security > Manage Users**. In this section the Administrator is allowed to create, edit, view and delete user accounts.

Creating a new user account

Procedure

1. Go to **Admin > Users & Security > Manage Users**.
2. Click **Create User**.
Alternatively, select an existing user and click **Save As** to create a copy of an existing user.
3. In the **Create User Properties** tab, update the information in the respective tabs:
 - a. In the **User Properties** tab, specify the name, logon name, role, authentication type and password.
 - b. If the user is to be authenticated by using LDAP, choose the LDAP authentication type.
 - c. In the **Report Preferences Preferences and Appearance** tabs, assign preferences and appearance settings. Note that the role you assign to the user determines which areas of DPA they can access.

- d. Click **OK** to confirm the settings.
- 4. Click **Close**.

Editing and deleting user accounts

The DPA Administrator can edit or delete any DPA user account except the default Administrator account.

Procedure

1. Go to **Admin > Users & Security > Manage users**
2. Select the user you would like to edit or delete.
 - Click **Edit** to customize desired items such as the user's name, role, password, or report and appearance preferences.
 - Click **Delete** and **Yes** to delete it.
3. Click **Close**.

Changing user account passwords

The DPA Administrator can change user account passwords in **Manage Users**. Non-Administrator users can change their password in **View User Properties** by clicking the gear icon on the top-right corner of the DPA web console.

Procedure

1. Go to **Admin > Users and Security > Manage Users**.
2. Select the user account for which you wish to change the password and click **Edit**.
3. Go to **Edit User Properties** and set the **Authentication Type** to **Password**.
4. Type the new password in the **Password** field, and then retype the password in **Confirm Password** field.

Note the following regarding DPA passwords:

- Blank passwords are not supported.
 - Minimum length is 9 characters.
 - The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character
5. Click **OK**.

Security Settings

You can configure user security settings. Go to **Admin > Users & Security**

Table 16 Password Policy

Setting	Description
Minimum number of characters	The minimum number of characters required for the password for the DPA web console. The minimum value is 9. There is no maximum limit of characters. DPA supports only Latin characters.

Table 16 Password Policy (continued)

Setting	Description
Must use uppercase characters	Requires uppercase characters in the password for the DPA web console. Enabled by default.
Must use lowercase characters	Requires lowercase characters in the password for the DPA web console. Enabled by default.
Must use special characters	Requires special characters in the password for the DPA web console. Enabled by default.
Must use numeric characters	Requires numeric characters in the password for the DPA web console. Enabled by default.

Table 17 Password History Policy

Setting	Description
Password History	Enables the limiting of password history.
Limit of password history	<p>The number of times that DPA allows the user to specify an identical password for that user from the previous password. The default value is 1. The maximum value is 10. Enabled by default.</p> <p>If Limit of password history is left at 1, that user cannot change the password to the current one.</p> <p>If Limit of password history is set to greater than 1, that user cannot change password to the current one and to the previous one. DPA displays a message indicating that the previous password was already configured and that the user must specify a new password.</p>

Table 18 Login Limit

Setting	Description
Login limit	Enables the limit for number of attempts to log in to the DPA web console.
Limit of login attempts	The number of attempts that DPA allows the user to log in to the DPA web console. The default value is 5. The range is 1–10.
Lockout timeout	The amount of time that DPA temporarily locks the user out of the DPA web console after exceeding the specified login limit. The default value is 3 minutes. The range is 1–60 minutes.

Table 19 Password Expiration

Setting	Description
Password Expiration	Enables password expiration. The default value is off.
Password expiration period (days)	The period in days that the DPA password is valid. The default value is 90. The maximum is 180.

User roles and privileges

Roles are used to handle the privileges that are allowed for users. Users gain their privileges by being assigned to the appropriate role.

Four roles are supplied by default within DPA: Administrator, Application Owner, Engineer, and User. The default user roles are set and cannot be changed.

The following table explains default role privileges.

Table 20 User roles

User roles	Privileges
Administrator	Can perform all configuration and reporting functions.
Application owner	Can perform all reporting functions and modify credential settings.
Engineer	Can perform all reporting functions and most configuration functions. Engineers cannot create or modify users or user roles, or modify system settings.
User	Can perform reporting functions only.

Creating a new user role

The DPA Administrator can create a new custom user role with custom permissions and settings.

About this task

Procedure

1. Go to **Admin > Users & Security > Manage Roles**.
2. Either click **Create Role**, or choose an existing role and click **Save As**.
Choose **Save As** to create a copy.
3. In the **User Role Properties** window,
 - a. Type a name and description for the new role in the **Name** and **Description** fields.
 - b. Set the Privileges, Accessible Groups, Dashboards, and Menus.
 - c. Click **OK** to confirm the settings.
4. Click **Close**.

Editing and deleting user roles

The DPA Administrator can edit or delete only custom user roles. Default user roles cannot be edited or deleted. It is not possible to delete a role unless the users within that role have first been assigned to alternative roles.

Procedure

1. Go to **Admin > Users & Security > Manage Roles**.
2. Select the custom user role you would like to edit or delete:
 - Click **Edit** to customize the privileges, accessible groups, dashboards, and menus.
 - Click **Delete** and **Yes** to delete the user role.

Viewing users within user roles

Go to **Admin > Users & Security > Manage Roles**. The DPA Administrator can review the users associated with a user role in the **Manage Roles** tab by selecting a specific user role name. A list of the default roles (administrator, application owner, engineer, user) is displayed, together with any new roles added since installation.

Limiting users to see only specific groups

You can configure a DPA user to be able to see specific groups or backup configuration items when running reports.

Before you begin

The groups must already exist.

About this task

By default, users can see the entire DPA object inventory. However, you may want to limit what certain users see of the DPA object inventory. For example, service providers may have groups configured in their DPA object inventory which correspond to their individual customers. The service providers may want to configure it so that the individual customers see and run reports against only the specific object inventory configured in their customer group when they log in to DPA.

Procedure

1. Go to **Admin > Users & Security > Manage Roles**.
2. Create the custom role that you require, or select the custom user role you would like to edit and click **Edit**.
3. Select the **Accessible Groups** tab.

The list of all available groups is displayed
4. Select the group that will be accessible by the role and click **>** or **>>** to move all groups.
5. Click **OK** to confirm the settings.
6. Click **Close**.

Restricting user groups

You can restrict user groups so that certain user groups or roles can set values for custom attributes without the ability to update system attributes or create or change groups.

Before you begin

- Ensure that you log in to the DPA server as Administrator.

Procedure

1. Create the Read Inventory role:
 - a. Go to **Admin > Users & Security > Manage Roles** and click **Create Role**.
The **User Role Properties** dialog appears.
 - b. Populate the fields accordingly:
In the **Name** field, type the name you would like to give to your role. For example, **Read Inventory**.
In the **Description** field, type a description if you would like to do so.
 - c. In the **Privileges** tab under **Inventory**, select **View existing objects and group management**.
 - d. In the **Accessible Groups**, select the groups that you would like to view, click **Move selected groups** and then click **Close**.
2. Create the Read Inventory user:
 - a. Go to **Admin > Users & Security > Manage Users** and click **Create Role**.
The **Create User Properties** dialog appears.
 - b. Populate the fields accordingly:
In the **Name** field, type the name you would like to give to your user. For example, **Read**.
In the **Logon** field, specify the logon the user should type. For example, **read**.
In the **Role** field, select the one that you created in step 1 for the Read Inventory role.
In the **Authentication** field, select the desired authentication type from the dropdown. If you choose Password, specify and confirm a password.
 - c. Click **OK**.
3. Create the Assign Attribute and Read Inventory Role:
 - a. Go to **Admin > Users & Security > Manage Roles** and click **Create Role**.
The **User Role Properties** dialog appears.
 - b. Populate the fields accordingly:
In the **Name** field, type the name you would like to give to your role. For example, **Assign Attribute and Read Inventory Role**.
In the **Description** field, type a description if you would like to do so.
 - c. In the **Privileges** tab under **Inventory**, select **Assign/unassign attributes**.
The **View existing objects and group management** privilege is selected automatically.
 - d. In the **Accessible Groups** tab, select the groups that you would like to view, click **Move selected groups** and then click **Close**.
4. Create the Assign Attribute and Read Inventory user:
 - a. Go to **Admin > Users & Security > Manage Users** and click **Create Role**.
The **Create User Properties** dialog appears.
 - b. Populate the fields accordingly:
In the **Name** field, type the name you would like to give to your user. For example, **Assign**.

In the **Logon** field, specify the logon the user should type. For example, `Assign`.

In the **Role** field, select the one that you created in step 3 for the Assign Attribute and Read Inventory role.

In the **Authentication** field, select the desired authentication type from the dropdown. If you choose Password, specify and confirm a password.

c. Click **OK**.

After you finish

External authentication, LDAP integration, and binding

DPA supports configuring an external authentication method via the Lightweight Directory Access Protocol (LDAP). DPA supports Microsoft Active Directory and OpenLDAP as LDAP servers

User account passwords are stored in the DPA Datastore only when the internal authentication method is configured. In the external authentication method the passwords are stored in the LDAP server. To enable LDAP authentication, select **Admin > Users & Security > Manage External Authentication**.

DPA supports two LDAP binding methods: anonymous bind and simple bind. To configure anonymous bind, ensure that the **anonymous bind** checkbox is checked in the **Manage External Authentication** tab. For simple bind, ensure that the **anonymous bind** checkbox is unchecked. Also, ensure that the username and password of a user with read base access is set.

LDAP authentication configuration

The following fields are used in the configuration of LDAP Authentication in DPA.

Table 21 LDAP Authentication configuration in DPA

Field	Description
Server	Hostname of the LDAP server. The hostname must be resolvable from the DPA server.
Use SSL	Select this option to connect to the LDAP server using an SSL connection.
Port	Port that the LDAP server listens on for requests: <ul style="list-style-type: none"> port 389 for non-SSL connections port 636 for SSL connections When you use Microsoft Active Directory configured as a Global Catalog server, specify the following in the Manage External Authentication dialog: <ul style="list-style-type: none"> port 3268 for non-SSL connections port 3269 for SSL connections
LDAP Version	Version of LDAP that is used on the server. DPA supports versions 2 and 3.

Table 21 LDAP Authentication configuration in DPA (continued)

Field	Description
Base Name	Location of all possible users. This location will be used as the starting point for all queries against the directory. The value entered must be the Distinguished Name of the base of the directory, for example, DC=eng,DC=company,DC=com.
Identification Attribute	The attribute in LDAP or Active Directory that is used to search for a user account, for example, sAMAccountName (Active Directory) or uid (OpenLDAP)
Anonymous Bind	DPA supports two different LDAP bindings: <ul style="list-style-type: none"> Anonymous Bind - Check the checkbox to connect to the LDAP server with Anonymous Bind Simple Bind - Leave the checkbox unchecked to use Simple Bind. This enables the Username and Password fields
Username	The bind DN of the user on the LDAP server permitted to search the LDAP directory within the defined search base.
Password	User password.
Validate	Click to test user authentication with the LDAP server. A message displays whether or not connection to the LDAP server successfully occurred.

Creating a new user account with LDAP authentication

As the DPA Administrator, once you have configured and tested the LDAP binding you can create or edit the user accounts that need to be authenticated by the LDAP server.

Procedure

1. Go to **Admin > Users & Security > Manage External Authentication**
2. Set the LDAP value in the **Authentication type** field.
3. Provide the user's Distinguished Name (DN) or the Identification Attribute value in the **External Name** field.

With the Active Directory integration, the Identification Attribute value is typically the sAMAccountName. With OpenLDAP, it is typically UID.

Automated user provisioning

Automated user provisioning is available in DPA when it is integrated with an LDAP server. Enabling the Auto Login feature makes it possible for DPA to automatically create a user account when a new user successfully logs in to DPA.

The user role assigned to the new user can be configured in the **Auto Login** tab. The Administrator can configure a default User Role or a role based on LDAP group mapping.

Auto-login—Default user role

When a default user role is set in the Auto login tab, this role is assigned to all new users automatically created by DPA. You can view the complete list of users created with the Auto login feature in the **Manage Users** tab. They will have the value *LDAPAUTO* in the Authentication type field.

Procedure

1. Configure and test LDAP integration in DPA.
2. Go to **Admin > Users & Security > Manage External Authentication > Auto-login Properties > Edit** and flag **Enable Auto-login**.
3. Select a role in the Default User Role drop-down list.
4. Click **OK** to confirm the settings.
5. Click **OK** in the **Manage External Authentication** tab to close.

When you successfully authenticate using Auto-login, DPA automatically creates a user account for you within DPA.

Auto login—LDAP group mapping

As DPA Administrator, you can map specific LDAP groups to DPA user roles in the Auto Login settings.

Procedure

1. Configure Auto login with a default user role.
2. Check the **Enable Group Mapping** checkbox to enable Group Mapping:
 - In the **Group Base** field, specify the Distinguished Name of the group. For example, `cn=users,dc=eng,dc=company,dc=com`
 - In the **Group Attribute** field, specify the LDAP attribute used for the group search. Typically this is either *CN* or *sAMAccountName* for Active Directory or *uid* for OpenLDAP.
 - In the **Group Member Attribute** field, specify the attribute that specifies members of the group. Typically this is either *member* for Active Directory or *memberUid* for OpenLDAP.
3. Click **Add** to add a new line to the **Group Mapping** section.
4. In the **LDAP Group Name**, set the name of the group to map with the user role.
5. In the **User Role**, choose one of the available roles from the dropdown list.
6. Use **Add**, **Remove**, **Up**, and **Down** to organize the Group Mapping.
7. Click **OK** to confirm the settings
8. Click **OK** in the **Manage External Authentication** tab to close.

Group Mapping

The group mapping feature allows DPA to map specified LDAP groups to DPA roles so that you can be assigned different DPA roles depending on the LDAP groups that you are in.

If you are a member of multiple LDAP groups, you are granted the DPA role that is mapped to the first group in the mapping table. Ensure that the LDAP group that maps to a DPA role with greater permissions is highest in the list. A user that is not a member of a group in the Group Mapping list is assigned the Default User Role. **Up** and **Down** buttons are provided to enable table entries to be moved to the desired positions in the table.

Configuring LDAP integration—scenario settings

In the following LDAP integration scenarios, the following settings are used. Note that these settings are example settings only.

Table 22 Open LDAP server settings

Setting description	Setting
Server name	lab.emc.com
LDAP administrator	cn=admin dc=lab,dc=emc dc=com
Groups	Administrators: cn=administrators,ou=groups,dc=lab,dc=emc,dc=com
	Users: cn=users,ou=groups,dc=lab,dc=emc,dc=com
	Support: cn=support,ou=groups,dc=lab,dc=emc,dc=com
Users	Paul Abbey: uid=PAbbey,ou=people,dc=lab,dc=emc,dc=com (Users member)
	John Smith: uid=JSmith,ou=people,dc=lab,dc=emc,dc=com (Support member)
	Tom Baley: uid=TBaley,ou=people,dc=lab,dc=emc,dc=com (Marketing member)

Scenario: Configuring LDAP integration with Simple Bind

Procedure

1. Go to **Admin > Users & Security > Manage External Authentication**.
2. Verify or type the following values in the User fields:
 - **Use LDAP Authentication:** selected
 - **Server:** lab.emc.com
 - **Use SSL:** selected (optional)
 - **Port:** 686
 - **LDAP Version:** 3

- **Base Name:** dc=lab,dc=emc,dc=com
 - **Identification Attribute:** uid (sAMAccountName for Active Directory integration)
 - **Anonymous Bind:** unselected
 - **Username:** cn=admin,dc=lab,dc=emc,dc=com
 - **Password:** <admin_password>
3. Click **Validate** to verify the LDAP binding.
If the validation fails, check the LDAP connectivity from the DPA Application server and the LDAP server parameters.
 4. Click **Test user** to verify the LDAP binding.
Use the following username and password:
Username: PAbbey
Password: <PAbbey_password>
 5. Click **OK** to verify the LDAP user authentication.
If the authentication fails, check if the username and password are correct in the LDAP server.
 6. Click **OK** in the **Manage External Authentication** to confirm settings and close.
 7. Go to **Admin > Users & Security > Manage Users** and click **Create User**.
 8. Type the following values in the **User Properties** tab:
 - **Name:** Paul Abbey
 - **Logon:** Pabbey
 - **External Name:** PAbbey
 - **Role:** User
 - **Authentication Type:** LDAP
 9. Click **OK** and verify that the account is in the user account list.
 10. Click **Close**.

Scenario: Configuring automated user provisioning with group mapping

Procedure

1. Go to **Admin > Users & Security > Manage External Authentication**.
2. Verify or type the following values in the User fields:
 - **Use LDAP Authentication:** selected
 - **Server:** lab.emc.com
 - **Use SSL:** selected (optional)
 - **Port:** 686
 - **LDAP Version:** 3
 - **Base Name:** dc=lab,dc=emc,dc=com
 - **Identification Attribute:** uid (sAMAccountName for Active Directory integration)
 - **Anonymous Bind:** unselected

- **Username:** `cn=admin,dc=lab,dc=emc,dc=com`
 - **Password:** `<admin_password>`
3. Click **Validate** to verify the LDAP binding.
If the validation fails, check the LDAP connectivity from the DPA Application server and the LDAP server parameters.
 4. Click **Test user** to verify the LDAP binding.
Use the following username and password:
Username: `PAbbey`
Password: `<PAbbey_password>`
 5. Click **Edit**.
 6. Check **Enable Auto Login**, and ensure that the Default User Role selected is **User**.
 7. Check **Enable Group Mapping** and verify or type the following values:
 - **Group Base:** `ou=groups,dc=lab,dc=emc,dc=com`
 - **Group Attribute:** `cn`
 - **Group Member Attribute:** `memberUid` (member for Active Directory integration)
 8. Click **Add**:
LDAP Group Name: `Support`
Role: `Engineer`
 9. Click **Close**.
 10. Log in as John Smith.
A new user account JSmith should be created with the Engineer role.
 11. Log out.
 12. Login as Tom Baley.
A new user account TBaley should be created with the User role.

System settings

You can modify the default system settings for DPA agents, the server, and the datastore.

Configuring backup and restore resolution fields

DPA allows you to create up to five custom backup and restore resolution fields that allow you to add a resolution to a failed job, and then at a later date view the resolution to see what caused the failure.

About this task

For example, you can create a field as a reference to an external ticketing system that includes further resolution information for failed backups. Administrators can control the format of a custom field and make the field mandatory or optional.

Procedure

1. Select **Admin > System > Manage Custom Resolutions**.

The Manage Custom Resolutions dialog box appears.

2. Select an available row from the list and click **Edit**.
The **Resolution Custom Field** dialog box appears.
3. Select **Active** to enable the custom field.
4. Type a Label for the field.
The field label will be used in the **Backup Resolution** and **Add Resolution** dialogs boxes.
5. Select the type of data that the custom field will hold from the **Input Cast** field.
Data types include:
 - Flag (True or False)
 - Integer value
 - Decimal value
 - Text
6. (Optional) Select **Mandatory** to force administrators to complete a field of type Text when creating or adding resolutions. For other field types, the default value is used in the resolution if the user does not specify a value.
7. Click **OK**.

After you finish

If desired, implement backup and restore resolutions in drilldown reports:

Add/View Backup Resolution actions can be used in all system reports that use the "Job Details Popup" drilldown menu.

1. Go to **ReportsReport Templates Custom Report Templates**, select the report you wish to add the backup resolution to and click **Edit**.
2. Select the **Preview** tab.
3. Click **Drilldowns** to display the drilldown reports menu and select **Same drilldown menu for all columns**.
4. Edit or create the pop-up menu with the resolution options:
 - a. Select **Action** and choose one of the backup and restore resolution options:
 - Add Backup Resolution
 - Add Restore Resolution
 - View Backup Resolution
 - View Restore Resolution

Other options include Show selected alerts, exclude edit, gap details, show related alerts, and request history.
 - b. Select **Automatic**.
 - c. Click **OK**.

Viewing and editing settings

To view or edit system settings, select **Admin > System > Configure System Settings**.

About this task

System Settings

The DPA system has settings for Data Collection Agents, Server, SharePoint, Replication Analysis, and Agentless Discovery. The following table describes each agent setting.

Table 23 Data Collection Agent settings

Setting	Description
Data Collection Agent Status	Enables collection of log files. Enabled by default.
Data Collection Agent Version	The version of the DPA data collection agent that is currently installed on the host.
Data Collection Agent Port	Port on which the data collection agent listens for requests.
Concurrency	Maximum number of threads the data collection agent uses to gather data. The default is five.
Log Level	Verbosity level when the data collection agent writes to the log file. For example, selecting Fatal writes only critical errors to the log file.
Log File	The location of the log file on the host.
Max Log File Size (MB)	Maximum size to which a log file can grow before the creation of a new log file (in MB). To set no limit for the size of the log file, set this value to 0.
Max Number of Log Files	Maximum number of log files maintained on the system. If a new file is created because the maximum file size of the current log file is exceeded, the oldest log file is removed.
Max Forward Queue Length	Maximum number of requests stored by the agent locally if the Server is offline.
Max Forward Queue Size (MB)	Maximum total size of all requests stored by the DPA data collection agent locally if the Server is offline (in MB). You can specify unlimited or specify a selected size.
Reload Data Collection Agent	Allows you to manually reload the data collection agent. This is done automatically when configuration changes are made in the DPA web console that affect a data collection Agent.
Remove Data Collection Agent	Removes the selected data collection agent.
Make Agent Default	Makes the selected data collection agent the default host.

Table 24 Server settings

Setting	Description
Global Data Collection Agent Settings	Binary Multiplier Switching this global setting on, defaults all Agents to use the binary multiplier. Binary multiplier converts all incoming data as 1024 KB= 1MB. Applies to NetWorker agents only where the incoming data from Backup server is converted as 1000 KB = 1MB. Binary

Table 24 Server settings (continued)

Setting		Description
		Multiplier is ignored when monitoring other applications.
	Timeout(s)	Time out setting that the server uses when talking to the agent. The default is 120 seconds.
Global Email Settings	Mail Server Hostname	Mail server to which email messages are forwarded when sent from DPA.
	Mail From Address	E-mail address assigned to email messages sent from DPA.
	Mail Server Port	Mail server port number.
Global Logging Settings	Global Logging Settings	Global logging settings for the Analysis Engine, Configuration, Listener, Publisher, Recoverability Analysis, Reporter, and REST API. Settings can be INFO, DEBUG, DEBUG LOW, WARN, ERROR, and FATAL.
Data Deletion	Data Deletion	Schedule to delete data gathered from your environment. The default is 9 a.m. to 5 p.m. every day.
Root Cause Analysis	Root Cause Analysis Settings	Option to enable Root Cause Analysis Summary.
		Option to enable Root Cause Analysis Deletion. The default deletion setting deletes data that is older than 200 days. The period is not user-configurable.
Generate Support Bundle	Generate Support Bundle	Option to generate support zip file.
	Include all logs	Option to include all logs. If unselected, DPA collects only the latest log files. If selected, DPA collects all historical log files. Unselected by default.
DB Export	Database Export Age Notification	Option to set a time period that the DPA Database export is considered up to date. The default value is one week. The minimum is one day. When the time period expires and there is no fresh DPA Datastore export during this period, an alert is issued.

Table 25 SharePoint settings

Setting		Description
Name	Name	The user-defined name of the SharePoint site created in DPA SharePoint Server Settings.
Site	Site URL	The SharePoint destination URL for publications. HTTP protocol defaults to port 80, HTTPS defaults to 443. You can also specify the port explicitly. For example, to set to http port 24438 site URL, type: <code>http://sharepoint-2013:24438/sites/demo2/</code> .
User	Username	The username associated with the SharePoint account

Table 26 Replication Analysis settings

Setting		Description
Replication Analysis	Client-Server Time Difference	The default is 10 minutes.
	Symmetrix and CLARiiON Log Level	Logging settings for Symmetrix and CLARiiON. Settings can be INFO and DEBUG.
	Support Symmetrix Masking Reports	Enables support for Symmetrix Masking Reports. Enabled by default.
	Support Application Discovery Impersonation	Enables support for Application Discovery Impersonation. Enabled by default.
Display Setting	Display dirty Recovery Points	Enables displaying dirty Recovery Points. Enabled by default.
	Aggregate Recovery Points	Enables aggregating Recovery Points. Enabled by Default.
	Minimum number of recovery points to aggregate	The default is 3. The minimum is 1. There is no maximum.

Agentless Discovery

The Agentless Discovery settings are described in the following table.

Table 27 Agentless Discovery settings

Setting	Description
Sudo Program Path	The sudo program path for Agentless discovery settings. The default path is <code>/usr/local/bin/sudo</code> . The sudo command can also be located in either <code>/sbin</code> or <code>/usr/sbin</code> .
Agent Response Timeout	The time that DPA waits for response from the agent before timeout.
Telnet/SSH Login Prompt Timeout	The time that DPA waits for Telnet/SSH session to be created before timeout.
Telnet/SSH Handshake Timeout	The time that DPA waits for Telnet/SSH handshake before timeout.
Delete files created on the client during agentless discovery	Defines if temporary files will be deleted from the analyzed object at the end of the discovery. The default is that the files will be deleted.

Server data deletion

DPA implements a default data deletion schedule for collected data and system-generated data. Collected data is the data gathered by the configured requests within Manage Data Collection Defaults. System-generated data is the data generated by the system processes, such as log messages, histories of reports, and alerts.

When data exceeds the retention period then the data is eligible for deletion. This data is then purged based on the data deletion schedule. Any unprocessed items remain in the queue until the next scheduled start time, at which point deletion of data continues.

You cannot delete a schedule that is currently used for scheduling a collected data deletion job. An error message is displayed if you attempt to do so.

Collected and system-generated data that is deleted is tracked in the `server.log`. For example:

```
Deleted 10 rows from table host_config
Deleted 10 rows from Request History
Deleted 10 rows from reportlogentry
Deleted 10 rows from dpa_request_statistics
Deleted 10 rows from reporterjob
```

The default data deletion schedule is from 9:00 a.m. to 5:00 p.m. daily.

Configuring Data Deletion Schedule

You can configure and specify a new schedule for use in Schedule Properties.

About this task

To configure data deletion, select **Admin > System > Configure System Settings > Server > Data Deletion**. The DPA Online Help provides more information.

Default retention periods

The following table provides information on default collected data retention periods.

Table 28 Default collected data retention periods

System information	Default retention period
Configuration data	365 days
Status data	90 days
Performance data	30 days
Job data	forever
Occupancy data	365 days

Default collected data retention periods are user-configurable within **Admin > System > Manage Data Collection Defaults**.

The following table provides information on default system-generated data retention periods. Default system-generated data retention periods are not user-configurable.

Table 29 Default system-generated data retention periods

Policy	Default retention period
alerts (analysisalert table)	365 days
report history (reporterjob table)	365 days
agent error log entries (reportlogentry table)	14 days
request statistics (dpa_request_statistics table)	28 days

Root Cause Analysis Settings

You can set the Root Cause Analysis Summary to calculate potential root causes on a regular schedule from within the Systems Settings. You can also schedule the system to delete Root Cause Analysis results data. The Root Cause Analysis Deletion setting deletes data that is older than 200 days. The period is not user-configurable. Root Cause Analysis Summary and Deletion are enabled by default.

Disabling Root Cause Analysis Summary

Select **Admin > System > Configure System Settings > Server > Root Cause Analysis Settings > Disable Root Cause Analysis**, and click **OK**.

About this task

Disabling Root Cause Analysis Deletion

Select **Admin > System > Configure System Settings > Server > Root Cause Analysis Settings > Disable Root Cause Analysis Deletion**, and click **OK**.

About this task

Gathering historical backup data using DPA web console

You can gather historical backup data on Avamar, BackupExec, DB2, HP DataProtector, NetWorker, NetBackup, Oracle RMAN, SAP HANA, and TSM.

About this task

Consider the following when you gather historical backup data using DPA web console:

- You cannot gather historical backup data at the host level. You must go one level down in the configuration tree, to the application object. For example, to collect historical data from NetWorker, choose the Networker application object below the host level object.
- You can only gather historical backup from the JobMonitor requests.

Procedure

1. In the web console, select **Inventory > Group Management**.
2. In the configuration tree, select the application object for which you'd like to gather historical backup data.

The application object **Details** window opens.

3. In the host details window, select the **Data Collection** tab.
4. In **Data Collection**, select the JobMonitor request.
5. Right-click **Run** and select **Gather historical data**.
6. In the **Gather historical data** window, click **OK**.

The same credentials and data options are available as for the request itself.

7. Click **Close** to the a dialog box that appears confirming that DPA is gathering the historical backup data.
8. Click **History** to view collected tests. The rows highlighted in orange indicate results from a historical backup gather.

Generate Support Bundle

The Generate Support Bundle option is a support tool. The Generate Support Bundle generates and saves a zip archive with provided resources in the file system directly from the DPA web console.

An EMC Technical Support Engineer might ask you to generate the Support Bundle and send it for analysis. The zip file is saved as the following local agent logs in the `support.zip` folder:

- `dpaagent.log`
- `dpaagent.log.0`
- `dpaagent.log.1`

The default location is user-configurable.

Generating the Support Bundle

About this task

Procedure

1. Select **Admin > System > Configure System Settings > Server > Generate Support Bundle** and click **OK**.
2. When prompted, enter your DPA Administrator credentials.

Digital certificate

DPA uses a self-signed digital certificate for identification and encryption. [Encryption of the DPA Application server](#) on page 63 provides information.

Time periods

When you run a report or create a scheduled report, you must decide the period of time over which the report is run, for example right now or last week. Several predefined time periods are provided by default and you can create custom time periods.

Creating custom time period for reports

To create a custom time period, select **Admin > System > Manage Time Periods**.

About this task

Time zones in DPA

DPA gathers data from the environment and stores it in the DPA database in UTC format.

If the timestamp that the DPA database receives from a backup server, application, host, or switch is in a local time zone, such as EST, the DPA Agent converts it to UTC before sending it into the DPA Server. For reporting on that data, there are several settings that can be set. [Time zone settings for reporting](#) on page 92 provides more information.

Time zone settings for reporting

You can set the following settings for time zones to ensure that the desired time zone displays in your DPA reports.

Discovered object details

After you discover an object using the Discovery Wizard, you can select the properties and choose to specify the time zone of where that object is located. In the discovered object **Details** window, select **Time zone** from the drop-down list.

User preferences

You can choose the time zone you wish data to be displayed in **User Preferences > View User Properties > Preferences**. Select the time zone from the Global Settings section, **Time zone** drop-down list.

Window properties

You can create a time period that is time-zone aware. In the **Window Properties** window, create a new time period and ensure that you select the **Adjust for time zone** option. If you select **Adjust for time zone** and the object is a backup client, DPA checks the parent backup server if the backup client doesn't have a time zone set explicitly on itself already, and creates a report that is time-zone aware.

Report Table Format

You can choose to configure a table style report to show the Time Zone object it was run on with the timestamp, by specifying which field to look at to find the name of the backup server. In the Report Editor, go to **Report Format > Table Format > Table Styles**. Under the Date Fields section ensure that the **Time Zone from Report Field** option is selected.

Example: Setting time zones for All Jobs report

This example shows how to set time zones on an All Jobs report for a NetWorker server that is located in the America/New York time zone for a database administrator that is located in the Europe/London time zone.

About this task

Before you change any settings, all report and display output is in UTC.

Procedure

1. Go to **User Preferences > View User Properties > Preferences** and select **Europe/London** from the Global Settings section, **Time zone** drop-down list. Then click **OK**.

When you run the All Jobs report on the NetWorker server, the DPA returns the report in UTC.

2. Update the time zone of the NetWorker server, which is located in New York:
 - a. Go to **Inventory > Object Library** and navigate to the NetWorker host.
 - b. Select the desired NetWorker host and in the **Details** window, select **America/New York** from the **Time zone** drop-down list.
 - c. Click **OK**.

The output remains in UTC because the user time zone setting or report setting are not yet changed. It does not change to America/New York time zone.

3. Create a custom time period that is time-zone aware. Go to **Window Properties** and create a custom time zone. Ensure that you select the **Adjust for time zone** option.

This makes the report query for the time relative to the time zone of the object. So because the NetWorker server is in New York, DPA runs the query for the custom time period in the New York time zone.

4. Edit the report Table Format so that the date fields display in the time zone from your Server field and a desired date format for the time zone:
 - a. In the Report Editor, go to **Report Format > Table Format > Table Styles**.
 - b. Under the Date Fields section, select the desired date format from the **Date Format** drop down, and ensure that the **Time Zone from Report Field** option is selected.
 - c. Click **OK**.

DPA refreshes the report with the time stamps of the time zone of the NetWorker server, in this case, America/New York.

Automatic report prioritization

The default number of reports to run concurrently per DPA Application server is 10. You can configure the default settings. The maximum number of reports to run concurrently per DPA Application server is 50; the minimum number is 2.

DPA automatically queues reports that are scheduled to run concurrently or that are running concurrently, and automatically retries reports when the previously scheduled reports have been

run. Additionally, any reports that you initiate from the web console take precedence over automated scheduled reports running from the server, including testing a scheduled alert.

In addition to giving priority to reports run from the web console, there is also a 30% minimum fixed concurrent space reserved for these reports on the server. For example, if the concurrency set is 10, three concurrent execution spaces on the server are reserved for web console reports. Hence, there can be three or more out of a maximum of 10 web console reports running at a particular instant. There can be only seven scheduled reports which can run concurrently.

Configuring concurrent report settings

Procedure

1. To configure concurrent report settings, select **Admin > System > Configure Report Settings > Concurrency**.

After you finish

After you change the concurrency setting in the DPA web console, ensure that you restart the DPA Application service. If it is clustered environment, restart all the DPA Application servers. This is so that the report-engine service picks up the new concurrency value.

Schedules

Schedules are used to define when to run a scheduled report or generate a dashboard view block, or to define the backup window specified in the Protection Policy. Several predefined schedules are provided by default and you can also create custom schedules.

A schedule is made up of components that define when each schedule produces certain results or runs certain reports. The Schedule Editor provides two ways to create schedules:

- Basic editor - allows you to create schedules on a weekly basis only and edit the day and time of the schedule.
- Advanced editor - allows you to create more complex schedules by manually editing the schedule parameters.

Schedules created in the basic editor can be edited using the advanced editor. However, schedules created and saved in the advanced editor cannot be edited in the basic editor.

Creating schedules

To create a schedule, select **Admin > System > Manage Schedules**.

About this task

Manage Data Collection Defaults

A DPA request contains data on how and when to gather data from an object. Data collection defaults are the template used by the Discovery Wizard to assign requests to objects. You can set the global default settings in **Admin > System > Manage Data Collection Defaults**.

All requests have a default data gathering frequency and a set of options associated with them. You can edit global data collection default values to be picked up by the Discovery Wizard for certain objects. The DPA online help provides information on editing requests.

You can gather certain types of data with DPA without deploying an agent on the monitored device. To do this, an agent on another computer (such as the DPA Server) gathers the data remotely. When gathering data remotely, the agent's host is referred to as a proxy server. The agent uses a protocol to gather data from the remote computer and forwards it back to the DPA server. The protocol used depends on the type of data being collected.

For certain device types, such as IP switches and Fibre Channel switches, data must always be gathered remotely as it is impossible to install an agent directly on a switch.

To configure remote data collection within DPA, configure the details when assigning requests. If the Discovery Wizard created the objects, this configuration is already created. However, if proxy or credential details have changed, modify the details as required. Retention Periods on Requests are set on individual request using the Edit Request dialog box. Table 15 provides information on default retention periods for Data Collection policies.

Data collection request options by module

Data collection request options by module are described in the following table.

Table 30 Data collection request options by module

Module	Option name	Value	Description
ARCserve	dateformat	%d/%m/%Y %T which is day,month, year and time.	The date format to be used. The <code>dateformat</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Job Monitor • Volume Status <p>Note 1. provides additional information on time formats.</p>
Avamar	capacityfactor	1.075	Avamar decimal capacity factor. The <code>capacityfactor</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	dbname	mcdb	Database name. The <code>dbname</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status requests
	dbport	5555	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status requests
	amount of seconds in the job collection	86400	Changes the maximum amount of time the request gathers job data in one run of Jobmonitor. Default is 86400, which is one day in seconds. The value is configurable.
Backup Exec	dbserver	No default value	Database server\instance. The <code>dbserver</code> option is present in the options for the following requests:

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Configuration • Job Monitor • Status • Volume Status
CLARiiON VNX	Connector	No default value	Indicates connector for the import clariion information request
	EventLog History Polling	21	The age of the data after which it is no longer included in polling in days for import clariion information request
Celerra	port	No default value	<p>HTTPS/HTTP port number in integers. The <code>port</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Performance • Status
	secure	True	<p>Indicates to send requests using HTTPS instead of HTTP. The <code>secure</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Status
	timeout	1800	<p>HTTP request timeout, in seconds. The <code>timeout</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Performance • Status
CommVault Simpana	appversion	0	<p>The version of CommVault Simpana to use. The <code>appversion</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Client Occupancy • Job Monitor • Status • Volume Status

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
	dbserver	No default value	DB server name. The <code>dbserver</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Client Occupancy • Job Monitor • Status • Volume Status
	setBackupJobsWithErrsToSuccess	False	If set to <i>True</i> , DPA reports commvault backup jobs' that were completed with <i>Completed w/one or more errors</i> status as successful jobs. The <code>setBackupJobsWithErrsToSuccess</code> option is present in the Job Monitor request.
Data Domain	timeout	10	SSH Timeout value in seconds. The <code>timeout</code> option for SSH is present in the options for the following requests: <ul style="list-style-type: none"> • Analysis • Configuration SSH • Performance SSH • Status SSH • SSH PCR
	timeout	10	SNMP Timeout value in seconds. The <code>timeout</code> option for SNMP is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
Data Protector	timeout	900	The timeout value in seconds for running commands for the Configuration request
	timeout	300	The timeout value in seconds for running commands. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Internal Database • Job Monitor

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Service Status • Status • Volume Status
	ignorefailedclones	False	Indicates not to collect information about source objects for failed clone jobs for the Job Monitor request
	nojobmedia	False	Indicates not to collect media information associated with each job for the Job Monitor request
	occupancy	False	Indicates to enable gathering of occupancy statistics for the Job Monitor request
	timeformat	No default value	omnidb time format for the Job Monitor request. Note 2. provides additional information on time formats.
DB2	amount of seconds in the job collection	86400	Changes the maximum amount of time the request gathers job data in one run of Jobmonitor. Default is 86400, which is one day in seconds. The value is configurable.
	database port	50000	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Job Monitor
	grace period for job	0	The grace period for the agent to look back and gather data for each request Job Monitor request, in seconds. The value is configurable. The minimum value is 0; the maximum value is 1000000.
EDL	timeout	10	SNMP timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
Fibre Channel Switch	timeout	10	SNMP timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests:

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Configuration • Performance • Status
Host System Monitoring	disk	True	Indicates to include host disk information for the Configuration and Replication request
	ESXRequestParameters.ESX_CREDENTIALS	No default value	ESX server credentials for the Configuration and Replication request
	ESXRequestParameters.ESX_SERVER	No default value	Name of the ESXServer server to be used the Configuration and Replication request
	fchba	True	Include host FC HBA information. The <code>fchba</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	fs	True	Include host filesystem information. The <code>fs</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	host	True	Include basic host information. The <code>host</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration and Replication • Status
	logical	False	Include logical network interfaces. The <code>logical</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	memory	True	Include host memory information. The <code>memory</code> option is present in the options for the following requests:

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	netint	True	<p>Include host network interface information. The <code>netint</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	remote	False	<p>Include remotely mounted filesystems. The <code>remote</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration and Replication • Performance • Status
	REPLICATION_MONITORING_OPTION	False	Enable Replication Monitoring for the Configuration and Replication request
	srm	True	Utilize srm libraries for disk/fs information for the Configuration and Replication request
	Time Offset(seconds)	0	Time Offset in seconds for the Configuration and Replication request
	disk	True	<p>Include host disk information. The <code>disk</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Performance • Status
	fullpath	False	Include the full path of the process name for the Status request
	process	True	Include host running processes information for the Status request
	specific	No default value	Monitor the named process only for the Status request; Windows only.
Illuminator clarapi Engine Discovery	TIME_OFFSET_OPTION	0	Time offset in seconds for the illuminator clarapi engine discovery request

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
HP Disk Array	port	5989	CIM provider port for HP EVA disk arrays. The <code>port</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
HP Virtual Library System	port	5989	Port to the HP VLS disk arrays. The <code>port</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	SSLflag	True	SSL flag is enabled for HP VLS disk arrays. The <code>SSLflag</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	timeout	600	Timeout in seconds for HP VLS disk arrays. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
Illuminator symapi Engine Discovery	Symapi Version	No default value	Indicates the SYMAPI version for the illuminator symapi engine discovery request
	TIME_OFFSET_OPTION	0	Time offset in seconds. The <code>TIME_OFFSET_OPTION</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • illuminator symapi engine discovery • import symmetrix information
	Allow Management over SRDF	False	Allows management over SRDF for the illuminator symapi engine discovery request.
	SYMAPI DB Path	No default value	Indicates the SYMAPI database path for the illuminator symapi engine discovery request

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
IP Switch	timeout	10	Timeout value in seconds for the Status, Performance, and Configuration requests
SQL Server Database	dbparams	No default value	XML specifying per database parameters/credentials. The <code>dbparams</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status
	dbport	1433	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status
	HomeDir	No default value	Application home directory information for the mssql application discovery request
	Tools Director	No default value	Tools directory property information for the mssql application discovery request
	Virtual Computer Name	No default value	Virtual computer name property information for the mssql application discovery request
	Number of (rotated) error log files to check for failed backups	Default value is 6	Use this option to determine the number of log files to be read. A larger number allows you to read older logs. The value must be greater than 0.
NearStore	timeout	10	SNMP timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
NetBackup	timeout	300	Command timeout in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Client Occupancy

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Configuration • Job Monitor • Media Server Status • Status • Volume Status • SLP Job Status
	EMMserver	No default value	<p>Hostname of Enterprise Media Manager (EMM) server; required only if not the Master Server host. The <code>EMMserver</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Status
	timeformat	No default value	License expiration date time format for the Configuration request
	timeformat	No default value	<p>bpdjobs time format for the Job Monitor request</p> <p>Notes 1. and 2. provide additional information on time formats.</p>
	partialasfailed	False	Mark partially successful jobs as failed for the Job Monitor request
	Whether to include container jobs	False	If set to <code>True</code> DPA also gathers the row for the parent/container job, in addition to the child jobs. Can be set for the Default request or on individual objects.
	Max data time range each request will gather	86400	Changes the maximum amount of time the request gathers job data in one run of SLP Job Status. Default is 86400, which is one day in seconds. The value is configurable.
NetWorker	command timeout	3600	<p>The timeout in seconds, used for running external commands to gather data.</p> <p>The <code>command timeout</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Status • ClientStatus

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> • JobMonitor • Occupancy • Volume Status
	mminfo timeformat	No default value	<p>The format that timestamps in the media database are returned in bpdjobs time format. This is used to decode the start time/end time for a Job. By default, this option is disabled and the module attempts to automatically calculate this value.</p> <p>The <code>mminfo timeformat</code> option is present in the Job Monitor request.</p>
	include jobs from media DB	True	<p>Allows you to switch off the search for successful jobs from the NetWorker Media database. DPA searches the Media database in addition to the NetWorker Jobs database for completed jobs. If you are not using an external scheduler to initiate backups, then set this to <code>False</code> to speed up running the Job Monitor request</p> <p>The <code>include jobs from media DB</code> option is present in the Job Monitor request.</p>
	max batch period of each request data poll	86400	<p>Changes the maximum amount of time the request gathers job data in one run of Jobmonitor. Default is 86400, which is one day in seconds. The value is configurable.</p>
NetWorker	individual ping timeout	10	<p>The timeout in seconds, used for timing out ping responses from backup clients.</p> <p>The <code>individual ping timeout</code> option is present in the Client Status request.</p>
	nsrexecd port	7937	<p>The NetWorker client process listen port.</p> <p>The <code>nsrexecd port</code> option is present in the Client Status request.</p>
	Number of concurrent pings	20	<p>The number of clients to ping at any one time.</p>

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			The <code>Number of concurrent pings</code> option is present in the Client Status request.
	List of critical clients to ping	No default value	The name of the file that holds a comma separated list of critical clients to ping instead of all clients . The <code>List of critical clients to ping</code> option is present in the Client Status request.
	Path and name of file used to store temporary occupancy data before processing	No default value	Path and name of file used to store temporary occupancy data before processing. The value should be a valid path on the Agent host. For example, on Windows use <code>C:\temp</code> and on UNIX/Linux use <code>/tmp</code> . You may need to restart the Agent after enabling for the option to take effect. The <code>Path and name of file used to store temporary occupancy data before processing</code> option is present in the Occupancy request.
	Forces short client names	true,false	Whether to return the short version of the client name or not. The <code>Forces short client names</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status • ClientStatus • JobMonitor • ClientOccupancy
	time format used to determine bootstrap time	False	Specifies the time format used to decode timestamps returned in the results from NetWorker. By default, this option is not enabled and the module use a best-guess method to decode the time format. The <code>time format used to determine bootstrap time</code> option is present in the Status request.

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
	time format used to determine volume access time	False	<p>Time format to use when decoding timestamps regarding the last time a volume was accessed. By default, this option is not set and the module attempts to calculate a time format automatically.</p> <p>The time format used to determine volume access time option is present in the Volume request.</p>
	time format used to determine volume retention period	False	<p>Time format to use when decoding timestamps regarding the retention time for a volume. By default, this option is not set and the module attempts to calculate a time format automatically.</p> <p>The time format used to determine volume retention period option is present in the Volume request.</p>
	Whether to include failed jobs which are retried	False	<p>The DPA Agent gathers only the final status of a backup job.</p> <p>If the job option is set to <code>False</code> and there is a long delay before the job runs, then the retry fails then this delays the reporting of the job failure. If the retry succeeds, the DPA Database has one entry for the job and that entry shows the job as a success.</p> <p>If the job option is set to <code>True</code>, the DPA Agent gathers all failed tries and the final status of the job and sends them to the DPA database. For example, if the job is retried once and succeeds, the DPA database records 2 entries for the job, 1 with a failure and 1 with a success. If both tries fail, then the DPA database records 2 job entries in the DPA database, both as failures.</p> <p>You can use the All Jobs - No Restarts report to filter out the failed attempts and show only the final status of the job.</p>

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
Oracle	dbparams	No default value	XML specifying per schema parameters/credentials. The <code>dbparams</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	dbport	1521	Database port integer. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	HomeDir	No default value	Application home directory information for the oracle application discovery request
	ArchivesPattern	No default value	Application archive pattern information for the oracle application discovery request
	LogPattern	No default value	Application log pattern information for the oracle application discovery request
	LogsDir	No default value	Application log directory information for the oracle application discovery request
PostgreSQL Database	dbparams	No default value	XML specifying per schema parameters/credentials. The <code>dbparams</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	dbport	5432	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status
	initialdb	postgres	Initial database to connect to this port.. The <code>initialdb</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Status

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
PureDisk	dbport	10085	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Client Occupancy • Configuration • Job Monitor
	dbserver	No default value	Database server host. The <code>dbserver</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Client Occupancy • Configuration • Job Monitor
RecoverPoint	scanforrecover	False	Scan for Recoverability for the configuration request
	Time Offset (in seconds)	0	Time offset in seconds for the configuration request
	timeout	300	SSH timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance cs • Performance
	filename	long_term_stats.tar.gz	Statistics filename for the performance cs request
	workdir	../tmp	Working directory for the performance cs request
RecoverPoint for VMs	Time Offset (in seconds)	0	Time offset in seconds for the configuration request
	timeout	300	REST API timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance CS • Performance • Status
RMAN	dbport	1521	Oracle TNS listener port. The <code>dbport</code> option is present in the options for the following requests:

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<ul style="list-style-type: none"> Job Monitor Control File Job Monitor Recovery Catalog
	amount of seconds in the job collection	86400	Changes the maximum amount of time the request gathers job data in one run of Jobmonitor. Default is 86400, which is one day in seconds. The value is configurable.
	RMAN Schema	No default value	
SAP HANA	database port	30115	Database port for the Job Monitor request
	amount of seconds in the job collection	86400	Changes the maximum amount of time the request gathers job data in one run of Jobmonitor. Default is 86400, which is one day in seconds. The value is configurable.
Symmetrix	Connector	No default value	Indicates connector for the import symmetrix information request
	Gather HBA Information	True	Gather HBA information for the import symmetrix information request
	Time Offset (in seconds)	0	Time offset in seconds for the Configuration request
	Symaudit History Polling	21	The age of the data after which it is no longer included in polling in days for Symaudit request
Tape Library	timeout	10	SNMP Timeout in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> Configuration Status
TSM	timeout	No default value	Internal timeout for commands sent to TSM server in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> Client Occupancy Job Monitor Process Monitor Volume Status
	timeout	3600	Internal timeout for commands sent to TSM server in seconds for the Configuration request

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
	timeout	900	Internal timeout for commands sent to TSM server in seconds for the Status request
	tsmhost	No default value	<p>Hostname of TSM server. The <code>tsmhost</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Client Occupancy • Configuration • Job Monitor • Process Monitor • Status • Volume Status
	tsmport	1500	<p>Port of TSM server. The <code>tsmhost</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Client Occupancy • configuration • Job Monitor • Process Monitor • status • Volume Status
	disableprivatevolumes	False	<p>Disable reporting of private volumes. The <code>disableprivatevolumes</code> option is present in the options for the following requests:</p> <ul style="list-style-type: none"> • Configuration • Volume Status
	backupsets	True	Whether to gather backup sets for the Job Monitor request.
	filterbynoderegtime	True	Filter Missed Jobs before node registration for the Job Monitor request
	Whether to gather failed jobs from the activity log	False	<p>If this is enabled and set to <code>True</code> any messages in the TSM activity log that indicates a failed backup has taken place are also reported as a failed job in DPA.</p> <p>The <code>Whether to gather failed jobs from the</code></p>

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
			<code>activity_log</code> option is present in the Job Monitor request.
	<code>processingtype</code>	No default value	The source of the processing jobs for the Job Monitor request. It can be either <code>SUMMARY</code> or <code>ACTLOG</code> .
	<code>OPTION_LIB_MANAGER_CRED</code>	OptionDefinition.Type.Credential	Library Manager Credentials for the Volume Status request
	<code>ignorewarnings</code>	No default value	Warning codes to treat Job Monitor request as successful. There are no character limitations.
VMware	<code>port</code>	443	Port of VMware server. The <code>port</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
	<code>timeout</code>	3600	Internal timeout for commands sent to VMware host in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • configuration • Performance • Status
	<code>usessl</code>	True	Use SSL over HTTP. The <code>usessl</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
	<code>vmwarehost</code>	No default value	Hostname of VMware server. The <code>vmwarehost</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
VMware vSphere Data Protection (VDP)	<code>capacityfactor</code>	1.075	Decimal capacity factor. The <code>capacityfactor</code> option is present in the options for the following requests:

Table 30 Data collection request options by module (continued)


Module	Option name	Value	Description
			<ul style="list-style-type: none"> • Configuration • Status
	dbname	mcdb	Database name. The <code>dbname</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status requests
	dbport	5555	Database port. The <code>dbport</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Job Monitor • Status requests
VPLEX	port	443	HTTPS/HTTP Port for the Configuration request
Webserver	page	No default value	Web page to get for the Response request
	port	80	Web server port. The <code>port</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Response
Xsigo	timeout	10	SNMP timeout value in seconds. The <code>timeout</code> option is present in the options for the following requests: <ul style="list-style-type: none"> • Configuration • Performance • Status
 Note: The following time formats are supported:	1.	<ul style="list-style-type: none"> • %c - Locale-specific • %x %X - Locale-specific - alternate format • %m/%d/%y %l:%M:%S %p - Hard- 	The meaning of the elements in the time and date formats is: <ul style="list-style-type: none"> • %c - Date and time using the current locale format • %x - Date using the current locale format • %X - Time using the current locale format

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
		coded 12-hour US date format <ul style="list-style-type: none"> • %m/%d/%Y %l:%M:%S %p • %d/%m/%y %l:%M:%S %p - Hard-coded 12-hour European date format • %d/%m/%Y %l:%M:%S %p • %m/%d/%y %r • %m/%d/%Y %r - Locale-specific • %d/%m/%y %r • %d/%m/%Y %r • %d/%m/%y %T • %d/%m/%Y %T • %m/%d/%y %T • %m/%d/%Y %T • %x - Locale-specific • %m/%d/%Y • %m/%d/%y • %d/%m/%y • %d/%m/%Y • %d.%m.%Y %T 	<ul style="list-style-type: none"> • %m - Month as an integer (1 - 12) • %d - Day of the month as an integer (00 - 31) • %y,%Y - Year without the century, as an integer (0 - 99) • %l - Hour in 12-hour format (1 - 12) • %M - Minute as an integer (0 -59) • %S - Seconds as an integer (0 - 59) • %p - Locale's equivalent of AM/PM • %r - Time in 12hr am/pm format • %T - Time - alias for hours:Minutes:Seconds.
	2.	<ul style="list-style-type: none"> • %c 	

Table 30 Data collection request options by module (continued)

Module	Option name	Value	Description
		<ul style="list-style-type: none"> • %x %X • %x, %X 	

Manage Sites

You can set the `Site` attribute in the object property dialog, similar to other attributes, like `Credentials` and `Schedule`. You can assign the `Site` attribute to all top-level and component objects. DPA does not support assigning the `Site` attribute to group objects. Objects are searchable by the `Site` attribute.

Go to **Admin > System > Manage Sites** to add, edit, and delete sites.

Creating, editing, and deleting Sites

Procedure

1. Go to **Admin > System > Manage Sites**

The **Manage Sites** window appears.

2. If you would like to:

- Create a site:
 - a. Click **Create Site**.
The **Create Site** dialog appears.
 - b. in the **Site Name** field, type a name for your site.
 - c. In the **Location** field, type three or more characters for the geographical location that is closest to your site, then select the location that is the best geographical match to your site.
 - d. Click **Select Location** and **OK**.
- Edit a site:
 - a. Select the site that you would like to edit from the list of sites. The **Edit Site** dialog appears.
 - b. Edit the desired field and click **OK**.
- Delete a site:
 - a. Select the site that you would like to delete from the list of sites. The **Delete Site** dialog appears.
 - b. Confirm or cancel deletion of the site, as applicable.

Application service administration

Running Linux DPA Application as non-root user

By default the DPA Application runs under root user on Linux. Carry out this procedure on the DPA Application Server to configure the DPA Application to run as a non-root user.

Procedure

1. Stop the DPA Application service. Type: `dpa app stop`
2. Create an OS user to be used for running DPA Application.
Alternatively, select the OS username `apollosuperuser` to be used for running the DPAApplication from the `dpaservices` group.
The `apollosuperuser` user is created during DPA Installation.
3. Transfer ownership of DPA services installation directory to the OS user, which will run the DPA Application. Type: `chown --dereference -IR <user_to_run_dpa>:<group_of_user> <dpa_install_dir>/services`
4. Modify the `<dpa_install_dir>/services/executive/applnsvc.sh` file. Change the line `RUN_AS_USER=` to `RUN_AS_USER=<user_to_run_dpa>`.
5. If performing an upgrade, repeat steps 3 and 4 above to restore the required FS permissions and owners.
6. Start DPA Application service. Type: `dpa app start`
7. (Optional) If you configured any third-party scripts, for example, pre-processing script for scheduled reports or post-processing script in publish settings or scripts for Analysis Policies, modify the scripts for the OS user to `<user_to_run_dpa>` as indicated in step 4.

The DPA Application may receive a permissions denial to run scripts under a new OS user that previously ran under the root user.

Setting TLS protocol version 1.2 only after installation or upgrade

You can set the TLS protocol version to 1.2 only after DPA installation or upgrade by using the `dpa application tlslevel` command.

Procedure

1. Stop the DPA Application server. Type:
`dpa app stop`
2. Run the `dpa application tlslevel` command to set the TLS protocol version to version 1.2 only. Type: `dpa app tls 1.2`
3. Start the DPA Application server. Type:
`dpa app start`

Customization of service information

This section provides information on the types of DPA service customization which only an administrator can do. You must have physical access to the host on which DPA is running.

The *Data Protection Advisor Product Guide* provides information on customizing viewlets, dashboards, and reports. Users can carry out these customizations.

VTL templates

When the Publisher process creates reports when publishing to HTML, it uses VTL templates located in the `vlttemplates` directory on the DPA Server to determine the report's default layout and style. By default, the DPA Server process uses the following template files: `reportcard.vtl`, `chart.vtl`, and `table.vtl` however, you can use another template file. You can create template files to change the appearance of reports that are published by the DPA Server process.

The template types are:

- Default uses the default VTL for the renderer.
- `pivot` is for generating pivot tables.
- `pivot.css` is for generating pivot tables using CSS.
- `pivot.controlpanel.css` is for generating pivot tables in control panels using CSS.

The following table lists the VTL templates.

Table 31 VTL templates

VTL template	Description	Template type
<code>chart.vtl</code>	Used by chart renderers that produces an image for the HTML output such as Area, Column, Line, Pie, Topology	Default
<code>chart.controlpanel.css.vtl</code>	Same as <code>chart.vtl</code> except this uses CSS.	N/A
<code>chart.css.vtl</code>	Same as <code>chart.vtl</code> except this uses CSS	css
<code>email.attach.vtl</code>	Used when sending the report as an attachment to the email	N/A
<code>email.image.embed.vtl</code>	Used for embedding the report inside of the email	N/A
<code>email.notification.vtl</code>	Used for creating the notification that can be sent out after a report was published	N/A
<code>healthstatus.vtl</code>	Used for Health Status	Default
<code>healthstatus.controlpanel.css.vtl</code>	Same as <code>healthstatus.vtl</code> except this uses CSS. Also does not contain the date and version at the bottom	N/A
<code>healthstatus.css.vtl</code>	Same as <code>healthstatus.vtl</code> except this uses CSS	css
<code>reportcard.vtl</code>	Used for ReportCard	Default
<code>reportcard.controlpanel.css.vtl</code>	Same as <code>reportcard.vtl</code> except this uses CSS. Also does not contain the date and version at the bottom	N/A

Table 31 VTL templates (continued)

VTL template	Description	Template type
reportcard.css.vtl	Same as reportcard.vtl except this uses CSS	css
table.controlpanel.css.vtl	Same as table.vtl except this uses CSS. Also does not contain the date and version at the bottom	N/A
table.vtl	Used for Table	Default
table.css.vtl	Same as table.vtl except this uses CSS	css
table.pivot.controlpanel.css.vtl	Same as table.pivot.vtl except this uses CSS. Also does not contain the date and version at the bottom	pivot.controlpanel.css
table.pivot.css.vtl	Same as table.pivot.vtl except this uses CSS	pivot.css
table.pivot.vtl	Used for Pivot Table	pivot
timeline.vtl	Used for timeline charts. HTML gets embedded in the VTL	Default
timeline.controlpanel.css.vtl	Same as timeline.vtl except this uses CSS. Also does not contain the date and version at the bottom	N/A
timeline.css.vtl	Same as timeline.vtl except this uses CSS	css

Example - Part 1: Adding a message and company details to the table VTL template

If you are required to send daily or weekly reports in HTML format to customers, and you accomplish this with scheduled reports, then you can add custom text (such as a message or company contact information) to the scheduled report by creating a custom VTL template. The custom text displays for all HTML reports using this template.

Procedure

1. In the `styles` or `vtltemplates` directory on the DPA Server, copy the table template, `table.vtl`, and rename it. For example, if you are creating a VTL template for table reports for the company EMC, use the naming standard of `table.<companyName>.vtl` then rename the table template to `table.emc.vtl`
2. Open the VTL in a text editor.
3. Using HTML tags, add text similar to the following within the body.

```
<body bgcolor="{$background}"><font face="Arial, Verdana, Helvetica, Sans-serif" color="{$foreground}">
```

```

<body>
Dear customer,
<p>
Your daily system status report is below.
<p>
Thank you,<br>
EMC Corporation
<p>
US Phone:1-800-555-5555<br>
Email:support@EMC.com<br>
Website: www.EMC.com
<p>
<table>
...
</table>
</body>

```

4. Save the VTL.

Example - Part 2: Using a custom VTL template in a scheduled report

Now that you have a custom VTL template, select this VTL in the Scheduled Report Wizard.

Procedure

1. In the DPA web console, create a new or update an existing scheduled report.
2. In **Publish Settings**, select the Web Page (.html) report format and complete the remaining fields.
3. In **Advanced**, select the EMC template and then click **OK**. The template named Default is the unedited table.vtl.
4. Click the test icon to send the scheduled report to the Publisher. If you publish to file, proceed to the default directory to view the report and then make any necessary updates to the VTL template. The default directory of the report is `<install-dir\services\shared\report-results\scheduled`.
5. If no further updates need to be made to the VTL template, save and close the Scheduled Report Editor.

Custom templates import and export

You can import and export custom report templates and custom dashboards from DPA 5.5.1 and later into DPA from a WDS file through the Custom Templates section. Importing and exporting to XML is not supported. You cannot import or export system templates. The imported reports must be supported on DPA.

You can import and export custom report templates and custom dashboards to fulfill the following needs:

- Import custom reports from DPA 5.x.
- Import custom reports that were created by EMC Professional Services.
- Export custom reports to back them up.
- Export a custom report that is not working to send it to EMC Customer Support for troubleshooting.

The *Data Protection Advisor online help system* provides more information on how to import and export custom report templates.

Clustering administration

Adding an Application server to a cluster after DPA deployment

Use this procedure to modify a DPA Application server that was installed as a standalone server, the installation default state, to be part of a cluster after DPA is deployed and operational using the DPA CLI.

Before you begin

- Stop the DPA agents.
- If you are running UNIX machines, ensure that you are a root user.

About this task

The commands in this procedure are formatted for UNIX.

Procedure

1. If you are not going to configure the node to be a Slave, proceed to step 2. If the standalone Application server is going to be a Slave node within the cluster, empty the message queues:
 - a. Stop the data collection agents.
 - b. Verify that the folder `/opt/emc/dpa/services/standalone/data/messaginglargemessages` has no messages. If it has no messages, proceed to step d.
 - c. If the `/opt/emc/dpa/services/standalone/data/messaginglargemessages` folder is not empty, run the following REST call on both the Application Master and Slave nodes:

HTTPS Operation : GET

REST URL: `https://<hostname>:9002/dpa-api/support/queues?name=DLQ`

The output should include a line such as the following:
`<currentTotalMessageCount>21</currentTotalMessageCount>`

For instance, in the example, `>21<` should match the number of files in the `messaginglargemessages` folder `/opt/emc/dpa/services/standalone/data/messaginglargemessages` folder. If the number of files does not match, wait until the messaging queue becomes empty.
2. Set the database connection pool size in all Datastore nodes. Run:


```
# dpa ds tune --connections xxx <RAM>GB
```

 where `xxx` is approximately 250 per each Application server. For example, 500 for a two-node cluster.

If the cluster is enabled with Datastore Replication, run this command for all Datastore Slaves.
3. If you are not running UNIX, proceed to step 4. If you are running UNIX machines, increase the number of file descriptors in the UNIX Application server:
 - a. Edit the `edit /etc/sysctl.conf` file to add the line `fs.file-max = 512000`
 - b. At the prompt, run `# sysctl -p`.
 - c. Edit the `/etc/security/limits.conf` file to add the line `* - nofile 65535`.
 - d. At the prompt, run `# ulimit -n 65535`.
4. Stop the Application server on the first node. Run:

```
# dpa app stop
```

- Promote the Application server to a Clusterable state. Run:

```
dpa app promote --role MASTER --bind <MASTER_IP> --path <Path to network share>
```

The `dpa app promote` command uses the default multicast port 239.1.2.10. You can specify a different multicast port as an optional parameter to this command. Ensure that all the cluster nodes use the same multicast address.

- Start the Application server. Run:

```
# dpa app start
```

- Verify in the `server.log` that node has started as Master.

A cluster can have only one Master node.

- Install additional Slave nodes.

After you finish

Apply the following configuration after upgrade:

- Report configuration settings
 - Log in to the DPA web console.
 - Go to **Admin > System** and then to **Configure Report Settings > Concurrency** .
 - Set the **Maximum Concurrent Reports per Application server** to **6** for the cluster.

Removing an Application server from a cluster

You can remove an Application server from a cluster using the DPA CLI to convert it back to standalone.

Procedure

- On the Application server, type `dpa application stop` to stop the Application service. The Application service must be stopped before removing from a cluster.
- On the Application server, type `dpa application demote` to demote the Application from a running cluster.
- On the Application server, type `dpa application configure` to verify that the Application is removed from the cluster.

It will show as type `STANDALONE`.

- On the Application server, type `dpa application start` to start the Application service and restore the Application server functionality.

[dpa CLI command](#) on page 127 provides more information on DPA Clustering CLI commands.

Clusters considerations for changing passwords

If the password for the Domain user is changed, you must uninstall and reinstall the DPA Application node.

- Run the following commands:

```
dpa app uninstall
dpa app install --user (DOMAIN\username) --password (password)
```


where:

- (*DOMAIN\username*) is the user account with which to run the Application service. The Log on as a service Windows permissions must also be enabled.
- <password> is the password for the user specified.

Datastore service administration

Note the following limitations for Datastore Replication:

- In busy environments, best practice is to stop the Application servers for a Datastore Replication export so that the export can complete and be imported to the Slave Datastore, and resync with the Master Datastore.
- DPA supports Datastore Replication exports from the Master Datastore only. DPA does not support Datastore Replication exports run from the Slave Datastore.

Backup of the Datastore

It is a best practice to back up the DPA Datastore regularly and particularly prior to making any major change to DPA such as upgrading to a newer version or migrating to new hardware. An export of the Datastore contents is part of an overall backup of the DPA instance.

Exporting and importing a DPA Datastore is supported only on the same version of the DPA Datastore.

Exporting the DPA Datastore

With this export command, a complete and consistent copy of the Datastore is exported to the local file system, in a location that can optionally be specified.

About this task

The default folder/subdirectory of the export is: `datastore-<version> <date and time>`.

For example, `datastore-6_3_0_90597-2017-10-01-1135`.

Type the following command from a command line prompt. `dpa datastore export [options]`

By default, the exported Datastore folder is saved to the same directory where the export command was run.

To save the exported Datastore folder to a specific directory, specify the location at the end of the command line. For example, the following command line exports the folder to `C:\` because that is the location specified: `C:\Program Files\EMC\DPA\services\bin>dpa datastore export C:\`

Exporting the DPA Datastore to Pipe

With this export format, a complete and consistent copy of the Datastore is streamed to a named pipe from a location where Avamar can read the contents.

About this task

Type the following command from a command line prompt. `dpa datastore export --pipeline`

For example, `dpa datastore export --pipeline /mydir/mypipe`

DPA supports backup up to Avamar using the `ds export` command and piping it directly to Avamar. For more information, see the Avamar documentation on how to pipe a backup into Avamar using "named pipes."

After you export Datastore

The `dpa ds export` command produces a folder that contains all the export files for the DPA Datastore. Read about the recommended actions for the folder.

You should back up the folder that contains all the export files for the DPA Datastore with Avamar or NetWorker or any other backup application.

If you are using Avamar, you must restore the content first and then carry out the import on DPA.

If you are using NetWorker, consider placing these Datastore export folders into a separate filesystem and use the NetWorker block-based backup method to backup this folder efficiently.

Importing the DPA Datastore

The `dpa datastore import` command line option is used to import the contents of a Datastore file to the DPA Datastore.

Procedure

1. Stop the DPA Application service.
2. Import the Datastore.
3. Start the DPA Application service.
4. From a command line prompt, type the following:

```
dpa app stop dpa datastore import [options] <filename> dpa app start where
<filename> is the previously exported Datastore file. The import command replaces the
existing datastore contents with the contents contained in the Datastore export file.
```

After you finish

For a complete list of DPA CLI commands, type `dpa --help` from a command line prompt. [DPA command line operations](#) on page 127 provides more information.

Datastore Replication administration

Configuring Datastore Replication after deployment

Use this procedure to configure Datastore replication on a system that is already installed and operational. Note that the CLI commands in this section are formatted for Linux RHEL.

Procedure

1. Confirm that the Datastore server is installed as a Slave. If it is not, configure the Datastore server as a Slave Datastore. Run `dpa.sh ds rep --role SLAVE <IP of master>` to make the Datastore server a Slave.
2. Follow the procedure [Integrating Slave Datastore after it has been offline](#) on page 125.

Configuring cascading Datastore Replication

You can configure cascading Datastore Replication after installation only with the DPA CLI. With cascading Datastore Replication, the Master Datastore replicates to a chain of Slave Datastores, one of which can be remote. Note that the CLI commands in this section are formatted for Linux RHEL.

Before you begin

- Stop all Application Servers. Type: `dpa.sh app stop`
- Stop all Datastore Servers. Type: `dpa.sh ds stop`

- The install directory for the Datastore must be the same on each Datastore machine for the import/export functionality to work.

Procedure

1. On the Master Datastore, run the following commands:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --role master
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --addSlave
<ip_of_replicating_slave> <DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

2. On the replicating Slave Datastore, run the following commands:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --role replicating_slave
<ip_of_master> <DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --addSlave
<ip_of_slave> <DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

3. On the Slave Datastore, run the following commands:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --role slave
<ip_of_replicating_slave> <DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

4. Synchronize the Slave Datastores with the latest Datastore copy from the Master Datastore:

- a. For each Datastore, create an empty directory on the Master Datastore to which to export the Master Datastore file set.

For example, /tmp/export.

- b. On the Master Datastore, run the following command, and ensure that you keep the Master Datastore running when you run the command

```
dpa.sh ds rep --export /tmp/export
```

- c. Use the appropriate platform to command copy the files to the empty directory on the Slave Datastore.

- d. On the replicating Slave Datastore, run the following commands:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --import /tmp/export
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

- e. On the Slave Datastore, run the following commands:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --import /tmp/export
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

5. Verify that replication is working on the Datastores. Run the command:

```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep
```

Output of the replicating Slave Datastore looks similar to the following:

```
<DPA_HOME>/emc/dpa/services/logs # /binary/emc/dpa/services/bin/dpa.sh
ds rep

Data Protection Advisor

[INFO] Replication State : REPLICATING_SLAVE (for 10.11.111.110)
[INFO] Defined Slaves
           : 10.11.111.111/12

[INFO]
           SLAVE          BYTES LAG
STATUS
[INFO]
           10.11.111.111          0
```

```
streaming
[INFO] SLAVE is behind the MASTER by 0 [HH:MM:SS]
Command completed successfully.
```

6. Start the Application Servers. Type: `dpa.sh app start`

After you finish

If the Master Datastore fails, you can make the replicating Slave Datastore or Slave Datastore into a new Master so that DPA can continue functioning. [Carrying out Datastore server failover](#) on page 124 provides more information.

Carrying out Datastore server failover

When the Master Datastore fails, carry out a failover to the Slave Datastore.

Before you begin

Ensure that the Slave Datastore is running.

Procedure

1. On the Slave Datastore, type:


```
dpa.sh ds rep --failover
```
2. Stop the Application server. Type:


```
dpa.sh app stop
```
3. Reconfigure the Application server to point to the new Master Datastore. Type:


```
dpa.sh app con -m <hostname/IP of new MASTER>
```
4. Verify that the Datastore is running. Type:


```
dpa.sh ds status
```

 Output is `INSTALLED`, `STOPPED`, or `RUNNING`.
5. If it is not running, start it. Type:


```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```
6. Start the Application server. Type:


```
dpa.sh app start
```

Reconfiguring Datastores

Use this procedure if you failed over to your Slave Datastore and want to reconfigure the former Master Datastore as a Slave Datastore.

Procedure

1. On the new Master Datastore, use the **addSlave** command with the IP of the new Master Datastore. Type:


```
dpa.sh ds rep --addSlave <ip_of_master>
```
2. Restart the new Master Datastore. Type:


```
dpa.sh ds restart
```
3. Export the new Master Datastore. Type:


```
dpa.sh ds rep --export /export
```

4. Configure the new Slave Datastore as SLAVE. Type:

```
dpa.sh ds rep --role SLAVE <ip of MASTER>
```

5. Stop the Slave Datastore. Type:

```
dpa.sh ds stop
```

6. Import the Master Datastore to the Slave Datastore. Type:

```
dpa.sh ds rep --import /import
```

7. Start the Slave Datastore server. Type:

```
dpa.sh ds start
```

Integrating Slave Datastore after it has been offline

This procedure is applicable if Datastore Replication was previously configured and the Slave Datastore goes down. This procedure is also applicable if you are introducing Datastore Replication into an already operational deployment. You then reintegrate a Slave Datastore.

About this task

Datastore Replication automatically resumes after short amounts of time offline, for example, after a restart of the Application server. The Datastore is configured to allow approximately 6 hours of downtime before it needs reinitialization. However, this value is approximate and a heavily loaded server may require reinitialization if down for less time. We recommend that you carry out testing to determine the threshold for your deployment.

This procedure is also applicable to resynchronizing a standalone Slave Datastore after isolation. Examples of isolation could be a network outage or break down in communications between the Master and Slave Datastores.

Procedure

1. Create an empty directory on the Master Datastore to which to export the Master Datastore file set. For example, `/tmp/export`
2. Export the Master Datastore file set from the running Master Datastore. Type:


```
dpa.sh ds rep --export /tmp/export
```
3. Create an empty directory on the Slave Datastore into which to copy the Master Datastore file set.
4. Use the appropriate platform to command copy the files to the empty directory on the Slave Datastore.
5. Import the Slave Datastore. Type:


```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds rep --import /tmp/import
```

 where `<DPA_Home>` is the location of the DPA installation.
6. Start the Slave Datastore server. Type:


```
<DPA_HOME>/emc/dpa/services/bin/dpa.sh ds start
```

 where `<DPA_Home>` is the location of the DPA installation. The status of the Slave Datastore at this point is STARTED.
7. Verify that replication is functioning. On the Master Datastore, type:

```
bin/dpa.sh ds rep
```

Output such as the following on the Slave Datastore appears: EMC Data Protection Advisor [INFO] Replication State : SLAVE (for 10.11.111.112)
Command completed successfully.

If the Slave has been down and is restarted, output such as the following indicating the bytes lag and status of *catchup* on the Master Datastore appears:

```
EMC Data Protection Advisor

[INFO] Replication State : MASTER
[INFO] Defined Slaves
           : 10.11.111.111/12

[INFO]
LAG           STATUS           SLAVE           BYTES
[INFO]
11245376      catchup           10.11.111.111

Command completed successfully.
```

Once the lag is caught up, output such as the following, with the status showing as *streaming*, appears:

```
EMC Data Protection Advisor

[INFO] Replication State : MASTER
[INFO] Defined Slaves
           : 10.11.111.111/12

[INFO]
LAG           STATUS           SLAVE           BYTES
[INFO]
0             streaming        10.11.111.111

Command completed successfully.
```

Stopping Datastore Replication

To stop Datastore Replication, stop the Slave Datastore. On the Slave Datastore, type `dpa.sh ds stop`.

DPA Database superuser password

One DPA Database superuser account is supplied within DPA Datastore: `apollosuperuser`. `apollosuperuser` is actually the user who owns the DPA Database and who can override all access restrictions within the DPA Database.

By default, the DPA Database is accessible by using that account only from the local machine. The `dpa datastore superpassword` CLI command provides the ability to change the `apollosuperuser` password. The `dpa datastore` commands section of the *Data Protection Advisor Installation and Administration Guide* provides information.

DPA command line operations

Sourcing the DPA config file for UNIX users

A Technical Support Engineer may ask you to source the DPA config file before running any agent binaries (including DPA Agent request in debug mode and bkupjob) and any command line operations on UNIX.

Procedure

1. Navigate to the `/etc` folder of the DPA installation directory.
2. Run the following command :

Results

```
cd <DPA install dir>/agent/etc
. ./dpa.config
```

The DPA config file sets up various environment variables and paths that the DPA agent uses. Running it when instructed ensures that the shell the user is working and has these set correctly. Failure to carry out this procedure when directed by a Technical Support Engineer could result in CLI command failure.

dpa CLI command

In a default DPA installation, the dpa CLI command can be found in `<install_dir>/services/bin` on UNIX and Linux and in `<install_dir>\services\bin` on Windows.

Use the following syntax:

For Windows:

```
dpa <service_part> <command> [options]
```

For Linux/UNIX:

```
dpa.sh <service_part> <command> [options]
```

Where `<service_part>` is Application, Datastore, Agent or service. The service component includes both the Application, Datastore, and Agent services.

```
dpa application <command> [options]
```

```
dpa datastore <command> [options]
```

```
dpa agent <command>
```

```
dpa service <command> [options]
```

The DPA server `start/stop/restart` command applies to whichever services are installed on the current host only. For example, if you run `dpa server stop` on the DPA Datastore, it does not stop services that may be running on the DPA Application server.

Examples of command and option abbreviations

The `dpa` command supports abbreviations of the commands. The following table provides some of the abbreviations. Refer to the specific `dpa` command for available options for that command.

Table 32 Command and option abbreviations

Command and option	Abbreviation
<code>--add</code>	<code>-a</code>
<code>--bind</code>	<code>-b</code>
<code>--cluster</code>	<code>-c</code>
<code>--delete</code>	<code>-d</code>
<code>--help</code>	<code>-h</code>
<code>--master</code>	<code>-m</code>
<code>--pipeline</code>	<code>-p</code>
<code>--platform</code>	<code>-p</code>
<code>tune</code>	<code>tun</code>
<code>dpa application</code>	<code>dpa app</code>
<code>dpa datastore</code>	<code>dpa ds</code>
<code>dpa service</code>	<code>dpa svc</code>

dpa agent commands

Use the `dpa agent` commands to manage the DPA Agent service. The `dpa agent` commands can be applied only to the local agent.

```
dpa agent start
dpa agent stop
dpa agent status
dpa agent restart
dpa agent install
dpa agent uninstall

dpaagent --set-credentials
```

After you start, stop, or restart a service, it may take a number of minutes to complete and may not result in an immediate state change.

dpa agent start

Starts the DPA Agent. The Agent service must be installed and stopped for this command to operate.

```
dpa agent start
```


dpa agent stop

Stops the DPA Agent. The Agent service must be installed and running for this command to operate.

```
dpa agent stop
```

dpa agent status

Displays the status of agent service. For example, RUNNING, STOPPED.

```
dpa agent status
```

dpa agent restart

Restarts the agent service. This command first stops the agent service and then starts the service. The agent service must be running for this command to operate.

```
dpa agent restart
```

dpa agent install

Installs the agent service. The agent service operates as a system-managed service, which is manageable through normal operating system service commands. Management of the lifecycle of the service can also be managed through this command line tool. This command installs the service but does not start the service automatically. If the agent service is already installed this command fails.

```
dpa agent install
```

dpa agent uninstall

Uninstalls the agent service.

```
dpa agent uninstall
```

dpaagent --set-credentials

Sets the DPA Agent Registration password. This command can be found in the following file locations:

- On Unix and Linux: <agent_install_dir>/agent/bin
- On Windows: <agent_install_dir>\agent\bin

```
dpaagent --set-credentials
```

Note the following regarding Agent password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol

- A minimum of 1 numeric symbol
- A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\agent\bin>dpaagent --set-credentials
```

```
DPA
Enter new password for the agent connection.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
- at least 1 special character
- at least 1 digit

Retype new password for the agent connection :
[INFO] Your new password has been applied to the configuration.
[INFO] For this new password to be used you must ensure that all agents use
the same new password value.

Command completed successfully.

Completed in : 1min 25secs
```

dpa application commands

Use the dpa application commands to manage the DPA Application service.

```
dpa application [options]
dpa application agentpwd [options]
dpa application adminpassword [options]
dpa application configure [options]

dpa application dspassword [options]
dpa application demote [options]

dpa application install [options]
dpa application importcertificate [options]
dpa application ping [options]
dpa application promote [options] [<Application Server_IP_Address>]
dpa application restart [options]
dpa application start [options]
dpa application status [options]
dpa application stop [options]
dpa application support [options] <ESRS_IP address>
dpa application tls [options]
dpa application tune <value>MB|GB [options]
dpa application uninstall [options]
dpa application version [options]
```

After you start, stop, or restart a service, it may take a number of minutes to complete and may not result in an immediate state change.

dpa application adminpassword

Resets the DPA Administrator password. You must run the command when the Datastore Service is running.

```
dpa application adminpassword [options]
dpa app pwd [options]
```

Command options

- help (-h) — Displays the help screen
- version — Displays the tool version information
- quiet — Suppresses all output except for warning and error messages

Note the following regarding the Administrator password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa app adminpassword
```

```
DPA
Enter new administrator password.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
- at least 1 special character
- at least 1 digit

Retype new admin password :
[INFO] Your new password has been set.
[INFO] You must restart all DPA application nodes for this new password to be
used.

Command completed successfully.

Completed in : 1min 25secs
```

dpa application agentpwd

Configures the DPA Agent Registration password on the Application side.

```
dpa application agentpassword [options]
dpa app agentpwd [options]
```

Command options

- help (-h) — Displays the help screen
- version — Displays the tool version information
- quiet — Suppresses all output except for warning and error messages

Note the following regarding Agent password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol

- A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa app agentpwd
```

```
DPA
Enter new password for the agent connection.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
- at least 1 special character
- at least 1 digit

Retype new password for the agent connection :
[INFO] Your new password has been applied to the configuration.
[INFO] For this new password to be used you must ensure that all agents use
the same new password value.

Command completed successfully.

Completed in : 1min 25secs
```

dpa application configure

Configures the Application service, including specifying the Datastore and cluster to communicate with. The Application service must be stopped for this command to operate.

```
dpa application configure [options]
dpa app con [options]
```

Command options

--master (-m) <IP_address> — Identifies the datastore with which to communicate.

--bind (-b) <IP_address> — Sets the bind address for the application service

--httpprotocol (-hp) <http status> — Turns HTTP protocol on or off. Possible values are TRUE to turn HTTP protocol on and FALSE to turn HTTP protocol off.

If you run the command without any options, the output shows information regarding how the Application server is currently configured. The Operation Mode in the output identifies whether the application is within a cluster or standalone.

Examples

Output for standalone cluster server:

```
C:\Program Files\EMC\DPA\services\bin>dpa app con
DPA
[INFO] Bind Address      : 0.0.0.0
[INFO] Datastore Service  : 127.0.0.1
[INFO] Operation Mode     : STANDALONE
```

Output for Master:

```
DPA
[INFO] Bind Address      : 0.0.0.0
[INFO] Datastore Service  : 127.0.0.1
[INFO] Operation Mode     : CLUSTER
[INFO] Cluster Role      : MASTER
```

```
[INFO] Cluster Address   : 10.64.213.61
[INFO] Multicast Address : 239.1.2.61
```

dpa application demote

Demotes the application service from a cluster environment. The application service will operate as a standalone object instance. The application service must be installed and stopped for this command to operate.

```
dpa application demote [options]
```

Command options

- help (-h) — Displays the help screen
- version — Displays the tool version information
- quiet — Suppresses all output except for warning and error messages

Examples

```
dpa application demote
dpa app demote
```

dpa application dspassword

Configures the DPA Datastore password.

```
dpa application dspassword [options]
dpa app dspwd [options]
```

Command options

- help (-h) — Displays the help screen
- version — Displays the tool version information
- quiet — Suppresses all output except for warning and error messages

Note the following regarding Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa app dspassword
```

```
DPA
Enter new password for the datastore connection.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
```

- at least 1 special character
- at least 1 digit

```
Retype new password for the datastore connection :
[INFO] Your new password has been applied to the configuration.
[INFO] For this new password to be used you must ensure that all datastore
nodes use the same new password value.
```

```
Command completed successfully.
```

```
Completed in : 1min 25secs
```

dpa application install

Installs the application service. The application service will operate as a system managed service, manageable through normal operating system service commands. Management of the lifecycle of the service can also be managed through this command line tool. This command will install the service, but will not start it automatically. If the application service is already installed this command will fail.

```
dpa application install [options]
```

Command options

--user (-U) (DOMAIN\username) User account having read and write access to the shared path specified. The specified user must have *Log on as a service* Windows permission enabled.

--password (-pass) <password> Password for the user specified (Windows only). If the user has changed the password, they must uninstall and install the Application service again.

--help (-h) Display help screen

--version Display tool version information

--quiet Display warnings and warnings and errors only

dpa application importcertificate

Allows you import your own certificate into the DPA application to encrypt the data rather than using the certificate provided by DPA.

```
dpa application importcertificate [options]
dpa app impcert [options]
```

Command options

--certificatefile (-cf) <certificatefile> — Sets the path of the certificate (X.509 format) to import.

--keystorefile (-kf) <keystorefile> — Sets the path of the keystore that contains the certificate to import.

--alias (-al) <alias> — Sets the certificate alias to use when accessing the given keystore.

--password (-pw) <password> — Sets the password to use when accessing the given keystore.

--quiet — Suppresses all output except for warning and error messages

Examples

```
dpa app impcert -kf "C:\work\new.keystore" -al newkey -pw password
```

dpa application ping

Tests the connection between the application object from which it is sent and the defined Master Datastore service.

```
dpa application ping [options]
dpa app pin [options]
```

Command Options

- help (-h) Display help screen
- quiet Display warnings and errors only

dpa application promote

Promotes the application service to a cluster environment. The application service will operate as a object within a cluster of objects. Management of the lifecycle of the service can also be managed through this command line tool. The application service must be installed and stopped for this command to operate.

```
dpa application promote [options]
```

Command options

- bind (-b) <IP_address> — Sets the bind address for the Application service
- user (-u) <username> — For UNIX: (username) is the user account that has read and write access to the shared folder. If omitted root user is used. For windows: (DOMAIN\Username) is the user account that has read write access to the shared folder. If omitted the local system user is used. This user account must have the Log on as a Service Windows permissions enabled.
- path (-p) <path> — Path that is shared among the clusters
- multicast (-m) <multicast address> Sets the multicast address used by the cluster application nodes to communicate with each other. All the application nodes in the cluster must use the same multicast address
- help (-h) — Displays the help screen
- role (-r) <role> Define the role of the application in cluster. Possible values are MASTER or SLAVE <MASTER_IP>
- quiet — Suppresses all output except for warning and error messages

Examples

```
dpa app promote --bind 192.168.1.0 --role MASTER --user user1 --path \\shared
```

dpa application restart

Restarts the application service. This command first stops the application service and then starts the service. The application service must be running for this command to operate.

```
dpa application restart [options]
```

Command options

- platform (-p) — Includes platform version information

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa application start

Starts the Application service. The Application service must be installed and stopped for this command to operate.


```
dpa application start [options]
```

Command options

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Delays when starting and stopping DPA services

You might experience delays in launching the web console when starting the DPA services. If the DPA services have just been installed, there is a delay of up to 10 minutes in launching the web console. Similarly, if the DPA services are restarted, there might be a delay of about 3 minutes in launching the web console.

 **Note:** The DPA services must be running if you want to launch the DPA web console.

dpa application status

Displays the status of application service. For example, RUNNING (STARTING...), RUNNING, STOPPED

```
dpa application status [options]
```

Command options

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Examples

```
# dpa application status
DPA
The status of the Application Service is RUNNING
```

dpa application stop

Stops the Application service. The Application service must be installed and running for this command to operate.

```
dpa application stop [options]
```

Command options

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa application support

Configures the DPA Application server with EMC Secure Remote Support (ESRS) Gateway.

If you are planning on using ESRS-VE for remote troubleshooting (recommended), ensure that you have the ESRS-VE environment installed and configured before DPA installation. The EMC Secure Remote Services landing page at https://support.emc.com/downloads/37716_EMV-Secure-Remote-Services-Virtual-Edition on EMC Online Support provides more information on ESRS-VE installations. .

```
dpa application support [options]
```

```
dpa app support [options]
```

Command options

- register (-r) <ESRS_IP address> — Registers the DPA Application with ESRS gateway
- update (-u) <DPA_new_IP address> — Updates the ESRS gateway with a new DPA server IP address
- deregister (-d) — Unregisters the DPA Application server from ESRS gateway
- ping (-p) <ESRS_IP address> — Pings to obtain the DPA Application server/node information
- help (-h) — Displays the help screen

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa app support --register 10.11.110.111
```

dpa application tlslevel

Sets the TLS protocol version for the DPA Application services. This command will install the service, but will not start it automatically. If the application service is already installed this command will fail.

```
dpa application tlslevel [options]
dpa app tls [options]
```

Command options

- 1.2 — Set the TLS protocol version for the DPA Application services to TLS version protocol 1.2 only
- 1.0 — Set the TLS protocol version for the DPA Application services to TLS version protocols 1.0, 1.1, and 1.2
- help (-h) — Display help screen
- version — Display tool version information
- quiet — Display warnings and warnings and errors only

Example

```
dpa app tls 1.2
```

dpa application tune

Configures tunable parameters of the Application service for the available host memory resources.

```
dpa application --tune <size> MB|GB
dpa app tune <size> MB|GB
```

Command options

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa application uninstall

Uninstalls the Application service.

```
dpa application uninstall [options]
```

Command options

- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa application version

Displays the version information for the various functional libraries that make up the application service. The functional libraries include Apollo, Controller, DPA (DPA), RemoteX, and UI.

```
dpa application version [options]
```

Command options

- platform (-p) — Includes platform version information
- help (-h) — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Examples

```
# dpa application version
[INFO] Version for      Apollo EAR is 1.0.0.3304
[INFO] Version for Controller RAR is 18.1.xxx
[INFO] Version for DPA EAR is 18.1.xxx
[INFO] Version for      Remotex EAR is 1.0.0.3304
[INFO] Version for          UI WAR is 18.1.x.local
```

dpa datastore commands

Use the dpa datastore commands to manage the DPA Datastore service.

```
dpa datastore [options]
dpa datastore configure [options]
dpa datastore dpassword [options]
dpa datastore export [options]
dpa datastore import [options] <import_filename>
dpa datastore install [options]
```

```

dpa datastore logtz <time zone>
dpa datastore recreate [options]
dpa datastore replicate [options]
dpa datastore restart [options]
dpa datastore start [options]
dpa datastore status [options]
dpa datastore stop [options]
dpa datastore superpassword [options]
dpa datastore support [options] <ESRS_IP address>
dpa datastore tune <size>MB|GB [options]
dpa datastore uninstall [options]

dpa datastore supportbundle [options] <directory of output file>
dpa datastore version

```

After you start, stop, or restart a service, it may take a number of minutes to complete and may not result in an immediate state change.

dpa datastore configure

Configures the Datastore service, including adding or removing an application service to the list of allowed connections to the datastore service.


```

dpa datastore configure [options]
dpa ds configure [options]

```

Command options

--bind <IP_address> — Set the bind address for the Datastore service. The default is 127.0.0.1

 **NOTICE** --bind cannot be specified with --add or --delete.

--add <IP_address> — Add an application service node as a valid Datastore client

--delete <IP_address> — Remove an application service node as a valid Datastore client

--help — Displays the help screen

--quiet — Suppresses all output except for warning and error messages

Examples

```

dpa datastore con --add 111.111.1.1

```

dpa datastore dpassword

Resets the DPA Datastore password. You must run the command when the Datastore Service is running.

```

dpa datastore dpassword [options]
dpa ds pwd [options]

```

Command options

--help (-h) — Displays the help screen

--version — Displays the tool version information

--quiet — Suppresses all output except for warning and error messages

Note the following regarding Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa ds dspassword
```

```
DPA
Enter new password for the datastore connection from the application node.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
- at least 1 special character
- at least 1 digit

Retype new password for the datastore connection from the application node:
[INFO] Your new password has been applied to the datastore.
[INFO] For this new password to be used you must ensure that all DPA
application nodes use the same new password value.

Command completed successfully.

Completed in : 1min 25secs
```

dpa datastore export

Exports the contents of the Datastore to the filename or pipeline specified. The Datastore service must be installed and running for this command to operate. Any existing filename present will be overwritten.

```
dpa datastore export [options]
```

```
dpa datastore export [options] <directory>
```

Command options

--pipeline — Export to pipe

--help — Displays the help screen

--quiet — Suppresses all output except for warning and error messages

Examples

```
C:\Program Files\EMC\DPA\services\bin>dpa datastore export C:\
```

The default filename of the export is: `datastore-<version> <date and time>`.

For example, `datastore-6_2_0_90597-2014-10-01-1135`.

dpa datastore import

Imports the contents of the Datastore export file to the Datastore. The import files must be available on the local filesystem. You will be prompted to stop all Application servers that

communicate with this Datastore prior running the command. The datastore service must be running for the import command to execute.

```
dpa datastore import [options] <filename>
```

Where <filename> is a previously exported datastore file. The import command replaces the existing Datastore contents with the contents contained in the Datastore export file.

Command options

--help — Displays the help screen

--quiet — Suppresses all output except for warning and error messages

<import_filename> — Filename of the exported file to import

Examples

```
# dpa datastore import datastore-2013-02-20-1205
DPA
Datastore imported from file : datastore-2013-02-20-1205
Imported to the datastore successfully
```

dpa datastore install

Installs the datastore service. The datastore service will operate as a system managed service, manageable through normal operating system service commands. Management of the lifecycle of the service can also be managed through this command line tool. This command will install the service, but will not start it automatically. If the datastore service is already installed this command will fail.

```
dpa datastore install [options]
```

Command options

--help — Displays the help screen --version — Displays the tool version information --quiet — Suppresses all output except for warning and error messages

dpa datastore logtz

Configures the DPA Database logs time zone

```
dpa datstore logtz <time zone>
```

```
dpa ds logstz <time zone>
```

Example

dpa datstore logtz 'Europe/Moscow' Configures the DPA Datastore logs time zone to Europe/Moscow

dpa datstore logtz DPA Datastore logs time zone to GMT

dpa datastore recreate

Recreates the datastore, reverting its content to factory settings.

Description

dpa datastore recreate [options]

```
dpa ds rec [options]
```

Command options

- force (-f) — Override prompt that the current Datastore data is going to be overwritten
- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Syntax

dpa datastore replicate

Configures the Datastore service to replicate to another instance.

Description

```
dpa ds rep [options]
```

Command options

- addSlave (-a) <hostname/IP of SLAVE> — Adds a Slave Datastore to a Master Datastore
- deleteSlave (-d) <hostname/IP of SLAVE> — Deletes a Slave Datastore from a Master Datastore
- role (-r) MASTER — Redefines the role of Slave Datastore to Master Datastore
- role (-r) SLAVE <IP of MASTER> — Redefines the role of Master Datastore to Slave Datastore
- failover — Initiates failover between Slave Datastore and Master Datastore
- import (-i) <import> — Initializes a SLAVE datastore with replica located in specified directory
- export (-e) <export> — Produces a clone of the MASTER datastore to specified directory
- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Syntax

dpa datastore restart

Restarts the Datastore service. This command first stops the Datastore service and then starts the service. The Datastore service must be running for this command to operate.

```
dpa datastore restart [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa datastore start

Starts the datastore service. The Datastore service must be installed and stopped for this command to operate.

```
dpa datastore start [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa datastore status

Displays the status of Datastore service. For example, RUNNING (STARTING...), RUNNING, STOPPED

```
dpa datastore status [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Examples

```
# dpa datastore status
DPA
```

The status of the Datastore Service is RUNNING

dpa datastore stop

Stops the Datastore service. The Datastore service must be installed and running for this command to operate.

```
dpa datastore stop [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa datastore superpassword

Resets the DPA Datastore superuser password. The superuser is the user that owns the DPA Database. You must run the command when the Datastore Service is running.

If you use Datastore Replication, you must run this command on all Datastore nodes. Run the command on Master node first, and only then on all other replication Slave nodes.

```
dpa datastore superpassword [options]
dpa ds superpwd [options]
```

Command options

- help (-h)** — Displays the help screen
- version** — Displays the tool version information
- quiet** — Suppresses all output except for warning and error messages

Note the following regarding Datastore password:

- Blank passwords are not supported.
- Minimum length is 9 characters.
- The following are required:
 - A minimum of 1 uppercase and 1 lowercase alphabetic symbol
 - A minimum of 1 numeric symbol
 - A minimum of 1 special character

Example

```
C:\Program Files\EMC\DPA\services\bin>dpa ds superpassword
```

```
DPA
Enter new password for the superuser owning the database.
The password must have:
- at least 9 characters
- at least 1 uppercase letter
- at least 1 lowercase letter
- at least 1 special character
- at least 1 digit

Retype new password for the superuser owning the database:
[INFO] Your new password has been applied to the superuser owning the
database.

Command completed successfully.
```

dpa datastore supportbundle

Gathers support information and stores the DPADPADPA Datastore support bundle zip to the specified directory.

```
dpa datastore supportbundle [options] <directory of output file>
dpa ds supbd [options] <directory of output file>
```

Command options

- help (-h)** — Displays the help screen
- quiet** — Suppresses all output except for warning and error messages

dpa datastore tune

Configures tunable parameters of the datastore service for the available host memory resources and database connections.

```
dpa datastore tune <size>MB|GB [options]
dpa ds tune <size>MB|GB [options]
```

Command options

- connections (-c) <connections> — Maximum number of concurrent Datastore connections allowed
- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa datastore uninstall

Uninstalls the Datastore service.

```
dpa datastore uninstall [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa datastore version

Queries the Datastore version and patch number

```
dpa datastore version [options]
```

```
dpa ds version [options]
```

Command options

- help (-h) — Displays the help screen

dpa service commands

Use the dpa service commands to manage the DPA Application, the DPA Datastore, and the DPA Agent services.

```
dpa service install [options]
dpa service restart [options]
dpa service start [options]
dpa service status [options]
dpa service stop [options]
dpa service uninstall [options]
```

dpa service install

Installs the Datastore service and then the Application service. The services operate as a system managed services, manageable through normal operating system service commands. Management

of the lifecycle of the services can also be managed through this command line tool. This command installs the services but does not start them automatically. If the services are already installed, this command fails.

```
dpa service install [options]
dpa svc install [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa service restart

Restarts the Application and Datastore services. This command stops the Application service, stops the Datastore service, and then starts the Datastore service and Application service. The services must be running for this command to operate.

```
dpa service restart [options]
dpa svc restart [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa service start

Starts the Datastore service and then Application service. The services must be installed and stopped for this command to operate.

```
dpa service start [options]
dpa svc start [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa service status

Displays the status of Application and Datastore services. For example, RUNNING (STARTING...), RUNNING, STOPPED

```
dpa service status [options]
dpa svc status [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Examples

```
# dpa service status
DPA
The status of the Datastore Service is RUNNING
The status of the Application Service is RUNNING (STARTING ...)
```

dpa service stop

Stops the Application service and then the Datastore service. The services must be installed and running for this command to operate.

```
dpa service stop [options]
dpa svc sop [options]
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

dpa service uninstall

Uninstalls the Application service and then the Datastore service.

```
dpa service uninstall [options] <certificate> <key>
dpa svc uninstall [options] <certificate> <key>
```

Command options

- help — Displays the help screen
- quiet — Suppresses all output except for warning and error messages

Loading historical backup job data

The preferred method to gather historical backup data is by using the DPA web console.

Before You Begin

[Gathering historical backup data using DPA web console](#) on page 91 provides more information. After a backup application object is created and requests are assigned, the agent immediately begins gathering data on backup jobs to store in the datastore. However, the agent also can gather data on backup jobs that were run prior to object creation in DPA.

Note: To commit the data to the DPA server, the installed agent must have previously been started and successfully registered with the DPA Server. However, it need not be currently running in order to load the historical data.

Each backup module has an equivalent executable in the installed Agent's bin directory, `<DPA_HOME>/emc/dpa/agent/bin` directory, where `<DPA_Home>` is the location of the DPA installation.

Description

The following example collects backup job data run on an NetWorker server:

Syntax

Example

```
<install_dir>/agent/bin/dpaagent_modnetworker -c -f jobmonitor -t
NetWorkerServer_IP -B "01/01/2012 00:00:00" -E "01/01/2012 00:00:00"
```

Running the executable with the `-?` parameter shows the valid command line options. Module options applicable to the request (eg. `timeformat`) may also need to be specified explicitly on the command line in order to ensure consistent behaviour with "normal" data collection. Specifically, in the case of the DataProtector jobmonitor request, the occupancy option must be specified explicitly if you want historic data to be included in occupancy calculations. The *Data Protection Advisor Data Collection Reference Guide* provides more information on options. The "Job Monitor" section provides more information on the occupancy option.

To load historical backup data, run the agent binary from the command line with the following parameters: You should specifically use:

- `-f <function name>` — Name of data gathering function to execute. Always `jobmonitor`. Mandatory.
- `-t <target host>` — Host address of backup application server. The default is `localhost`.
- `-B <start time>` — Start time from which to gather backup jobs. The format is `dd/mm/yyyy hh:mm:dd`.
- `-E <end time>` — End time from which to gather backup jobs. The format is `dd/mm/yyyy hh:mm:dd`.
The start and end times can also be in Unix epoch time format.

If `<start time>` is specified and `<end time>` is not, `<end time>` is set to the current time. This includes all the backup jobs that ended after `<start time>`.

If `<end time>` is specified and `<start time>` is not, `<start time>` is set to 0. This includes all the backup jobs that end before `<end time>`.

- `-i` — TSM instance name (TSM only).
- `-l <log file name>` - Name and path of the log file to generate when running the command to load historical data.
The default log file location is the location from which the command is run.
- `-U` — Username to connect to the backup application (TSM and Avamar only).
- `-P` — Password to connect to the backup application (TSM and Avamar only).
- `-c` - Commit — Instructs the module to send the data to the DPA Server. Mandatory.

The following example collects backup job data run on an Avamar server:

Example

```
dpaagent_modavamar.exe -f jobmonitor -t De-dup-muc.corp.emc.com -U viewuser -
P viewuser1 -c -B "01/01/2012 00:00:00" -l /tmp/mod_avamar.log
```

Job summary reports

The job summary reports provide overviews of the totals of backup and maintenance jobs (such as all jobs, successful jobs, failed jobs) that have occurred on backup servers. The summary reports rely on the most up-to-date data in the datastore to produce accurate summary results.

Description

While historical backup job data is loading using the agent command line options, summary reports might display inaccurate totals. It is best to wait until all historical job data is loaded before running summary reports for the loaded historical periods.

Syntax

CHAPTER 4

Environment discovery in DPA

This chapter includes the following sections:

- [Configuring the environment for discovery](#) 152
- [Discovering a host or object manually](#) 196
- [About job data gathering after discovery](#) 197
- [Monitored objects and groups](#) 198
- [Configuring policies, rules, and alerts](#) 203

Configuring the environment for discovery

Discovery overview

The diagram below shows the relationship between the DPA Application object and the DPA Agents deployed to monitor your data protection infrastructure.


Some types of devices need to be monitored by using a DPA Agent deployed as a proxy. A proxy is used typically where the object being monitored is hardware and access for agent installation is not possible. Most types of backup managers can be monitored by an agent directly installed on the same host as the backup manager, or remotely by using proxy agent if the backup manager is resource constrained.

DPA is case insensitive with regard to backup pool names. For example, if you define the pools

- test_name
- Test_name
- Test_Name

DPA creates one object in the configuration tree. When you run a report on the scope and select this object, you will see only one set of numbers.

Figure 3 Relationship between DPA Application nodes and DPA Agents monitoring applications



Defining objects to be monitored

To define objects to be monitored in DPA, follow the steps in the following table.

About this task

Table 33 Data monitoring setup summary

Step	Description
Check licenses	Check that the licenses to monitor your device, host, or environment have been purchased and installed.
Install the agent	If you are monitoring the object from a host other than the DPA server host, you need to install the DPA agent. See DPA Agent installation on page 51.
Install third-party binaries or define the object for monitoring	<p>This step is required for remote or agentless (proxy) data collection.</p> <p>You might need to install binaries on the DPA host or the remote agent host to connect to the monitored object. You also might need to define an account or connection on the monitored object.</p> <p>The following sections describes the prerequisite configuration for all objects:</p> <ul style="list-style-type: none"> • Configuring for Replication Analysis on page 184 • Configuration of storage arrays for replication analysis on page 186 • Monitoring of backup applications on page 156 • Monitoring of Databases on page 169 • Monitoring of RecoverPoint on page 185 • Monitoring operating systems on page 181 • Monitoring of tape libraries on page 191 • Monitoring of switches and I/O devices on page 192 • Monitoring of file servers on page 186 • Monitoring of protection storage on page 188 • Monitoring of StorageTek ACSLS Manager on page 190 • Monitoring of disk management servers on page 188 • Monitoring of VMware environment on page 194

Table 33 Data monitoring setup summary (continued)

Step	Description
Create or modify the DPA credential	A credential stores the information used to connect to the monitored object. You might need to modify the default credential or create a new one with the account details from the previous step.
Run the Discovery Wizard	Use the Discovery Wizard to define objects to be monitored. Select Inventory > System > Run Discovery Wizard .
Modify data collection default settings	Review the default retention times for all requests and modify if required. Data collection requests are assigned to the object created by the Discovery Wizard. If you want to modify the default data collection, select Admin > Systems > Manage Data Collection Defaults .
Test data collection	After at least 10 minutes of letting the request run, run a report from the object that should include data (for example, Backup Job Summary or a configuration report).

Before you run the Discovery Wizard

Procedure

1. Check the installed licenses. In the DPA web console, go to **Admin > System > Manage Licenses**.

The options that are available for configuration in the Discovery Wizard depend on the types of licenses that you have installed with DPA. If you do not have the correct license installed, the option to create that device or host is disabled in the wizard.

2. If you are performing discovery on a Linux host, ensure that the *libstdc++.so.6* library is installed on the host.
3. Ensure that you take note of the connectivity details outlined in the following table.

Table 34 Connectivity details for configuring data collection through the Discovery Wizard

Item	Value to note for input in Discovery Wizard
Network Configuration Information for DPA Server or agent if agent is remote to DPA server	
Hostname	Value:
IP Address	Value:
Network mask	Value:
Primary DNS server address	Value:
Secondary DNS server address	Value:

Table 34 Connectivity details for configuring data collection through the Discovery Wizard
(continued)

Item	Value to note for input in Discovery Wizard
Gateway Address	Value:
Time zone	Value:
Credential Information Needed for Discovery of Virtual Disks through SSH	
IP Address of ESX Server	Value:
ESX Server Root Credential	Value:
Credential Information Needed for Discovery of Servers and Arrays	
Server Name/IP	
SSH Credentials	Value:
RPC Credentials	Value:
WMI Credentials	Value:
Solutions Enabler Host Credentials Requires root/administrator credentials	Value:
RPA Credentials	Value:
Credential Information Needed for Monitoring of Oracle Databases	
Oracle username and password required	Value:
Oracle Service Name and Port, specifically the Oracle SID and TNS port	Value:
Oracle Monitor RMAN An oracle user with catalog access to the RMAN schema and the username and password is required	Value:
Oracle Host Name	Value:
Oracle Monitor Schema If multiple RMAN schemas are present on one Oracle SID, then each RMAN schema owner and username and password are required.	Value:
Credential information needed for SQL Server databases	
SQL Database User Account	Value:
SQL Server Instance	Value:
SQL Database Name	Value:
PostgreSQL Credentials	
PostgreSQL User Account (must be a super user)	Value:
Credential information for Backup Servers, Tape Libraries, I/O Devices	
CommVault User Account	Value:

Table 34 Connectivity details for configuring data collection through the Discovery Wizard (continued)

Item	Value to note for input in Discovery Wizard
<p>Avamar User Account</p> <p>As of Avamar 7.1, Avamar no longer ships with a default password for the viewuser account, and the viewuser account password is set by the user during installation Avamar installation. If you are discovering Avamar 7.1 or later, and it was not upgraded from a previous version, you must create a new set of credentials within DPA. Go to Admin > User > Set Credentials.</p>	Value:
HP Data Protector User Account	
IBM TSM host, TSM Instance Name, TSM port and TSM username and password for each TSM instance is required	Value:
Symantec Backup Exec User Account	Value:
Symantec PureDisk User Account	Value:
<p>SNMP community string for Data Domain</p> <p>SSH username and password for Data Domain, preferably a separate username and password than the Data Domain's system administrator default credentials.</p> <p>Both are required because data is collected using both of the mechanisms</p>	Value:
SNMP Community String for EDL	Value:
SNMP String for Fibre Channel Switch	Value:
SNMP Community String for Tape Libraries	Value:
SNMP Community String for IP Switch	Value:

Monitoring of backup applications

This section describes how to monitor backup applications.

Monitoring of CA BrightStor ARCserve

CA BrightStor ARCserve servers are monitored from an agent running on the CA BrightStor ARCserve server or from an agent running on any other Windows computer in the environment.

Before starting the Discovery Wizard for monitoring CA BrightStor ARCserve

Before you begin

- You must know the resolvable hostname or IP address of the ARCserve server.
- When running ARCserve 11.x, the hostname must be the host short name. You cannot use aliases.

Procedure

1. Install the ARCserve Manager on the computer on which the agent is running.
The agent credentials must match the existing ARCserve account.
2. If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for ARCserve, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Monitoring of CommVault Simpana

Monitor CommVault Simpana servers from an agent running on the CommVault Simpana database or from an agent running on any other Windows computer in the environment.

Before starting the Discovery Wizard for monitoring CommVault Simpana

The DPA Agent service must run with a named account if the CommVault SQL Server is using Windows authentication. The named account chosen for the DPA Agent service must have permission for read access to the CommVault SQLServer Database.

About this task

Alternatively, if SQL authentication is used, you must define DPA credentials for the CommVault requests; for example, username: cvadmin; password: password of cvadmin user.

You need to know:

- The resolvable hostname or IP address of the CommVault server.
- The database hostname and instance name if the CommVault database is remote to the server.

If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for CommVault Simpana, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Monitoring of Avamar

Monitor Avamar servers using a DPA agent installed on any remote computer in the environment, including the DPA Server. Do not install a DPA Agent on the Avamar server or storage object.

To enable monitoring of basic Avamar grid on version 7.2 and later, by the supported DPA deployment, ensure that you select **Remote Data Collection Unit**.

To enable the Clone Operations report to display data when the source grid is selected as the scope for the report, you must monitor the source Avamar grid using the Job Monitor request from an Avamar replication setup.

Before starting the Discovery Wizard for monitoring Avamar

No additional software is required to monitor an Avamar server remotely.

Before you begin

Before you start the Discovery Wizard, you need to know the resolvable hostname or IP address of the Avamar server.

Procedure

1. To gather data from Avamar, DPA connects directly to the Avamar database. It connects to the mcdb database on the default port for Avamar, which is 5555. If these parameters were modified, edit the Avamar Configuration, Avamar Job Monitor and Avamar Status request options to specify the database name and port in use. In the DPAA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

2. If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for Avamar, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.
3. If you are discovering Avamar 7.1 or later, and it was not upgraded from a previous version, you must create a new set of credentials within DPA. Go to **Admin > User > Set Credentials**.

As of Avamar 7.1, Avamar no longer ships with a default password for the viewuser account, and the viewuser account password is set by the user during installation Avamar installation.

4. Create new credentials in the Default Avamar Credentials in the DPA web console from **Admin > System > Manage Credentials** as username / password get reset on upgrade.

When DPA connects to the database, it uses the viewuser account to log in to the database.

About job data gathering after Avamar discovery

Read about Avamar job data gathering after you discover Avamar within DPA.

- When a new Avamar server is discovered, DPA gathers job data from 14 days before.
- Each time the Jobmonitor request is run DPA gathers at most a "batch period" amount of data. This value is configurable and defaults to one day's worth of data.
- After multiple Jobmonitor requests have been run, the time period of gathered jobs catches up to the present time and new backups are gathered.
- The default time between the end of the last Jobmonitor to when a new Jobmonitor request is run, is 5 minutes. This is configurable as with all requests.

[Data Collection Request Options by Module](#) provides more information.

Monitoring of NetWorker

Monitor NetWorker either from an agent running on the backup server or remotely using an agent running on the DPA Server or any other remote computer in the environment.

Before starting the Discovery Wizard for monitoring NetWorker

If monitoring NetWorker remotely, the NetWorker client package must be installed on the agent's host. The NetWorker module uses commands such as `jobquery` and `nsradmin` to communicate with the NetWorker server and requires access to the binaries within the NetWorker client package.

Before you begin

- Before you start the Discovery Wizard, you need to know the resolvable hostname or IP address of the NetWorker server.
- If you are monitoring NetWorker 9.0.0.4 and later, ensure that you have the NetWorker server credentials. You will be prompted to enter the NetWorker server credentials to allow the DPA Agent to issue an `nsrauth` and to run `nsradmin`.

Procedure

1. If you are monitoring NetWorker 9.0.0.4 and later remotely, install NetWorker Client and NetWorker Extended Client. The NetWorker 9 Client and Extended Client must be installed on the DPA Agent host. If you have a previous version of the NetWorker Client, then you need to upgrade. If you are monitoring older versions of NetWorker, use the NetWorker9 Client and Extended Client to monitor those other versions if the DPA Agent is used to also monitor a NetWorker 9 server.

2. If you are monitoring NetWorker 7.6 or later remotely, the DPA user and the proxy host must be added to the Users list of the NetWorker Administrators User Group. For example, if you are monitoring NetWorker remotely from the host DPA Agent Host and the agent is running as the Windows user DPAAgent, you must add the following line to the Users list of the properties for Administrators:

```
user=DPAAgent,host=DPAAgentHost
```

3. If you DPA to collect job data from 14 days before, and for the reports show data straight away for NetWorker, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

About job data gathering after NetWorker discovery

Read about NetWorker job data gathering after you discover NetWorker within DPA.

- When a new NetWorker server is discovered, DPA gathers job data from 14 days before.
- After you run multiple Jobmonitor requests the time period of gathered jobs catch up to the present time and new backups are gathered.

As a result of this operation, it will take 7 hours for the jobmonitor request to start gathering current job data. This is because each request is scheduled by default to run every 30 minutes and in each request a maximum of 1 day's data is gathered. [Data Collection Request Options by Module](#) provides more information.

Monitoring of HP Data Protector

An agent can monitor HP Data Protector servers running on the HP Data Protector Cell Manager or remotely from another computer.

Before starting the Discovery Wizard for monitoring HP Data Protector

About this task

If monitoring a Cell Manager remotely, follow the same instructions as documented in [Monitoring HP Data Protector remotely](#) on page 161.

Note: You cannot assign the status request when monitoring the HP Data Protector server remotely because it relies on a the `omnisv` command. The command is only available on the Data Protector server.

If you are monitoring a Data Protector environment that uses the Manager of Managers option, you must configure DPA as if monitoring a remote Data Protector server.

To monitor HP Data Protector remotely, you must install the HP Data Protector client software on the agent's host and configure the client on the Data Protector Cell Manager so that it has permission to run reports. [Monitoring HP Data Protector remotely](#) on page 161 provides information on testing connectivity from the agent host.

If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for HP Data Protector, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.


Gathering occupancy data

Gathering occupancy data is not enabled by default for HP Data Protector. To enable occupancy data gathering, you must enable the occupancy option for the DataProtector Jobmonitor request

and assign the the DataProtector Client Occupancy request to the Data Protector client in the **Edit Request** dialog.

About this task

You can use the `DP_OCCUPANCY_DB_PATH` environment variable for the DPA Agent to control where the occupancy data is stored when you run the jobmonitor request. If you do not use the `DP_OCCUPANCY_DB_PATH` environment variable, then the system stores the occupancy data in the temporary directory.

 **Note:** Gathering occupancy information for HP DataProtector can have a significant performance impact on the Data Protector server.

Changing the location of Occupancy database on Linux Procedure

1. Stop the DPA Agent.
2. Use the `cd` command to access the `/opt/emc/dpa/agent/etc` directory.
3. Edit the `dpa.custom` file. Add the following to the end of the file:

```
COLLECTOR_DP_OCCUPANCY_DB_PATH=/your/absolute/path/
export COLLECTOR_DP_OCCUPANCY_DB_PATH
```

Ensure that you include the trailing backward slash (/) character in the path.

4. Restart the DPA Agent

Changing the location of Occupancy database on Windows

About this task

Procedure

1. Stop the DPA Agent.
2. Run the `regedit.exe` as the administrator user.
3. Expand the `HKEY_LOCAL_MACHINE` registry key.
4. Expand the `SOFTWARE` registry key.
5. Create an EMC registry key if one does not already exist.
6. Create a DPA registry key if one does not already exist.
7. Ceate an Agent registry key if one does not already exist.
8. Create a new String registry value with name `DP_OCCUPANCY_DB_PATH` and set the value to the desired directory path.

For example: `C:\DPA\OccupancyData\` Ensure that you include the trailing slash (\) character in the path.

9. Restart the DPA Agent.

omnirpt patch

HP has released a patch for Data Protector 6.1 that must be installed on a Data Protector 6.1 installation before it can be supported by DPA.

The following table lists the required patch ID by platform.

Table 35 HP Data Protector 6.1 patch IDs

Platform	Patch ID
Windows	DPWIN_00417
HPUX PA-Risc	PHSS_39512
HPUX IA64	PHSS_39513
Linux	DPLNX_00077
Solaris	DPSOL_00371

The patch is available for General Release from HP from www.hp.com. Type the patch ID into the Search field of the HP home page. You are directed to the patch download page.

Configuring restore job data and updated occupancy retention times

Carry out the following procedure to obtain Jobmonitor function restore job data and updated occupancy retention times.

Procedure

1. In the HP Data Protector Manager UI, go to **Internal Database > Global Options**.
2. Add the following options:

Option	Description
EnableRestoreReportStats	Enable extended restore session data
LogChangedProtection	Log occupancy changed retention

Ensure that you set the Value for both options to **1** and select **In Use** for both.

3. Restart the HP Data Protector services with the `omnisv` command for the changes to take effect.

Monitoring HP Data Protector remotely

You must install the client software on the computer that monitors the Cell Manager:

Procedure

1. Launch the Data Protector Manager administration GUI to add a client.
2. When selecting the software components to install on the client, ensure that the **User Interface** option is selected.

The DPA Data Protector module requires access to commands such as `omnirpt` and `omnicellinfo` to gather data from the Cell Manager. These components are only installed when the user interface component is installed, so it is essential to select this option.

3. Configure the client to have permissions to run reports on the Cell Manager. First determine the user for which the Agent process will be running:
 - On UNIX systems, the Agent always runs as the root user.
 - On Windows systems, the Agent runs as the DPA Agent service user. To verify the user for the service on a Windows system, launch the Windows service control manager and view the details of the DPA Agent service.
4. Create a user on the Cell Manager that matches the Agent's username. Type the name of the host in the **user definition** field.

5. Add the user to a Data Protector User Group that has **Reporting and Notifications** and **See Private Objects** permissions.

Typically, this means adding the user to the admin group. However, to restrict a user from inheriting other administrator privileges, create a new group with Reporting and Notification and See Private Objects permissions and add the user to that group.

6. Verify that remote authentication privileges are set up correctly by running the following command from the Agent's host:

```
omnirpt -tab -report list_sessions -timeframe 06/01/01 12:00
06/01/30 12:00
```

If successful, this command returns a list of all the sessions that have run on the Data Protector server during the time period specified. If an error indicating insufficient permission to run reports appears, review the configuration settings on the Data Protector server.

Monitoring of IBM Tivoli Storage Manager (TSM)

Monitor a TSM server from an agent running on the TSM Server or remotely from an agent running on a different host, such as the DPA server. If you are monitoring TSM remotely, follow the instructions in [Monitoring TSM remotely](#) on page 163 before configuring the server in DPA.

Before starting the Discovery Wizard for monitoring TSM

The TSM Credential must use the name and password of a TSM Administrator. The Administrative user does not need full system privileges: Analyst or Operator privileges are sufficient.

Procedure

1. If the Server being monitored is a shared Library Client, set the agent using the following DPA environment variables (UNIX) or registry settings (Windows) to query the Server's Library Manager to gather certain data:
 - AGENT_TSM_LIBMGRUSERNAME
 - AGENT_TSM_LIBMGRPASSWORD

By default, the agent uses the same credentials used to query the Library Client to query the Library Manager.
2. If you want DPA to collect job data from 14 days before, and for the reports show data straight away for TSM, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.
3. Select **Admin > System > Manage Credentials** to modify the TSM Credentials that are created after you have used the Discovery Wizard to create a TSM object.

Gresham Clareti EDT

In Tivoli Storage Manager environments that use Gresham Clareti EDT for device control, DPA communicates with EDT to gather device configuration information by reading information from two files:

- elm.conf
- rc.edt

DPA reads from `elm.conf` at the following location:

- On Windows, an environment variable called *EDT_DIR* is set by EDT. DPA looks up the location specified in *EDT_DIR*.
- On Unix, DPA looks first in `/opt/GEsedt-acsls/bin` for `elm.conf`. If not found, on AIX DPA looks in `/usr/lpp/dtelm/bin`. On other flavours of UNIX/Linux, DPA looks in `/opt/OMIdtelm/bin`.

If the `elm.conf` file is not present in these directories, the registry variable (Windows) or environment variable (UNIX) `AGENT_TSM_ELMCONF_FILENAME` can be set to the location of `elm.conf` if required.

DPA reads from the `rc.edt` file at the following location:

- On Windows, DPAA looks up the location specified in the environment variable *EDT_DIR*.
- On UNIX, DPA looks first in `/opt/GEsedt-acsls/SSI` for `rc.edt`. If not found, on AIX DPA looks in `/usr/lpp/dtelm/bin`. On other flavours of UNIX/Linux, DPA looks in `/opt/OMIdtelm/bin`.

If the `rc.edt` file is not present in these directories, the registry variable (Windows) or environment variable (UNIX) `AGENT_TSM_RCEDT_FILENAME` can be set to the location of `rc.edt` if required.

Note: Because a TSM environment using EDT requires the agent to read from these files to collect configuration data, the agent must be on the same server as the TSM server.

Monitoring TSM remotely

When monitoring a TSM instance remotely, you must install the TSM client software on the host that will monitor the TSM instance. The TSM module uses the `dsmaadm` command included with the TSM client software to connect to the TSM instance and gather data.

About this task

In a default TSM Client installation on a Windows computer, the administrative components required by DPA are not installed. To install the administrative components:

Procedure

1. Click **Custom** when prompted during the TSM client installation.
2. Select **Administrative Client Command Line Files** and click **Next**.
The TSM client installation continues.
3. After the TSM client installation is complete, initialize the client for the first time by starting the TSM Backup-Archive GUI from the **Start** menu. Use the wizard to configure the client.
4. To configure the client, accept the default **Help me configure the TSM Backup Archive Client** value and click **Next**. Either import an existing options file or create a new one when prompted.
5. Accept the default value **Create a new options file**. You must create a blank options file called `ddsm.opt` in the `baclient` directory under the install directory for TSM (default `C:\Program Files\Tivoli\TSM`).
6. Continue to progress through the wizard. Complete all of the windows in the wizard until a new options file is created.

About job data gathering after TSM discovery

Read about TSM job data gathering after you discover TSM within DPA.

- When a new TSM server is discovered, DPA gathers job data from 14 days before.
- The next time the Job Monitor request runs, the current poll time is set to the next day and data is collected for the next day.

- The current poll time is advanced one day at a time from 14 days back every time the Job Monitor request runs, collecting the data for that day until two weeks of data has been collected. Data collection resumes as normal from then on.
- The poll time default value is 1 day and is user-configurable under the TSM Job Monitor request options section.

[Data Collection Request Options by Module](#) provides more information.

Monitoring of Symantec Backup Exec

Monitor Symantec Backup Exec servers from an agent running on the Backup Exec server or from an agent running on any other Windows computer in the environment. The DPA Agent service needs to run with a named account that can authenticate with the BackupExec server.

Monitoring of backup servers in a Symantec Cluster Server and Microsoft Cluster Server environment

This section provides configuration information for monitoring backup servers in Symantec Cluster Server and Microsoft Cluster Server (MSCS) environments.

Supported platforms

- Symantec Cluster Server is supported on Linux and Solaris
- MSCS is supported on Windows

The *Data Protection Advisor Software Compatibility Guide* provides more information on supported platform versions.

Monitoring backup applications configured as part of a cluster

You can monitor your backup applications that are configured as part of a cluster in a couple of ways.

About this task

To monitor to a backup application in a cluster environment:

Procedure

1. Install a remote Agent on a system outside of the cluster. Ensure that:
 - the Agent can access the virtual server of the cluster using the required ports.
 - the Agent has any required backup application binaries installed.
2. Discover the virtual server of the cluster by using the DPA Discovery Wizard.
3. Collect data by using the remote Agent.

Results

In this configuration if the server fails over, the cluster name always resolves and provides the backup data.

Alternative procedure for monitoring backup applications configured as part of a cluster

To monitor a backup application in a cluster environment as well as monitor the local host resources

Procedure

1. Install a local agent on each host in the cluster for host monitoring only.
2. Select one of the agents on the physical servers to monitor the virtual server.

Before starting the Discovery Wizard for monitoring Symantec Backup Exec

To monitor a Symantec Backup Exec backup server remotely, the agent must run as a named user account rather than the Local System account. When installing the agent, you are prompted to specify whether the agent runs using the Local System account or as a named user.

About this task

The Backup Exec Credentials must use the username and password of a Windows administrator account on the Backup Exec server.

Select **Admin > System > Manage Credentials** to modify the Backup Exec Credentials that are created after you have used the Discovery Wizard to create a Backup Exec object.

Monitoring Backup Exec Remotely

To verify that the agent is running, launch the Windows Service Control Manager (**Start > Settings > Control Panel > Administrative Tools > Services**). Right-click on the DPA agent service and select Properties:

About this task

Procedure

1. Select the **Log On** tab of the Service Properties panel.
2. Select **This Account**.
3. Type the username and password of the local administrator account to run the service.
4. Modify the service account details and click **OK**.
5. Restart the service to activate the changes.

Monitoring of Symantec NetBackup

Configure a Symantec NetBackup server to be monitored from an agent running on the NetBackup Master Server or from an agent running on a different host, such as the DPA server.

When monitoring Symantec NetBackup from a proxy Agent, a proxy Agent can monitor NetBackup master servers that are within the same NetBackup Media Manager (EMM) domain. This means that an Agent is required for each EMM Domain.

Before starting the Discovery Wizard for monitoring Symantec NetBackup

Media Server Status data can only be collected if an agent is installed on the Media Server itself. It cannot be collected through proxy.

About this task

You must specify the `timeformat` option in the jobmonitor request for gathering openfiles, errors, and mount information. For example, "%m/%d/%Y %T"

If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for NetBackup, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Configuring NetBackup authentication for remote data collection

To gather data remotely, the following must be configured:

About this task

- The NetBackup Remote Administration Console, a component of the NetBackup Server software, must be installed on the agent's host.
- The agent's host must be able to successfully resolve the NetBackup Master Server.

- The NetBackup Master Server must be able to successfully resolve the agent's host.

The following sections describe how to resolve the agent host from the NetBackup Master Server on UNIX and Windows.

Configuring NetBackup authentication for remote data collection on UNIX

If the NetBackup Master Server is running on a UNIX computer, you must add the name of the host on which the agent is running to the `bp.conf` file on the NetBackup Master Server.

About this task

To add the host:

Procedure

1. Open `/usr/opensv/netbackup/bp.conf` for editing and add the following line:

```
SERVER = Agenthost
```

where *Agenthost* is the agent's hostname. The agent's hostname must be resolvable by the Master Server.

2. Restart NetBackup on the Master Server for the changes take effect.

Configuring NetBackup authentication for remote data collection on Windows

If the NetBackup Master Server is running on a Windows computer, add the name of the agent host through the NetBackup Administration Console:

Procedure

1. On the NetBackup Server, launch the **NetBackup Administration Console** and open the **Master Server Properties** dialog box:
 - Select **Netbackup Management > Host Properties > Master Servers**.
2. Double-click **Host** in the right-hand panel.
3. In **Master Servers Properties**, **Servers** field, type the name of the agent host to the list of additional servers that are allowed to access the Master Server.
4. Click **OK**.
5. Restart the NetBackup services. Alternatively, reboot the machine to activate the changes.

Monitoring of Symantec PureDisk

Configure a Symantec PureDisk server to be monitored from an agent running on the PureDisk Server or from an agent running on a different host. Symantec PureDisk can only be monitored on SUSE Linux 10. The root user cannot be used to gather data from PureDisk.

Before starting the Discovery Wizard for monitoring Symantec PureDisk

PureDisk servers implement a firewall that might prevent DPA from gathering data from PureDisk or from communicating with an agent installed on the PureDisk server. To ensure successful data gathering and communications, the following sections describe how to configure the PureDisk server before configuring the server in DPA.

About this task

The configuration process depends on the version of PureDisk being monitored.

Manually configuring the firewall (versions of PureDisk earlier than 6.5)

Procedure

1. Log on to the PureDisk server as the root user.
2. Stop the PureDisk firewall by running the following command:

```
/etc/init.d/pdiptables stop
```

3. Edit the file `/etc/puredisk/iptables-rules` by inserting one of the following lines directly after this line in the file:

```
-A INPUT -p icmp -j ACCEPT
```

Note: It is important that the line is inserted at the correct location in the file, otherwise it might not take effect.

- If you are monitoring PureDisk with an agent installed on the PureDisk server, add the following line:

```
-A INPUT -p tcp -m tcp --dport 3741 -j ACCEPT
```

- If you are monitoring PureDisk from an agent running on a different host, add the following line:

```
-A INPUT -p tcp -m tcp --dport 10085 -j ACCEPT
```

4. Restart the PureDisk firewall by running the following command:

```
/etc/init.d/pdiptables start
```

Updating the IP tables rules (PureDisk version 6.5)

Manually configuring the firewall will not work for PureDisk version 6.5. To update the PureDisk IP table:

About this task

Procedure

1. Open the following file in a text editor:

```
/etc/puredisk/custom_iptables_rules
```

2. If the DPA agent is installed on the PureDisk server, add the following line to the rules file (three columns separated by a tab):

```
tcp      {controller_host_ip}    3741
```

This allows connections from the controller host to the DPA agent on port 3741 on the PureDisk server.

3. If the DPA agent is installed on a remote host, add the following line to the rules file (three columns separated by a tab):

Results

```
tcp      {agent_host_ip}          10085
```

This allows connections from the agent host to the postgres database on port 10085 on the PureDisk server.

You can specify a single host or an entire subnet (by including a `/mask`), as in the following example:

```
tcp      10.64.205.0/24            10085
```

The `/etc/puredisk/custom_iptables_rules` file provides additional information on configuring this file.

Monitoring of VMware vSphere Data Protection

Monitor VMware vSphere Data Protection (VDP/A) servers using a DPA Agent installed on any remote computer in the environment, including the DPA Server.

Do not install a DPA Agent on the VMware vSphere Data Protection server.

Before starting the Discovery Wizard for monitoring VDP/A

No additional software is required to monitor a VMware vSphere Data Protection server remotely.

Before you begin

Ensure that you know the resolvable hostname or IP address of the VMware vSphere Data Protection server.

About this task

To gather data from a VMware vSphere Data Protection server, DPA connects directly to the VDP/A database. It connects to the database on the default port, which is 5555. The port is not configurable.

For monitoring of VDP 5.5, 5.8, and 6.0

Procedure

1. Edit the `postgresql.conf` file. Uncomment line in the following and change `localhost` to `localhost, Agent_IP_Address`

```
vi /data01/avamar/var/mc/server_data/postgres/data/postgresql.conf
listen_addresses='localhost,Agent_IP_Address'
```

2. Edit the `pg_hba.conf` file. Add the second line:

```
vi /data01/avamar/var/mc/server_data/postgres/data/pg_hba.conf
host all all Agent_IP_Address/0 trust
```

3. Edit the `firewall.base`, `vi /etc/firewall.base`.
 - a. Enable remote access to Postgres db service.
 - b. Add the following lines to the bottom of the `firewall.base` file:

```
iptables -I INPUT 1 -p tcp --dport 5555 -j ACCEPT
iptables -I INPUT 1 -p tcp --dport 5558 -j ACCEPT
```

4. Reboot the VDP appliance.

Monitoring of Data Domain Backup Enterprise Applications

DPA supports Data Domain Backup Enterprise Applications (DDBEA) for backing up databases without the use of another backup application, such as backing up Oracle RMAN without the use

of NetWorker. The EMC Data Protection Advisor Software Compatibility Guide provides information on supported databases.

If monitoring the Enterprise App for backing up Oracle RMAN, follow the procedure provided in [Monitoring of Oracle and Oracle RMAN](#) on page 172.

If monitoring the Enterprise App for backing up Microsoft SQL Server, follow the procedure provided in [Monitoring of Microsoft SQL Server](#) on page 171.

If monitoring the Enterprise App for backing up PostgreSQL, follow the procedure provided in [Monitoring of PostgreSQL](#) on page 178.

If monitoring the Enterprise App for backing up SAP HANA, follow the procedure provided in [Monitoring of SAP HANA](#) on page 179.

Monitoring of Databases

This section describes how to monitor databases.

Monitoring of DB2

A DB2 database can be monitored from an agent running on the same host as the DB2 server, or from an agent running on a different host, such as the DPA server. The DPA Agent must be run on Windows or Linux.

Before starting the Discovery Wizard for monitoring DB2

For DPA Agent to collect data from DB2 database, you must copy the DB2 client .jar file to the DPA plugins directory.

Procedure

1. Create a directory called *plugins* under `<DPA_install_dir>\agent\`.
2. Copy the DB2 client jar file *db2jcc4.jar* to the *plugins* folder under `..\EMC\dpa\agent\`.

For the custom location or path add following tag: `<PLUGINS_DIR>path </PLUGINS_DIR>` in `dpaagent_config.xml` located under `<DPA_install_dir>\agent\etc`

where *path* is the path of the directory created in step 1.

For example `<PLUGINS_DIR>c:\program files\emc\dpa\agent\plugins</PLUGINS_DIR>`

3. If you DPA to collect job data from 14 days before, and for the reports show data straight away for DB2, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Permissions

Ensure you have the correct permissions to gather data on DB2.

About this task

Ensure that you have Select operations privileges on:

- the `sysibmadm.db_history` view.
- the `<user_name>.UTILSTOP_DPABACKUP` and `sysibm.syscolumns` tables. This is required for version DB2 version 11.1.1.1 and later.

Configuring DB2 to show size field in Backup All Jobs report

You must create the DB2 EVENT MONITOR DPABACKUP on the DB2 database itself for the DPA Agent to send data to the DPA server with the DB2 backup size value.

Before you begin

- DPA supports calculating the backup size only for DB2 version 11.1.1 and later.
- The event monitor must be created by the same user whose credentials are assigned to the DB2 Jobmonitor request.

About this task

Carry out this procedure on the DB2 database itself. For information on how to carry out these steps on DB2, consult vendor documentation.

Procedure

1. Create event: `CREATE EVENT MONITOR DPABACKUP FOR CHANGE HISTORY WHERE EVENT IN (BACKUP) WRITE TO TABLE autostart`
2. Turn on the event monitor.
3. Set the event monitor to `DPABACKUP state 1`.
4. Verify that the event has been created correctly. Carry out the backup database online.
`Type: backup database sample online`
The new record should be present in the table.
5. Select `*from UTILSTOP_DPABACKUP`.

About job data gathering after discovery

Read about job data gathering after you discover some applications within DPA.

The information in this section applies to the following applications:

- NetWorker
- Avamar
- TSM
- HP DataProtector
- Commvault Simpana
- NetBackup
- ArcServ
- DB2
- SAP HANA
- RMAN
- MSSQL

With regard to the applications above, note the following:

- When a new server is discovered, DPA gathers job data from 14 days before if you enable this feature.
- The next time the Job Monitor request runs, the current poll time is set to the next day and data is collected for the next day.
- The current poll time is advanced one day at a time from 14 days back every time the Job Monitor request runs, collecting the data for that day until two weeks of data has been collected. Data collection resumes as normal from then on.

- The poll time default value is 1 day and is user-configurable under the Job Monitor request options section.
- When setting data collection, the **Frequency** must always be a lower value than **max data time range each request will gather from**. Otherwise, request does not catch up to the current time and each time the request runs, it falls further behind and does not gather remaining data.

[Data Collection Request Options by Module](#) provides more information.

Monitoring of Microsoft SQL Server

Monitor Microsoft SQL Servers from an agent running on the SQL Server database, or from an agent running on any other Windows computer in the environment. The DPA Agent service needs to run with a named account that can authenticate with Microsoft SQL Servers.

Ensure that you specify the firewall inbound rules to allow incoming connections to SQL Server Browser service SQLBrowser.exe. It uses UDP port 1434.

Before starting the Discovery Wizard for monitoring Microsoft SQL Server

To connect to SQL Server using Windows Authentication, the DPA Agent must run as a named user with MS-SQL access and not as the Local System Account. Verify that the service is running as the correct user before proceeding with the configuration of the database.

About this task

To monitor clustered SQL Server installations, set DPA to monitor it as a remote target even if the DPA Agent is installed locally on a physical node of the cluster. The target name should be set to the cluster alias name.

Ensure that the DPA Agent has read access to both the DPA Master and the MSDB databases during the DPA discovery test, even if you do not select database monitoring.

Agent requirements for monitoring Microsoft SQL Server

The agent needs to be able to connect to the SQL Server master database in order to gather the data required. The agent can either:

- Use SQL Server Authentication using the credentials of the request (if set).
- Use SQL Server Authentication using the credentials against an explicit master database in the list of databases to be monitored (if set)
- If these are not set, the agent uses Windows Authentication using the logon ID of the agent process.

If none of these are sufficient to connect to the master database, the request will not gather data.

User account requirements for monitoring Microsoft SQL Server

To gather data successfully, the user account used to connect to the SQL Server database must be granted specific privileges. Any SQL Server user with dbo access will have the correct privileges by default.

If you do not want to connect with a user with dbo access, configure a user with the following:

- Map the user to the database with the public role.
- Grant explicitly the VIEW SERVER STATE and VIEW DEFINITION privileges (SQL Server 2005 only).

The VIEW SERVER STATE privilege is granted at the server level. The VIEW DEFINITION privilege might be granted at the server level (under the name VIEW ANY DEFINITION) or at the database, schema, or individual object level.

- Grant explicitly the EXECUTE permission of the system stored procedure `xp_readerrorlog`.

SQL Server 2005 and 2008

To grant server-wide privileges to the SQL Server login used by the agent, including VIEW DEFINITION privileges for all database tables, connect to the SQL Server as an administrator and run:

```
GRANT VIEW SERVER STATE TO <login\domain> GRANT VIEW ANY DEFINITION TO <login\domain>
```

However, to grant VIEW DEFINITION privileges for only the specific databases that you want to monitor, connect to the SQL Server as an administrator and run:

```
GRANT VIEW SERVER STATE TO [login\domain] GRANT VIEW DEFINITION ON DATABASE :: <dbname> TO <username>
```

To grant the EXECUTE permission of the system stored procedure `xp_readerrorlog` run:

```
USE Master GO GRANT EXECUTE ON OBJECT::sys.xp_readerrorlog TO ddDBO GO
```

Monitoring Microsoft SQL Server for replication analysis

The DPA server must connect as a database user with connect privileges for all of the databases and write privilege for the TEMPDB database. For Windows authentication, the user must be able to connect to all SQL Server databases and should have write privilege for the TEMPDB database.

Enable support of TLS 1.2 only

To enable TLS 1.2 only, the DPA Agent must use ODBC driver, which supports TLS version 1.2.

To use concrete ODBC driver, add the new string value `MSSQLSERVER_DRIVER` in registry:`HKEY_LOCAL_MACHINE\SOFTWARE\EMC\DPA\AGENT` . That value must contain the name of the installed ODBC Driver which supports version TLS 1.2. For example `SQL Server Native Client 11.0`

Monitoring of Oracle and Oracle RMAN

DPA can collect data from two parts of Oracle: from the Oracle database itself, where it collects metrics about the database instance; and from Oracle RMAN. In both cases, you must install Oracle client software.

DPA does not ship Oracle client (OCI) libraries with the DPA Agent. You can download the Oracle Instant Client software from oracle.com for the platform/OS you are installing on. Ensure that the architecture version matches your OS as well as Oracle versions. For example, to collect data from Oracle 12c database, use the Oracle 12c instant client version. If you are collecting from mixed Oracle versions, use the latest version in your environment for the instant client. For the DPA Agent to collect data from an Oracle database or Oracle RMAN, DPA requires the following libraries for Oracle:

- `libociei.so`
- `libocci.so`
- `libclntsh.so`

You must create a symbolic link for the `libclntsh.so` library to the current Oracle build directory. [Creating symbolic link for current Oracle build directory on UNIX](#) on page 173 provides information.

You must manually copy it into `AGENT_ORACLE_CLIENT_PATH` in order to work with the DPA Agent.

On Windows this is `OCI.DLL` and on UNIX, it is `libclntsh.so`.

Note: The library must be for the same platform as the DPA Agent. Example, if a 64-bit Windows DPA agent is installed, then you must use the 64-bit Windows Oracle library.

You can download the Oracle Database Instant Client at <http://www.oracle.com/technetwork/database/features/instant-client/index.html>

While installing the DPA Agent, you are prompted to specify if you want to utilize the Agent to monitor Oracle and if so, provide the location of the Oracle client libraries. On Windows, this action sets a registry setting and on UNIX modifies an environment variable in the `dpa.config` file. If you change the location of the libraries after the install process is completed, then you need to perform these steps manually.

Creating symbolic link for current Oracle build directory on UNIX

You must create a symbolic link for the `libclntsh.so` library to the current Oracle build directory. You must manually copy it into `AGENT_ORACLE_CLIENT_PATH` in order to work with the DPA Agent.

Procedure

1. Install using `rpm` command. Run: `rpm -i oracle.instantclient<version.build.architecture>.rpm`.
For example: `rpm -i oracle.instantclient12.1-basic-12.1.0.2.0-1.x86.rpm`
The output of `/usr/lib/oracle/12.1/client64/lib` shows -shows the latest Oracle client. For example, `libclntsh.so.12.1`.
2. Create the symbolic link for `libclntsh.so` and add execution permission on the files. Run: `ln -s libclntsh.so<version.build.architecture> libclntsh.so chmod 755 *`
For example: `ln -s libclntsh.so.21.1 libclntsh.so chmod 755 *`
3. Verify that the current Oracle build is created in `/usr/lib/oracle` (http://docs.oracle.com/cd/B19306_01/server.102/b14357/ape.htm)

Windows

Procedure

1. Update the registry entry with the location of the Oracle instant client software:
 - a. Navigate to the folder where the Oracle client software is located.
 - b. Use `regedit` to manually edit the location of the Oracle instant client software.

Manually configuring DPA Agent to monitor Oracle database and Oracle RMAN

About this task

- To manually configure the DPA Agent to monitor Oracle RMAN: On Windows, set the "HKLM/Software/EMC/DPA/Agent" registry of value type `REG_SZ` as follows:

Value name: `ORACLE_CLIENT_PATH`


Value data: `<directory containing the Oracle client libraries - oci.dll>`

Note: The registry key is created if you have selected the Oracle database to be monitored option while installing the DPA Agent. If the registry key is not created, you must create it manually.

- On UNIX, modify the `dpa.config` file

The `dpa.config` file is available in `<installdir>/agent/etc/dpa.config`. Search for line `AGENT_ORACLE_CLIENT_PATH=` and set the variable to the directory containing the Oracle client libraries - `libclntsh.so`.

Restart the Agent service if you have changed the `dpa.config` file to include the Oracle client path.

 **Note:** Ensure that you discuss RMAN licensing requirements with your EMC Account Representative.

Before starting the Discovery Wizard for monitoring Oracle

To monitor an Oracle database for data protection data, the agent must connect to the database as an Oracle user.

Before you begin


DPA does not require the operating system password to the Oracle server. DPA requires the Oracle username/password used for the RMAN catalog or system catalog queries only.

About this task

To gather data successfully for Oracle databases, this user must be able to create and drop global temporary tables, and to perform selects on the following tables and views:

- V_\$INSTANCE
- V_\$PROCESS
- V_\$DATABASE
- V_\$PARAMETER
- DBA_DATA_FILES
- V_\$SYSTEM_PARAMETER
- V_\$DATAFILE
- V_\$SESS_IO
- V_\$SESSION
- DBA_FREE_SPACE
- V_\$SESSMETRIC (Oracle 10 only)
- DBA_TABLESPACES
- DBA_TEMP_FILES
- DBA_EXTENTS
- USER_EXTENTS
- V\$LOGFILE
- V\$LOG
- AUDIT_ACTIONS
- V\$CONTROLFILE

Any user with the SYSDBA role will have these privileges by default, so we recommend that you specify a user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a user with the SYSDBA role to connect, then you can create a separate user and explicitly grant permissions on those tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege, and grant permissions to create and drop global temporary tables, as the following example shows:

 **Note:** The following information is required to get Oracle data from a cluster setup.

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT ON V_$INSTANCE TO limited_user;
GRANT SELECT ON V_$PROCESS TO limited_user;
GRANT SELECT ON V_$DATABASE TO limited_user;
GRANT SELECT ON V_$PARAMETER TO limited_user;
GRANT SELECT ON DBA_DATA_FILES TO limited_user;
GRANT SELECT ON V_$SYSTEM_PARAMETER TO limited_user;
GRANT SELECT ON V_$DATAFILE TO limited_user;
GRANT SELECT ON V_$SESS_IO TO limited_user;
GRANT SELECT ON V_$SESSION TO limited_user;
GRANT SELECT ON DBA_FREE_SPACE TO limited_user;
GRANT SELECT ON DBA_TABLESPACES TO limited_user;
GRANT SELECT ON DBA_EXTENTS TO limited_user;
GRANT SELECT ON USER_EXTENTS TO limited_user;
GRANT SELECT ON DBA_TEMP_FILES TO limited_user;
GRANT SELECT ON V_$LOGFILE TO limited_user;
GRANT SELECT ON V_$LOG TO limited_user;
GRANT SELECT ON AUDIT_ACTIONS TO limited_user;
GRANT SELECT ON V_$CONTROLFILE TO limited_user;
exit;
```

Or

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT_CATALOG_ROLE TO limited_user;
exit
```

On Oracle database 12c RAC installation, to connect to the container database (CDB), you can use a common user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a user with the SYSDBA role to connect, then you can create a separate user. This must be prefixed with "c###" or "C###" and explicitly grant permissions on those tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege, as in the above example.

To connect to a pluggable database (PDB), you can use a common user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a common user with the SYSDBA role to connect, then you can create a PDB specific local user and explicitly grant permissions on those PDB tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege of the PDB.

The GRANT CREATE ANY TABLE command allows this user to create and drop global temporary tables. Global temporary tables must be created and dropped during some DPA Agent requests. DPA does not create or drop any other tables. To prevent the limited_user from inserting records into any table, you can execute the following SQL statement for increased security:

```
ALTER USER limited_user QUOTA 0M ON <TABLESPACE>; where <TABLESPACE> shall be replaced by the name of the tablespace for limited_user .
```

Before starting the Discovery Wizard for monitoring RMAN

To monitor an RMAN database for data protection data, the agent must connect to the database as an Oracle user.

Before you begin

Ensure that you have the following information connection parameters from the Oracle DBA or the RMAN catalog or system catalog queries:

- Oracle SID for RMAN Catalog
- Oracle TNS port being used for RMAN Catalog
- Oracle RMAN username/password with required privileges. These are SELECT only privileges or SELECT_CATALOG_ROLE privileges. In the case of multiple RMAN catalogs on one Oracle Server, you must have a username/password into each schema. Best practice is to use the same username/password across all RMAN catalogs/schemas.
- RMAN schema owner name, and if there are multiple RMAN catalogs on one Oracle Server, every RMAN schema owner name

About this task

To gather data successfully for Oracle RMAN Job Monitor Recovery Catalog, this user must be able to perform selects on the following tables and views:

- V_\$RMAN_CONFIGURATION
- RC_BACKUP_SET
- V\$PROXY_DATAFILE
- RC_RMAN_BACKUP_JOB_DETAILS
- RC_BACKUP_DATAFILE
- RC_BACKUP_PIECE
- RC_DATAFILE
- RC_DATABASE
- RC_BACKUP_CONTROLFILE
- RC_BACKUP_CONTROLFILE_DETAILS
- RC_BACKUP_DATAFILE_DETAILS
- RC_RMAN_STATUS
- RC_BACKUP_ARCHIVELOG_DETAILS
- RC_BACKUP_REDOLOG
- RCVER
- PRODUCT_COMPONENT_VERSION

To gather data successfully for Oracle Job Monitor Control File, this user must be able to perform selects on the following tables and views:

- V_\$RMAN_CONFIGURATION
- V_\$RMAN_STATUS
- V_\$BACKUP_DATAFILE
- V_\$BACKUP_PIECE
- V\$BACKUP_SET
- V\$PROXY_DATAFILE
- V\$RMAN_BACKUP_JOB_DETAILS
- V\$DATABASE
- V\$DATAFILE
- V\$BACKUP_DATAFILE_DETAILS
- V\$BACKUP_ARCHIVELOG_DETAILS
- V\$BACKUP_REDOLOG

- RCVER
- PRODUCT_COMPONENT_VERSION

Any user with the SYSDBA role will have these privileges by default, so we recommend that you specify a user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a user with the SYSDBA role to connect, then you can create a separate user and explicitly grant permissions on those tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege, as the following example shows:

Note: The following information is required to get Oracle data from a cluster setup.

For Oracle RMAN Job Monitor Recovery Catalog :

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT ON V_$RMAN_CONFIGURATION TO limited_user;
GRANT SELECT ON RC_BACKUP_SET TO limited_user;
GRANT SELECT ON V_$PROXY_DATAFILE TO limited_user;
GRANT SELECT ON RC_RMAN_BACKUP_JOB_DETAILS TO limited_user;
GRANT SELECT ON RC_BACKUP_DATAFILE TO limited_user;
GRANT SELECT ON RC_BACKUP_PIECE TO limited_user;
GRANT SELECT ON RC_DATAFILE TO limited_user;
GRANT SELECT ON RC_DATABASE TO limited_user;
GRANT SELECT ON RC_BACKUP_CONTROLFILE TO limited_user;
GRANT SELECT ON RC_BACKUP_CONTROLFILE_DETAILS TO limited_user;
GRANT SELECT ON RC_BACKUP_DATAFILE_DETAILS TO limited_user;
GRANT SELECT ON RC_RMAN_STATUS TO limited_user;
GRANT SELECT ON RC_BACKUP_ARCHIVELOG_DETAILS TO limited_user;
GRANT SELECT ON RC_BACKUP_REDOLOG TO limited_user;
exit;
```

Or

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT_CATALOG_ROLE TO limited_user;
exit
```

By default, a virtual catalog user has no access to the base recovery catalog. The following privileges should be granted for him to get access to metadata:

```
GRANT RECOVERY_CATALOG_OWNER to limited_user;
GRANT CATALOG for DATABASE db to limited_user;
```

For Oracle Job Monitor Control File:

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT ON V_$RMAN_CONFIGURATION TO limited_user;
GRANT SELECT ON V_$BACKUP_DATAFILE TO limited_user;
GRANT SELECT ON V_$BACKUP_PIECE TO limited_user;
GRANT SELECT ON V_$RMAN_STATUS TO limited_user;
GRANT SELECT ON V_$BACKUP_SET TO limited_user;
GRANT SELECT ON V_$PROXY_DATAFILE TO limited_user;
GRANT SELECT ON V_$RMAN_BACKUP_JOB_DETAILS TO limited_user;
```

```
GRANT SELECT ON V_$DATABASE TO limited_user;
GRANT SELECT ON V_$BACKUP_DATAFILE_DETAILS TO limited_user;
GRANT SELECT ON V_$DATAFILE TO limited_user;
GRANT SELECT ON V_$BACKUP_ARCHIVELOG_DETAILS TO limited_user;
GRANT SELECT ON V_$BACKUP_REDOLOG TO limited_user;
GRANT SELECT ON V_$PROXY_DATAFILE TO limited_user;
GRANT SELECT ON V_$RMAN_BACKUP_JOB_DETAILS TO limited_user;
exit;
```

Or

```
CREATE USER limited_user IDENTIFIED BY password;
GRANT CREATE SESSION TO limited_user;
GRANT CREATE ANY TABLE TO limited_user;
GRANT SELECT_CATALOG_ROLE TO limited_user;
exit
```

On Oracle database 12c RAC installation, to connect to the container database(CDB), you can use a common user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a user with the SYSDBA role to connect, then you can create a separate user. This must be prefixed with "c##" or "C##" and explicitly grant permissions on those tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege, as in the above example.

To connect to a pluggable database (PDB), you can use a common user that has the SYSDBA role when configuring the database for monitoring. If you do not want to use a common user with the SYSDBA role to connect, then you can create a PDB specific local user and explicitly grant permissions on those PDB tables or grant "create session" followed by SELECT_CATALOG_ROLE privilege of the PDB.

The GRANT CREATE ANY TABLE command allows this user to create and drop global temporary tables. Global temporary tables must be created and dropped during some DPA Agent requests. DPA does not create or drop any other tables. To prevent the limited_user from inserting records into any table, you can execute the following SQL statement for increased security:

```
ALTER USER limited_user QUOTA 0M ON <TABLESPACE>; where <TABLESPACE> shall be replaced by the name of the tablespace for limited_user .
```

If you would like DPA to collect job data from 14 days before, and for the reports show data straight away for Oracle RMAN, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Monitoring of PostgreSQL

A PostgreSQL database can be monitored from an agent running on the same host as the PostgreSQL database or from an agent running on a different host, such as the DPA server.

Before starting the Discovery Wizard for monitoring PostgreSQL

To monitor a PostgreSQL database, the agent must connect to the database as a PostgreSQL super user. A super user has the correct privileges by default. We recommend that you specify a super user when configuring the database for monitoring.

About this task

To create a super user, the PostgreSQL administrator must be a super user, and create the account as in the following example:

```
CREATE ROLE xxxxx WITH login superuser password yyyyyy ;
```

where xxxxx is the new username and yyyyyy the new user's password.

The following parameters will not be populated in the database server parameters table unless you are connecting to the database as a super user:

- config_file
- data_directory
- dynamic_library_path
- external_pid_file
- hba_file
- ident_file
- krb_server_keyfile
- log_directory
- log_filename
- preload_libraries
- unix_socket_directory

The following items are also unavailable unless you are connecting as a super user:

- In the datafile configuration table, the full path to the datafiles cannot be shown, as the path of the file is found in the data_directory parameter. The string (postgres data directory) is shown instead.
- In the connection status table, the f_command and f_status fields will not be populated with the right information. These fields will be set to <insufficient privileges>.

Connecting to the database as a super user populates all fields.

Monitoring of SAP HANA

A SAP HANA database can be monitored from an agent running on the same host as the SAP HANA server, or from an agent running on a different host, such as the DPA server. The DPA Agent must be run on Windows or Linux.

Before starting the Discovery Wizard for monitoring SAP HANA

For DPA Agent to collect data from SAP HANA database, you must copy the SAP HANA client .jar file to the DPA plugins directory.

Procedure

1. Create a directory called *plugins* under <DPA_install_dir>\agent\.
2. Copy the SAP HANA client jar file *ngdbc.jar* to the *plugins* folder under ..\EMC\dpa\agent\.

For the custom location or path add following tag: <PLUGINS_DIR>path </PLUGINS_DIR> in dpaagent_config.xml located under <DPA_install_dir>\agent\etc

where *path* is the path of the directory created in step 1.

For example <PLUGINS_DIR>c:\program files\emc\dpa\agent\plugins</PLUGINS_DIR>

3. If you want DPA to collect job data from 14 days before, and for the reports show data straight away for SAP HANA, enable the default historical data from the Job Monitor request. In the DPA web console, go to **Inventory > Object Library > [select object] > Data Collection**.

Permissions for discovering data for SAP HANA

To gather data on SAP HANA, the database user must have certain privileges that allow the user to run SELECT queries.

The credentials are used by the DPA Agent to get access to the following tables:

- M_BACKUP_CATALOG view
- M_BACKUP_CATALOG_FILES view

Usually, the privileges granted to the PUBLIC role are sufficient to read that data. For more information, refer to vendor information on privileges required for running SELECT queries.

Monitoring of applications using cloud-based solutions

This section describes how to monitor applications using DPA that is deployed on cloud-based solutions.

Monitoring applications on Amazon Web Services

DPA supports deployment of DPA within Amazon Web Services as well as on premises for discovery and monitoring of supported backup and monitoring applications on premises or within Amazon Web Services. The *Data Protection Advisor Software Compatibility Guide* provides information on supported versions of backup and monitoring applications.

Before you begin

- Ensure that you configure the DPA Data Collection Agent in the same Amazon Web Services space as the objects that you plan to monitor by using Amazon Web Services.
- If you are configuring DPA to monitor applications deployed on cloud-based solutions using a VPN, ensure that ports and protocols are available across the VPN. If you are using nonstandard ports work with your Cloud services provider or with Amazon Web Services to open nonstandard ports. [DPA port settings](#) provides information on standard DPAA ports.

Procedure

1. Deploy DPA in your Amazon Web Services environment.

[Installing DPA](#) on page 29 provides information on DPA installation. Refer to Amazon Web Services documentation for specific product requirements.

2. Discover the supported application on the DPA instance within Amazon Web Services.

The sections in this chapter provide information. For example, to discover and monitor NetWorker, [Monitoring of NetWorker](#) on page 158 provides information.

Monitoring applications on Microsoft Azure

DPA supports deployment of DPA within Azure for discovery and monitoring of supported backup and monitoring applications. The *Data Protection Advisor Software Compatibility Guide* provides information on supported versions of backup and monitoring applications.

Procedure

1. Deploy DPA in your Azure environment.

[Installing DPA](#) on page 29 provides information on DPA installation. Refer to Azure documentation for specific product requirements.

2. Discover the supported application on the DPA instance within Azure.

The sections in this chapter provide information. For example, to discover and monitor NetWorker, [Monitoring of NetWorker](#) on page 158 provides information.

Monitoring of hosts

This section describes monitoring of hosts.

DPA provides two options during host discovery:

- Host System monitoring, to monitor configuration, performance, and status of the operating system.
- Replication monitoring, to perform Storage Replication Analysis.

Monitoring operating systems

Use the Discovery Wizard Host System to monitor configuration, performance, and status of the operating system. There are several DPA modules that gather different types of information, as described in the following table.

Table 36 System monitoring modules

Module	Description
Host	Gathers basic information about the operating system type.
Disk	Gathers configuration, status, and performance information on the disks attached to the host.
Fibre Channel HBA	Gathers configuration, status, and performance information on Fibre Channel HBAs configured on the computer.
File system	Gathers configuration, status, and performance information on the file systems mounted to the host.
Memory	Gathers configuration, status, and performance information on memory in the host.
NetInt	Gathers configuration, status, and performance information on network interface cards in the host.
Process	Gathers information on any processes running on the host.
Processor	Gathers configuration, status, and performance information on all CPUs on the host.

Gathering of data from UNIX operating systems

To perform system monitoring on UNIX computers, install an agent on the host that is to be monitored. It is not possible to gather system information remotely from UNIX computers.

Discovering agent hosts for UNIX for gathering data

UNIX hosts are discovered using SSH or telnet/ftp with root access.

About this task

If security requirements do not allow for root credentials to be supplied to DPA, sudo is a workaround that can temporarily elevate a user's credentials to root for specific commands configured in the sudoers file.

Modifying sudoers file for DPA storage discovery

A user can log in to a UNIX host as a non-root user, and use sudo to run SCSI commands successfully to discover storage related information for the host. The following is an example of what needs to be added to the sudoers file

About this task

```
# sudoers file.
#
# This file MUST be edited with the 'visudo' command as root.
#
# See the sudoers man page for the details on how to write a sudoers file.
#
# Host alias specification
# User alias specification
# Cmnd alias specification
# Defaults specification
# User privilege specification
root    ALL=(ALL) ALL
# Uncomment to allow people in group wheel to run all commands
# %wheel    ALL=(ALL)    ALL
# Same thing without a password
# %wheel    ALL=(ALL)    NOPASSWD: ALL
# Samples
# %users    ALL=/sbin/mount /cdrom,/sbin/umount /cdrom
# %users    localhost=/sbin/shutdown -h now
user alias ALL = (ALL) PASSWD: /var/tmp/IllumAgent/apollloreagent
# Defaults specification
# User privilege specification
root    ALL=(ALL) ALL
CMGU    ALL=NOPASSWD:CMGEMC
# Uncomment to allow people in group wheel to run all commands
# %wheel    ALL=(ALL) ALL
# Same thing without a password
# %wheel    ALL=(ALL) NOPASSWD: ALL
# Samples
# %users    ALL=/sbin/mount /cdrom,/sbin/umount /cdrom
# %users    localhost=/sbin/shutdown -h now
```

```
#cmguser ALL=(ALL) NOPASSWD: ALL
```

Gathering of data from Windows operating systems

To gather performance data from a Windows host, you must install Windows Management Infrastructure (WMI) on the Windows host you are monitoring.

It is possible to gather all system monitoring information remotely from Windows computers, with the exception of Fibre Channel HBA information. To gather Fibre Channel HBA information, the agent must be installed on the computer. [Monitoring a Windows host remotely](#) on page 183 provides more details on the steps required to monitor a Windows host remotely.

To set up system monitoring for a system on which an agent is installed, assign the system monitoring requests to the host or group to monitor.

Discovering agent hosts for Windows for gathering data

If application discovery is being performed without an agent, Windows host discovery uses Remote Procedure Calls (RPC) for replication analysis and WMI for System information.

About this task

Checking RPC Communication Procedure

1. Open the Run dialog box from the Windows **Start** menu.
2. Type:


```
net use \\<servername>\admin$ /user:<username>
```
3. Click **Enter**. Type the password.
4. A successful connection should return the following message: The command completed successfully.
5. Delete the network map. Type:


```
net use \\<servername>\admin$ /delete
```

Checking WMI Communication Procedure

1. Open the Run dialog box from the Windows **Start** menu.
2. Type WBEMtest and click **Connect** in the Windows Management Instrumentation Tester dialog box.
3. In the **Connect** field, type \\<servername>\root\cimv2.
4. In the **Credentials** fields, type the username and password used to connect to the application host you are monitoring.
5. Click **Connect** to return to the Windows Management Instrumentation Tester dialog box. Click **Query**.
6. In the **Enter Query** field, type:


```
select * from win32_processor
```
7. Click **Apply**.

If WMI can connect, data from the application host is displayed.

Monitoring a Windows host remotely

All system information can be gathered remotely from a Windows computer with the exception of Fibre Channel HBA information. To monitor a Windows computer remotely, you must install an agent on another Windows computer. You cannot remotely monitor a Windows computer from an agent running on a UNIX computer.

About this task

To monitor a Windows host from another Windows computer, the DPA agent service must run as administrator on the computer performing the monitoring. [Modifying the login parameters of the agent service](#) on page 183 provides more information.

Modifying the login parameters of the agent service

Checking if this is required. To modify the login parameters of the agent service:

Procedure

1. Launch the Windows Services control manager: **Start > Settings > Control Panel > Administrative Tools > Services**).

2. Select the DPA Agent service.
3. Right-click and select **Properties** from the menu.
4. Select the **Log On** tab in the **Properties** dialog box.
5. Select **This Account**.
6. Type the username and password of the administrator that the service to run as.
7. Click **OK** and restart the service.

Monitoring activity on a remote computer

Procedure

1. Create a host object for the computer to monitor in the web console. The name of the object is the hostname of the remote host. The hostname must be resolvable from the computer on which the agent that will be monitoring the object is running.
2. Assign requests to that object to specify the data to gather.
3. Mark each request as a proxy request and complete the details.
4. To complete the proxy details, type the name of the host for the agent in the **Proxy Host** field.
5. Create a Windows credential for the Administrator account on the computer being monitored. This account can be the name of a Local Administrator or that of a Domain Administrator.
6. Notify the agent that will monitor the server of the changes by reloading the agent.

Monitoring of a host for system data

Monitor an application host for system data from an agent running on the host or another host in the environment.

Before starting the Discovery Wizard for monitoring a host for system data

System data can only be gathered from UNIX systems by an agent local to the UNIX host.

About this task

Configuring for Replication Analysis


Use the Discovery Wizard to perform Storage Replication Analysis.

Before you begin

- For ProtectPoint backup and recovery configuration., ensure that you have application discovery ability or that you have set the Replication Monitoring flag.
- For ProtectPoint backup and configuration, ensure that you synchronize the time, within a maximum of 1-minute difference, of the host that is protected by ProtectPoint with the Solutions Enabler host that manages the storage array that the application is mapped to.
- Ensure that communication between the monitored host and the recoverability process is enabled:
 - For monitoring Windows servers remotely, you must enable RPC services and ensure that they are accessible to the recoverability agent.
 - For UNIX/Linux remote application monitoring, you must enable SSHD and ensure that it is accessible to the recoverability agent.
 - For UNIX/Linux remote application monitoring, you must enable FTP/Telnet services and ensure that they are accessible to the recoverability agent.

Monitoring of Microsoft Exchange Server

To discover Microsoft Exchange Server, you must discover the host that Microsoft Exchange Server runs on. An Exchange Server can be monitored for recoverability from an agent installed on the same host as the Exchange Server or an agent installed remotely.

 **Note:** Microsoft Exchange can only be monitored for replication analysis, and for system information from the Exchange server host.

Before starting the Discovery Wizard for monitoring Microsoft Exchange Server

The account used to connect DPA to the Exchange server must be a domain user with Exchange read-only administrator rights and local administrator rights. DPA does not support replication analysis for two Exchange information stores on a cluster. To connect to the exchange application you must have Exchange read-only administrator rights. To retrieve the disks information from Windows you must be an operating system user with local administrator rights.

About this task

Monitoring Oracle for Replication analysis

To monitor an Oracle database for replication analysis, the agent must connect to the database as an Oracle user able to perform selects on the following tables and views:

- DBA_DATA_FILES
- DBA_TEMP_FILES
- DBA_TABLESPACES
- V_\$DATAFILE
- V_\$LOGFILE
- V_\$CONTROLFILE
- V_\$LOG_HISTORY
- V_\$ARCHIVED_LOG
- V_\$INSTANCE
- V_\$DATABASE
- V_\$PARAMETER
- DICT
- DBA_TAB_COLUMNS

When monitoring Oracle on a Windows platform, the operating system user specified in the Credential must belong to the group ORA_DBA. On UNIX, if you use UNIX authentication, you need not define the credentials in the database.

Updating Oracle statistics

To gather accurate figures on the number of rows and size of tables and indexes, it is important that Oracle statistics are updated on a regular basis. The Oracle documentation contains more details on how to set up a job to update Oracle statistics.

About this task

One method to update Oracle statistics on a Schema is to run the following command:

```
exec dbms_stats.gather_schema_stats(ownname => '***SCHEMANAME***',
estimate_percent => 5, cascade => true, options => 'GATHER');
```

Monitoring of RecoverPoint

You must monitor RecoverPoint from an agent installed remotely, the DPA server, for example.

When discovering RecoverPoint, DPA supports discovering only one management IP. Additionally, DPA supports monitoring only the management IP and not the RPA IP. Ensure that you monitor the Management IP and not the RPA IP.

Monitoring of primary storage

This section describes how to monitor primary storage.

DPA breaks primary storage out to the following categories:

- File Servers
- Storage Arrays for Replication Analysis
- Disk Management Servers

Monitoring of file servers

This section describes how to monitor file servers.

Monitoring of EMC File Storage

EMC File Storage must be monitored from an agent running on a remote computer, for example, the DPA server.

 **Note:** EMC File Storage is interchangeably referred to as Celerra File Storage.

Configuration of storage arrays for replication analysis

DPA monitors VNX Block, CLARiiON, Symmetrix, and VPLEX storage arrays. If these storage arrays are replicated with RecoverPoint, additional configuration is required to enable complete replication analysis.

Port for EMC VPLEX arrays

DPA connects to the VPLEX on TCP port 443.

Discovery of VPLEX arrays

VPLEX storage arrays can be monitored from the DPA Server or remotely from any host that has DPA agent installed.

DPA discovers all of the storage arrays that are being managed and creates objects in the object library inventory.

Port for VNX Block /CLARiiON arrays

DPA connects to the VNX Block/CLARiiON on TCP port 443. However, if the VNX Block/CLARiiON is configured to use port 2163, use port 2163.

Discovery of VNXBlock/CLARiiON arrays

VNXBlock/CLARiiON storage arrays must be monitored remotely from a proxy server or, as a last resort, from an agent that runs on a different host, such as the DPA server. This is also known as the SE host or Connector.

The SE host can be used for discovery through a DPA Agent installed on it or through an agentless mechanism that requires a privileged user's credentials.

DPA discovers all of the storage arrays that are being managed and creates objects in the object library inventory.

You will need to supply the name of the host on which EMC Solutions Enabler is installed.

Discovery of Symmetrix arrays

Symmetrix storage arrays must be monitored remotely from an agent running on a different host, such as the DPA server.

To configure multiple hosts and multiple storage arrays, use the Discovery Wizard. DPA discovers all of the storage arrays that are being managed and creates objects in the object library inventory.

You must supply the name of the host on which Solutions Enabler is installed.

In order for Solution Enabler to see Device Groups which are stored locally by default on Solution Enabler, you must open `Global Name Services` in Solution Enabler options file, as follows:

1. Open option file under `./emc/API/symapi/config/`
2. Find the line

```
#SYMAPI_USE_GNS = ENABLE
```

and unremark it so that it is: `SYMAPI_USE_GNS = ENABLE`

3. Save the file.
4. Verify that the GNS service is running by running the `stordaemon list` command.
5. Run the `symcfg disco` command.

Performing hostless discovery on Symmetrix and VNX/CLARiiON

Host discovery with replication monitoring requires either the installation of a local agent on the host or the deployment of a remote agent with credentials for host access. Either method might be prevented by customer security policies.

About this task

To use the agentless option, you must provide the Solutions Enabler host credentials. The prerequisites for hostless discovery are the same as those described in [Discovery of Symmetrix arrays](#) on page 186.

Configuring storage arrays that use RecoverPoint to gather replication data

If your VNX/CLARiiON or Symmetrix storage arrays are replicated with EMC RecoverPoint, DPA provides replication analysis for RecoverPoint replication operations.

About this task

To perform replication analysis for RecoverPoint, you have to configure the VNX/CLARiiON or Symmetrix storage arrays and the RecoverPoint host in DPA in the correct order.

Procedure

1. Use the Discovery Wizard to create the host object for the Solutions Enabler host that is connected to the storage array replicated with RecoverPoint.
2. Discover the arrays attached to the host.
3. Configure the Symmetrix or VNX/CLARiiON arrays using the Discovery Wizard.
4. Import replication policy data from the storage arrays.
5. Configure the EMC RecoverPoint appliances' data monitoring, as described in [Monitoring of RecoverPoint](#) on page 185.
6. Ensure that the RecoverPoint Configuration request has been assigned to the RecoverPoint appliance object that handles replication for the storage array. Run this request.
7. After the RecoverPoint Configuration request is run and sufficient time has passed, DPA should have begun gathering replication analysis data for RecoverPoint. Reports can be run from the storage array objects and the **Replication Analysis** area will show the mapping of storage and recovery points.

Before starting the Discovery Wizard for Monitoring EMC File Storage

The EMC File Storage module gathers information from EMC File Storage through an XML API and directly from the EMC File Storage Control Station. You must create an administrator with specific privileges on the EMC File Storage:

Procedure

1. Log in to the EMC File Storage Manager web browser interface as an administrator.
You can also use the command line interface to create a DPA administrator.
2. Navigate to **Security > Administrators**.
3. Create a new administrator, with a username of DPA, for example.
4. Select **Local Only Account** and type and confirm a password for the administrator.
5. Select a **Primary Group** of at least opadmin level of privilege. DPA does not need greater privileges than those assigned by opadmin.
6. Enable the following client access options:
 - XML API v2 allowed
 - Control Station shell allowed
7. Click **OK**.

Results

The DPA Credential used to connect to the EMC File Storage must contain the username and password of the EMC File Storage administrator you created.

Monitoring of disk management servers

This section describes how to monitor disk management servers.

Monitoring of HP Command View

Monitor a HP EVA Disk Array through HP Command View from an agent running on the Command View host, or remotely from an agent running on a different host, such as the DPA server.

The username and password used to gather data must match a valid username and password defined in the CommandView CIM server. You can configure this from the CommandView management interface.

DPA gathers data from HP Command View using SMI-S on the default secure port of 5989.

Monitoring of protection storage

This section describes how to monitor protection storage.

About this task

Monitoring of Data Domain

DPA monitors Data Domain backup appliances. For DDOS 4.8, only Tape Drive and Tape Library Status and Configuration information is returned. You must enable the Data Domain analysis request on the Data Domain systems on which you wish to gather the data.

Before starting the Discovery Wizard for monitoring Data Domain

You must enable SNMP on port 161 and SSH on port 22 on the Data Domain backup appliance. You also need to set the SNMP community string. You can do this from the command line.

Before you begin

- Ensure that you have user role rights to run SSH requests on the Data Domain system.
- Ensure that you have user admin privileges to run PCR (Physical Capacity Reporting) for monitoring Data Domain OS 5.7 or higher.

Procedure

1. Log on to the Data Domain appliance console using the sysadmin account.
2. Type the following command to check the existing configuration:

```
snmp show ro-communities
```

```
snmp add ro-community <string> hosts <host IP address>
```

where *<string>* is the selected community string (for example, public) and *<host IP address>* is the IP address of the DPA Agent that you are using to monitor the Data Domain. You will have to disable and re-enable SNMP for the new string to take effect.

```
snmp disable
snmp enable
```

If you are not using a community string of public, you must change the community string used in the Data Domain Credential.

You can also set SNMP settings through the **System Settings** tab of the Data Domain Enterprise Manager interface.

3. Edit the DPA Data Domain SSH Credential to specify an SSH username and password configured on the Data Domain device. Go to **Admin > System > Manage Credentials** in the DPA web console.

This is required:

- to ensure configuration of SSH PCR data collection when monitoring Data Domain OS 5.7 or higher.
 - When the request runs, it gathers statistics for the command polling period time, and then it creates the physical capacity measurement schedule on the Data Domain. The Data Domain then gathers the statistics. The statistics are gathered, collected, and sent to the DPA server when the subsequent request runs. As a result the first time the request runs no data is collected on the reports; data is collected and reported only at the second run of the request. [DPA postinstallation steps](#) on page 60 provides more information.
 - The command polling period is rounded up to a full day times. The command polling period value will be set to twice the polling period value with the proviso that the command polling period will be at least 2 days' time. For example, if the polling period is set to 24hours or less, DPA gathers statistics for 2 days. If the polling period is set to 3 days, the DPA gather statistics for 6 days.
 - to get LUN information from Data Domain such as devices, device-groups, pools, static-images, and access groups for ProtectPoint SnapVX Backup and Recovery. [Configuring DPA for ProtectPoint SnapVX Backup and Recovery](#) on page 190 provides information.
- among other information.

Configuring DPA for ProtectPoint SnapVX Backup and Recovery

You must configure DPA to associate the information collected on the host in the DPA environment to the information collected on the VMAX3 in the DPA environment, and in turn associate that information to the information collected on the Data Domain in the DPA environment.

Before you begin

- Ensure that you synchronize the time, within a maximum of 1-minute difference, of the host that is protected by ProtectPoint with the Solutions Enabler host that manages the storage array that the application is mapped to.
- The *Data Protection Advisor Software Compatibility Guide* provides information on supported versions of and OS requirements for:
 - ProtectPoint
 - Solutions Enabler
 - VMAX3
 - Data Domain

Procedure

1. Configure the host for replication analysis.

provides information. Ensure that you have application discovery ability or that you have set the Replication Monitoring flag. This is required for ProtectPoint backup and recovery configuration.

2. Discover the VMAX3 and SE host.

[Discovery of Symmetrix arrays](#) on page 186 provides information.

3. Discover the Data Domain host.

[Monitoring of Data Domain](#) on page 188 provides information. Ensure that you provide SSH credentials at the Data Domain discovery wizard. This is required to get LUN information from Data Domain such as devices, device-groups, pools, static-images, and access groups.

After you finish

If desired, add new protection rules to your protection policy so Linked, StaticImage, and SnapVX Missing Recovery Point alerts are generated.

Monitoring of StorageTek ACSLS Manager

StorageTek ACSLS Manager cannot be monitored remotely. A DPA agent must be installed on the ACSLS AIX or ACSLS Solaris host.

Before starting the Discovery Wizard for Monitoring StorageTek ACSLS Manager

The agent must be installed and running on the StorageTek ACSLS Manager server that you want to monitor.

About this task

After installing the agent, verify that the ACS_HOME value in the DPA.config file matches the location in which ACSLS is installed. Verify that the ACSDBDIR value in the DPA.config file matches the path to the ACSLS DB folder (the default is export/home/ACSDB 1.0).

Monitoring of tape libraries

DPA can gather information about tape libraries and the drives within those tape libraries. When you specify a hostname, ensure that the name of the tape library is resolvable from the host that is monitoring the tape library.

Before starting the Discovery Wizard for monitoring tape libraries

The tape library credentials must contain the read-only community string for the tape library in the **Password** field of the **Credential Properties** dialog box. Unless the community string was modified on the tape library, set the community string to **Public**.

About this task

Select **Admin > System > Manage Credentials** to modify the tape library credentials that are created after using the Discovery Wizard to create a tape library object.

Monitoring the IBM System Storage TS 3500 tape library

Use the Tape Library Specialist web interface to enable Simple Network Management Protocol (SNMP) requests for the IBM System Storage TS 3500 Tape Library. To enable SNMP requests:

About this task

Procedure

1. Type the Ethernet IP address on the URL line of the browser.
2. Select **Manage Access > SNMP Settings**. In the **SNMP Trap Setting** field, view the current setting then click to enable SNMP requests.
3. Ensure that the **SNMP Requests Setting** field is set to **Enabled**.

Monitoring the IBM TotalStorage 3583 tape library

Configure the Remote Management Unit (RMU) to enable SNMP for the IBM TotalStorage 3583 Tape Library. To enable SNMP:

About this task

Procedure

1. In the RMU, click **Configuration**.
2. In the SNMP Configuration region, perform the following:
 - To enable the feature, select **ON** in the **SNMP Enabled** field.
 - To enable or disable SNMP alerts, select **ON** or **OFF** in the **Alerts Enabled** field.
 - In the **Manager** field, type the SNMP server address.
 - In the **Public Name** field, type the name of the read-only SNMP community.
 - In the **Private Name** field, type the name of the read/write SNMP community.
3. Click **Submit** and review the changes.
4. Type the password and click **Confirm**. Redirect the browser if required.
5. Click **Done** to reboot.

Monitoring the IBM TotalStorage 3584 tape library

To enable SNMP from the web interface of the IBM TotalStorage 3584 tape library:

About this task

Procedure

1. From the Welcome screen of the Tape Library Specialist Web Interface, select **Manage Access > SNMP Settings**.
2. In the **SNMP Trap Setting** field, view the current setting, and select the button to enable or disable SNMP requests.

Alternately, to enable SNMP requests from the operator panel:

3. From the Activity screen of the tape library operator panel, select **MENU > Settings > Network > SNMP > Enable/Disable SNMP Requests > ENTER**.

The screen displays the current status of SNMP requests.

4. Press **UP** or **DOWN** to specify ENABLED or DISABLED for SNMP messaging, and click **ENTER**.

To accept the new setting and return to the previous screen, click **BACK**.

The Enable/Disable SNMP Requests screen redisplay the new setting.

Monitoring the Oracle SL24 Tape Autoloader and SL48 tape library

Configure the Remote Management Interface (RMI) to enable SNMP for the Oracle StorageTek SL24 Tape Autoloader or SL48 Tape Library. To enable SNMP:

About this task

Procedure

1. In the RMI, navigate to **Configuration > Network**.
2. Ensure the **SNMP Enabled** checkbox is enabled.
3. The **Community Name** string must be contained in the credentials used to connect to this Tape Library in DPA.
4. Click **Submit** and review the changes.

Monitoring the HP StorageWorks tape library

Configure the NeoCenter utility to enable SNMP for the tape library. To enable SNMP:

About this task

Procedure

1. Launch the NeoCenter utility from the host.
2. Select **Configure** from the Main screen menu. The **Configure** dialog box appears.
3. Select the **SNMP Traps** tab.
4. In one of the available **Trap Address** fields, type the IP address of the DPA server.

Monitoring of switches and I/O devices

This section describes how to monitor switches and I/O devices.

Monitoring of Fibre Channel switches

DPA gathers information about ports on Fibre Channel switches, including configuration, connectivity status, and throughput.

When you specify a hostname, ensure that the name of the switch is resolvable on the agent's host.

Before starting the Discovery Wizard for monitoring Fibre Channel switches

To ensure that Brocade switches return all data, verify that the Fibre Channel Alliance MIB is loaded and enabled on the switch. This MIB might not be installed on the switch by default. To enable FA-MIB support on Brocade switches, log in as an administrator and run the `snmpmibcapset` command. Change the FA-MIB parameter to Yes. Click Enter to accept the default for the other settings.

About this task

For example:

```
telnet <switch>
> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB SW-MIB FA-MIB SW-TRAP FA-TRAP
FA-MIB (yes, y, no, n): [yes]
SW-TRAP (yes, y, no, n): [enter]
FA-TRAP (yes, y, no, n): [enter]
SW-EXTTRAP (yes, y, no, n): [enter]
>
```

Monitoring of IP switches

When you are specifying a hostname, ensure the name of the switch is resolvable on the agent's host.

Before starting the Discovery Wizard for monitoring IP switches

The IP Switch Credentials must contain the SNMP community string for the IP switch in the **Password** field of the **Credential Properties** dialog box. Unless the community string was modified on the IP switch, set the community string to public.

About this task

Select **Admin > System > Manage Credentials** to modify the IP Switch Credentials that are created after you have used the Discovery wizard to create an IP switch object.

Monitoring of Xsigo I/O Director

When you are specifying a hostname for the Xsigo I/O Director, ensure the hostname or IP address of the Director is resolvable on the agent's host.

Before starting the Discovery Wizard for monitoring Xsigo I/O Director

The Xsigo Director SNMP credentials must contain the SNMP community string for the Director in the **Password** field of the Credential. Unless the community string was modified on the Director, set the community string to public.

About this task

Select **Admin > System > Manage Credentials** to modify the default Xsigo Director SNMP Credentials if required, or to create a new credential.

Virtualization management

This section describes how to monitor a virtualized environment.

Monitoring of VMware environment

Monitor your VMware environment from an agent running on the VirtualCenter Server or remotely from an agent running on a different host, such as the DPA server.

- The Discovery Wizard can be used to add a vCenter server to DPA. Go to **Admin > System > Discovery Wizard > Virtualization Management**.
- To add a vCenter server, you must provide the vCenter hostname and credentials for a vCenter user.
- You can select whether to monitor the vCenter host only or to also monitor the virtual machines connected to the vCenter host.
 - If you select to monitor virtual machines, DPA queries the vCenter Server and displays a list of virtual machines. The discovery process can take a while if there are a large number of virtual machines configured on the vCenter server.
 - For each virtual machine you can select whether you wish to discover the host in DPA. Discovering the host adds the host to the DPA inventory.
 - For each virtual machine selected for discovery, you can select whether to enable Host System Monitoring, which gathers configuration, performance and analysis data; and Replication Monitoring, which enables replication analysis.
 - For each virtual machine selected for Host System Monitoring, you can specify which DPA Agent should be used to monitor the virtual machine. You can change the DPA Agent for multiple machines simultaneously by using CNTRL-Click or SHIFT-Click to select multiple systems.
 - Windows virtual machines can have Host System Monitoring performed using a remote DPA Agent such as the DPA Agent installed on the DPA Server; or a local agent, such as DPA Agent installed on each Windows virtual machine.
 - UNIX/ Linux virtual machines must have a DPA Agent installed on the virtual machine for Host System Monitoring, on a local agent.
 - If you choose to do host monitoring for each VM, you must provide Windows credentials for each Windows Virtual Machine being monitored with a remote agent. The credentials can either be a local administrator or a domain administrator. You can change the credential for multiple machines simultaneously by using CNTRL-Click or SHIFT-Click to select multiple systems. You need not provide these credentials if you are monitoring the vCenter
 - Discovered virtual machines are displayed under the vCenter object in DPA and by default will also be added to Configuration / Servers / Application Servers group. You can change and add groups for the virtual machines to appear. Go to **Admin > System > Discovery Wizard > Destination Group**.
- The final screen of the vCenter Discovery Wizard displays a summary of options selected. If you click **Finish**, it adds the objects to DPA and enables monitoring options selected.

Monitoring of RecoverPoint for VMs

You must monitor RecoverPoint for VMs from an agent installed remotely; the DPA server, for example. The DPA Agent must be run on Windows or Linux.

When discovering RecoverPoint for VMs, DPA supports discovering only one management IP. Additionally, DPA supports monitoring only the management IP and not the RPA IP. Ensure that you monitor the Management IP and not the RPA IP.

Before starting the Discovery Wizard for monitoring RecoverPoint

DPA needs to be able to connect to the RecoverPoint environment Command Line Interface (CLI) through a secure SSH connection on port 22. DPA connects to the RecoverPoint appliance using the default CLI user admin, but any defined user with sufficient privileges to run a CLI command remotely using SSH is possible; the monitor account is sufficient.

About this task

However, DPA must not connect with the RecoverPoint user boxmgmt because user boxmgmt is reserved for starting the RecoverPoint installation manager automatically.

If you are running RecoverPoint 4.1 where the default user is "monitor," then you must create a new user because the default user specified in DPA no longer exists. If you do not create a new user after installing RecoverPoint 4.1, the request with RecoverPoint Credentials from DPA fails.

Monitoring of clusters

This section describes how to monitor clusters.

Monitoring of Microsoft Server Failover Cluster

To discover Microsoft Server Failover Cluster, you must install the agent on each machine which is in the cluster. The *Data Protection Advisor Software Compatibility Guide* provides information on supported versions.

You must discover Microsoft Server Failover Cluster by a remote agent within the DPA Discovery Wizard. This agent should be installed on one of machine from the cluster. DPA provides two discovery options:

- **Monitor Cluster and hosts which are included in cluster**—If you select this, DPA automatically selects Clustered Server the Cluster with Cluster Configuration and Cluster Status requests.
DPA assigns the Host Monitoring, Host Configuration, and Host status requests to all hosts which are included in cluster.
- **Monitor only Cluster**—If you select this, DPA automatically selects the Cluster with Cluster Configuration and Cluster status requests.

 **Note:** Hosts that are included in cluster will not have the assigned requests.

Monitoring of Veritas Cluster Server and Veritas Infoscale Availability

To discover Veritas Cluster Server and Veritas Infoscale Availability, you must install the agent on each machine which is in the cluster. The *Data Protection Advisor Software Compatibility Guide* provides information on supported versions.

You must discover Veritas Cluster Server and Veritas Infoscale Availability by a remote agent within the DPA Discovery Wizard. This agent should be installed on one of machine from the cluster. DPA provides two discovery options:

- **Monitor Cluster and hosts which are included in cluster**—If you select this, DPA automatically selects Clustered Server the Cluster with Cluster Configuration and Cluster Status requests.
DPA assigns the Host Monitoring, Host Configuration, and Host status requests to all hosts which are included in cluster.
- **Monitor only Cluster**—If you select this, DPA automatically selects the Cluster with Cluster Configuration and Cluster status requests.

 **Note:** Hosts that are included in cluster will not have the assigned requests.

Discovering a host or object manually

About this task

This procedure does not apply when discovering CLARiiON, Symmetrix, VNX, or VPLEX arrays.

Note: The steps that display vary based on the object that you are discovering.

Procedure

1. Perform one of the following:
 - Go to **Inventory > Object Search**.
 - Go to **Admin > Run Discovery Wizard**.
2. In **Objects to Discover**, select one of the following:
 - **Host** and then select Host.
 - **Primary Storage** and then select File Storage or VPLEX.
 - **Protection Storage** and then select Data Domain, Disk Library, NetApp NearStore, or Tape Library.
 - **Switch** and then select Fibre Channel switch, IP switch, or Xsigo switch.
3. Select the option to discover the host or object manually.
4. Identify the application host by hostname or IP address, alias, operating system, Credential, remote data collection agent, or ports. If you are discovering Primary Storage, Protection Storage, or Switches, then proceed to step 8.
5. Select **Host System Monitoring** or **Replication Monitoring** for each host that you want to discover. The Replication Monitoring option is available only if you have a storage capacity license. If you do not select an option during discovery, you can later Add Requests and its options.
6. Select whether a Local or Remote data collection agent will gather data for this application. If you selected **Host System Monitoring** and your host is running Linux, UNIX, or other non-Windows platforms select local data collection agent. For Remote data collection agents, select the host that has the agent installed.

Note: If you had specified a Data Collection Agent for RecoverPoint, RecoverPoint for VMs, or VPLEX in the **Viewing and editing Data Collection defaults** area, the agent is displayed here by default.

To add or edit an agent, specify the fields described in the following table.

Field	Description
Hostname	Name of the host with the data collection agent installed.
Display Name	Name of the host with the data collection agent installed that displays.
Operating System	Operating system of the host with the data collection agent installed.
Host System Monitoring	Select to monitor configuration, performance, and status for this host.
Replication Monitoring	Select to perform Replication Analysis for this host.

7. If you selected Host System Monitoring and Remote data collection agent or agentless, select or set the application host credential.

Note: If you had specified a credential for RecoverPoint, RecoverPoint for VMs or VPLEX in the **Viewing and editing Data Collection defaults** area, the credential is displayed here by default.

8. (optional) Test the connection to the object. If the test fails for host or credential errors, click **Back** to resolve and then retest.
9. (optional) Add the object to a group or to multiple groups. Press **Ctrl** or **Shift** and click to select multiple objects.
10. (optional) If you have defined custom attributes, select the attributes that you want to apply to the discovered objects. Create attributes in **Admin > Manage Custom Attributes**.
11. Click **Finish** to start the Discovery Job, which adds the objects to the Object Library and selected destination groups.

About job data gathering after discovery

Read about job data gathering after you discover some applications within DPA.

The information in this section applies to the following applications:

- NetWorker
- Avamar
- TSM
- HP DataProtector
- Commvault Simpana
- NetBackup
- ArcServ
- DB2
- SAP HANA
- RMAN
- MSSQL

With regard to the applications above, note the following:

- When a new server is discovered, DPA gathers job data from 14 days before if you enable this feature.
- The next time the Job Monitor request runs, the current poll time is set to the next day and data is collected for the next day.
- The current poll time is advanced one day at a time from 14 days back every time the Job Monitor request runs, collecting the data for that day until two weeks of data has been collected. Data collection resumes as normal from then on.
- The poll time default value is 1 day and is user-configurable under the Job Monitor request options section.
- When setting data collection, the **Frequency** must always be a lower value than **max data time range each request will gather from**. Otherwise, request does not catch up to the current time and each time the request runs, it falls further behind and does not gather remaining data.

[Data Collection Request Options by Module](#) provides more information.

Monitored objects and groups

Objects overview

DPA discovers the applications and devices in your data protection environment and stores these logical and physical entities as objects in the object library. Discovered objects are grouped into the following categories in the object library:

- Applications
- Hosts
- Storage
- Switches

The following rules apply to objects:

- No two objects can share the same name
- No object can share its name with an alias of another object

The object library enables you to view objects and their attributes.

Searching for objects

About this task

You might search for objects to change Data Collection Requests for multiple objects at once.

Procedure

1. **Select Inventory > Object search .**
2. Type the search criteria:
 - In the **Name** field, type the object name. For example, hostname, application name, switch name.
 - In the **Types** field, select the object type. You can choose top-level object types, like Host and Switch; Backup Server, Backup Client, Backup Pool under Backup Application; and all Application object types.
 - In the **Groups** field, select the object group or Smart Group.
 - In the **Groups** field, select **Not In** if you would like to search for objects that are not included in a group, including Smart Groups. Note that **In** is selected by default.
 - In the **Requests** field, filter by request. If you want to search by requests not assigned, select **Not Assigned**. Note that **Assigned** is selected by default.
 - In the **Agent** field, select the Agent from the Data Collection Agent.
 - In the **Attributes** field, select the attribute. In the **Select Attributes** dialog, if you want to search by attributes not assigned, select **Not Assigned**. Note that **Assigned** is selected by default. If you select Not Assigned, the Value and Clear columns are disabled.

Note the following regarding search for Backup Client, Backup Pool under Backup Application:

- The **Requests** and **Agent** search options are not available with the search for backup clients and pools.
- Data Collection requests and assignments are not available on results of backup clients and pools searches.

The **Types** and **Groups** fields are organized the same as within the Report Scope Configuration tree. If you enter multiple search criteria, they are joined by AND.

3. Click **Search**.

The search displays up to 500 items. To limit the number of items below 500, restrict your search criteria.

Viewing objects

Select **Inventory > Object Library** .

Viewing and editing attributes for multiple objects

Use this procedure to select multiple objects returned from an object search and view and edit the attributes assigned to multiple objects in one action.

Procedure

1. Search for the objects that you would like to view or edit the attributes.
[Searching for objects](#) on page 198 provides information.
2. Select the objects that are returned in the search, and right-click to select **Set Attributes**.
The **Attributes – Multiple Objects** window appears.
3. To edit the attributes for the selected objects, select the check boxes next to the **Name** column and then click **OK**.

Editing data collection for objects

As part of the discovery process, the DPA Discovery Wizard assigns data collection requests directly to an object during object creation. To edit the default data collection requests for a specific object:

About this task

[Searching for objects](#) on page 198 provides additional information on editing data collection requests.

Procedure

1. Select **Inventory > Object Library**.
2. **Select a host and then click the > Data collection > tab.**
3. Click **Properties**.
4. Select a request and then click **Edit**.

Results

[Manage Data Collection Defaults](#) on page 94 provides information on default data collection requests. The *Data Protection Advisor online help system* set provides procedures to add, edit and view data collection requests.

Groups

A group is a collection of objects. For example, you can create a group of objects that are used by an application. This way, when you apply a policy to the group, the policy is applied to all of the objects within the group.

 **Note:** An object can exist in more than one group.

Configuration group

The Configuration group is created by default. The Configuration group is created with an initial structure that groups the data protection environment into Servers, Switches, and Storage. All data protection hosts, devices, and applications discovered by the Discovery Wizard are first added to the Configuration group. Objects that are removed from the Configuration group are not deleted. Objects removed from Configuration group appear under Objects Not In Groups..

Creating groups

Procedure

1. Go to **Inventory > Group Management**.
2. In the object inventory, select **Groups** and click **Create Group**.
3. Type a name for the new group.
4. From the object inventory, select the host or group of hosts that you would like to be in the group.
5. Copy and paste the hosts into the new group you have created.

Ensure that you do not cut or delete the hosts from their original object inventory location.

Object attributes

Object attributes extend the information that DPA holds about an object. After a custom attribute is created, the attribute can be enabled for any valid objects as per custom attribute settings and a value can be assigned.

When creating or editing an object, attributes are filtered to be associated with one or more specific types of objects, and only to objects with an existing attribute that matches a given value.

For example, an Asset Tag attribute might be created to represent an asset identifier for the physical components of an operating environment (such as hosts, storage arrays, and switches). The Asset Tag attribute need not be assignable to logical components like database instances or processes.

In the attribute definition, the Asset Tag is configured to be associated with a subset of physical object types. You can further configure this attribute to only be associated with physical object types that have an attribute of Business Unit, for example.

Smart Groups

Smart Groups allow users with administrative privileges to create groups that are populated dynamically with information from the results of DPA reports. A Smart Group runs a custom report and then creates objects based on the results of the report.

The main benefit of Smart Groups is that they provide high levels of flexibility. Administrators can set up Smart Groups to dynamically create lists of objects that match specific business and technical criteria.

Creating Smart Groups

About this task

The *Data Protection Advisor online help system* provides more information on creating Smart Groups. [Multilevel Smart Group](#) on page 201 and [Single-level Smart Group](#) on page 202 provide more information on these options.

Procedure

1. Select **Inventory > Group Management**.
2. Click **Create Group** and then **Create Smart Group**.
3. Specify a name for the Smart Group in the **Smart Group Name** field.
4. Specify the Time Zone for the Smart Group.
5. Select an option: **Single-level Smart Group** or **Multilevel Smart Group** and click **Configure Smart Group Level**.
6. Specify the **Generation Frequency**:
 - If you would like DPA to generate the Smart Group at a scheduled time, select frequency type **Once a day at** or **Schedule**.
 - If you would like to generate the Smart Group when you create or edit it, select frequency type **On demand**.
7. Specify the fields for each report object chosen and click **OK**.
8. If you would like to configure the Smart Group to store and report on the content nodes historically, set **Enable History** to **On**.
By default **Enable History** is configured to **Off**.
9. Click one of the following:
 - **Save and Run** if the Generation Frequency type is set to **Once a day at** or **Schedule**.
 - **OK** if the Generation Frequency type is set to **On demand**.

Multilevel Smart Group


Unlike Single-level Smart Group, which returns only 1 level of child objects based on the Smart Group, the Multilevel Smart Group can create multiple levels of child objects from a single Smart Group. It also allows you to configure which fields you want to be used in which level, and what type of object you want to be created. There is no limit to the number of levels you can configure. If desired, you could have a complete mapping of your DPA environment using multilevel Smart Groups.

For example, a report used in the Smart Group that returns the data in the following table could be configured to return the object configuration shown in the figure below when run.

Table 37 Multilevel Smart Group example

Customer	Cost Center	Client
Cust1	CC1234	Client1
Cust1	CC1234	Client2
Cust1	CC5678	Client3
Cust1	CC5678	Client4
Cust2	CC1234	Client5
Cust2	CC1234	Client6
Cust2	CC5678	Client7
Cust2	CC5678	Client8

Figure 4 Object library Multilevel Smart Group configuration example



You can assign chargeback and data protection policies to either the Smart Group or to the child objects returned, and see when the structure was last refreshed or generated. By default, the Smart Group generates daily. Additionally, because hierarchical groups can integrate with external data sources, you can create a single hierarchy Smart Group to create the object structure that may already exist in an external system or database.

Only users with permissions to see the Smart Group can see it, expand it, and run reports on it.

Single-level Smart Group

Single-level Smart Group a single set of objects from a report contained in one level of hierarchy. You can assign the same items that you can assign to typical objects, including analyses and scheduled reports. DPA can then generate alerts and reports for a Smart Group outputting objects.

For example, a financial firm might have a convention where the first two characters of each backup client indicate the business unit to which the client is assigned. If the first two characters are a and m, then the backup client belongs to the asset management group. Due to the nature of the business, a large number of clients are created, renamed, or removed daily. Rather than spend a lot of time updating the group configuration each day, the DPA administrator can create a Smart Group that uses the existing Backup Client Configuration report to list each backup client. In the Smart Group, the administrator can filter the results to only contain clients that start with a and m.

As DPA automatically updates the client configuration list every time it obtains data from the backup server, this list is kept up-to-date with whatever changes are made within the backup environment.

Other examples include:

- All backup clients containing exch.
- All hosts with an E: drive.
- All objects with severity 1 alerts in the last day.

Smart Group History

Smart Group History enables you to store and report on the content nodes historically.

The Smart Group History setting allows you to report on changes within Smart Groups, so service providers can provide accurate historical billing.

If the Enable History setting is turned on, then every time the Smart Group is generated subsequently, the history is stored. If the setting is turned off, then all history is deleted and only the current state is stored when the Smart Group is regenerated. By default, the Enable History setting is set to **Off**.

Gathering historical backup data using DPA web console

You can gather historical backup data on Avamar, BackupExec, DB2, HP DataProtector, NetWorker, NetBackup, Oracle RMAN, SAP HANA, and TSM.

About this task

Consider the following when you gather historical backup data using DPA web console:

- You cannot gather historical backup data at the host level. You must go one level down in the configuration tree, to the application object. For example, to collect historical data from NetWorker, choose the Networker application object below the host level object.
- You can only gather historical backup from the JobMonitor requests.

Procedure

1. In the web console, select **Inventory > Group Management**.
2. In the configuration tree, select the application object for which you'd like to gather historical backup data.

The application object **Details** window opens.

3. In the host details window, select the **Data Collection** tab.
4. In **Data Collection**, select the JobMonitor request.
5. Right-click **Run** and select **Gather historical data**.
6. In the **Gather historical data** window, click **OK**.

The same credentials and data options are available as for the request itself.

7. Click **Close** to the a dialog box that appears confirming that DPA is gathering the historical backup data.
8. Click **History** to view collected tests. The rows highlighted in orange indicate results from a historical backup gather.

Configuring policies, rules, and alerts

Policies and alerts overview

DPA contains customizable policies and rules that control how DPA generates alerts, measures backup and replication performance and determines values for chargeback reporting.

Policies

DPA policies are a collection of user data about how backup and replication should operate in the environment (recoverability and data protection policies) or about the cost of storage and data protection operations (chargeback policies).

Recoverability, backup, and service level management reports then show how the operations in the environment compare to the policy settings, for example, gaps in the recoverability chain for a storage array, or if a backup server is not meeting a Recovery Point Objective.

DPA provides the following policy types:

- **Analysis policies** - are a collection of one or more rules that are used primarily for generating alerts. Alerts are displayed by default in the **Alerts** section. You can edit the policy to send events to emails, scripts, SNMP traps, or Windows Event Logs. [Policies and generating events](#) on page 230 provides more information.
- **Protection policies** - are a collection of user data about how backup and replication should operate in the environment. These policies consist of recoverability and protection rules. These are used primarily for generating alerts. Alerts are displayed by default in the **Alerts** section.
- **Chargeback policies** - are used to determine the cost of storage and data protection operations for chargeback reports.

By default, analysis, protection, and chargeback policies are off for all objects and groups.

Analysis policies

An analysis policy is a collection of one or more rules that is assigned to an object or group. Rules contain the logic for when to issue an alert. The analysis engine compares monitored data to the conditions in a rule, and triggers alerts when a rule is matched. Event-based rules trigger an alert in response to data that is streaming into the DPA server. Schedule-based rules periodically compare data in the DPA Datastore against rules to detect a match. Alerts can contain dynamic textual information and might include populated links to reports. Only analysis policies can generate alerts.

Analysis rule template

An analysis rule template is a set of instructions that defines the rules logic. When a rule template is added to an analysis policy, the Analysis Engine carries out certain operations and then displays the resulting events in the **Alerts** section of the web console.

A rule template consists of the name of the rule along with details that specify how that rule is run.

For example, a rule template can be created to monitor whether a file system is likely to exceed 90% utilization in the next hour.

An Analysis Policy contains multiple rules that apply to different object types. The Analysis Engine only runs the rules that are applicable to a given object. For example, if the object is a switch, then the Analysis Engine will only run the rules in the policy that apply to switches.

Event-based rules versus schedule-based rules

Event-based rules work in response to data that is streaming into the DPA server in real time and triggers alerts. There are five types of conditions for event-based rules:

- **Condition filter**—Alert on a set condition; for example, backup failed. Condition filter is the most common condition for event-based rules.
- **Lack of event**—Alert if an event does not occur for defined period of time; for example, Agent is down.
- **Prediction**—Alert if an event occurs in a defined period of time; for example, Filesystem is filling up.
- **Configuration change**—Alert if there is any type of change in your configuration; for example, active or inactive, version, OS type, specific fields, increase or decrease by a certain percentage.
- **Inventory change**—Alert if there is new type of node is auto-created; for example, new RMAN instances.

Schedule-based rules run periodically to check whether to issue an alert. Depending on the type of schedule you have set to collect the data, the alerts could be sent hours after issue was detected in the DPA server.

For both schedule-based rules and event-based rules, you must create a policy that contains a rule, apply the policy to a group of applicable nodes, and ensure that new data that is received for the nodes with the applied policy contains entities that fulfil rules conditions. DPA web console

provides a rich rule editor that allows you to create, edit, and customize both event- and schedule-based rules according to your needs. [Creating an analysis rule](#) on page 205 provides more information.

Guidelines for analysis rules components

Consider the following main components when you are creating analysis rules: the category of the rule that you are setting the alert for, object type that you want to monitor and create the alert for, and the object attributes that you are alerting on.

DPA contains a robust repository of analysis rules system rule templates. Before you create a custom analysis rule, check that one does not exist that fits your needs. Go to **Policies > Analysis Policies > System Rule Templates**. If you select a System Rule Template and edit it, DPA clears out the customizations used to build the policy, which means you do not see how DPA builds the policy.

Analysis rule category

Categories are a way for DPA to store the analysis rules. They are also a way for you to filter and locate analysis rules that you have created. There is no hard and fast rule about choosing a category for analysis rules that you create. If you create a custom analysis rule, select a category from the dropdown that best fits a way that you will remember or find the rule that you are setting. The *Data Protection Advisor online help system* provides information about the analysis policy categories.

Object type and attributes

The object type and attributes you select depend on the scenario on which you want to trigger the alert; for example, the objects you are monitoring and data being gathered about them. If you need assistance with the data being gathered on the objects that DPA monitors, the *Data Protection Advisor Data Collection Reference Guide* provides information on objects and attributes, where the table names within each module function map to an object, and the field name within each table map to an attribute. Within the object type and alert trigger you can configure and further filter this information for the rule.

Creating an analysis rule

Use the DPA rule editor to create an analysis rule template. The following is a high-level overview of the process. The *Data Protection Advisor online help system* provides detailed instructions on how to create, edit, or copy an Analysis Rule template.

About this task

This is a general procedure for creating an analysis rule. Specific examples for event-based and schedule-based analysis rules follow.

Procedure

1. In the DPA web console, navigate to **Policies > Analysis Policies > Rules Templates**.
2. Click **Create Rule Template**.
This opens the rules editor.
3. Provide a name and description for the alert that is triggered by this rule.
4. Select a category associated with the rule.

The *Data Protection Advisor online help system* provides information on rule categories and descriptions.

5. Specify whether the rule is event based or a scheduled rule.

An event-based rule triggers an alert in response to data that is streaming into the DPA server. A Schedule-based rule runs periodically to check whether to issue an alert.

If the rule is a Schedule-based rule, set the **Report Parameters Default Values**.

6. Select the appropriate object types:
 - by hierarchy
 - by function
7. Define when and how the alert must be triggered.

Note that DPA does not support the option to test the `Lack of event` trigger for `Number of samples`, even though the option still appears as valid in the DPA web console. DPA supports the `Number of samples` option for `Time window`.

Creating event-based rules for condition filter

Event-based rules work in response to data that is streaming into the DPA server in real-time and triggers alerts on a set condition; for example, backup failed. The condition filter is the most common condition for event-based rules.

About this task

The procedure below focuses on creating a rule to alert for a failed backup.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to the condition you are setting for the rule.

For example, `backup failed`

You can enter a condition description as well, if you like. This is optional.

There is an existing system template rule called Backup Failed, which you can edit if you like. This example shows you how to create it from scratch.

3. In the **Category** field, select the relevant category from the dropdown that best fits the rule you are setting.

For example, **Data Protection**.

In this case, Data Protection is the most appropriate category because you want to alert on data that is not protected.

4. Configure the object type. For this example, we want to alert on backups that have failed on each backup client, so we select the Backupjob object.

- a. In **Object Type**, click **Select**.

The **Select Object Types** window opens.

- b. Expand Backup Applications, expand the BackupClient object, and then select **Backupjob** from the **Select Object Type** list and click **Select Object Type**.

You can use the filter function to easily find the object you would like to monitor on.

The object type you select depends on the scenario on which you want to trigger the alert.

5. Configure the alert trigger. For this example, we want to look only at failed jobs, so we select the trigger and set conditions filters to find only failed jobs:

- a. In **Alert Trigger**, click **Select**.

The **Select Alert Trigger** window opens.

- b. Select **Conditions Filter** radio button and then click **Select and Edit Filter**.

The **Edit Filter** window opens.

- c. Click **Select Attribute**.

The **Select Attribute** window opens.

- d. Ensure that the **Attribute** radio button is selected and click **Browse**.

The **Browse Attributes** window opens.

- e. From the Backupjob Category select the row with the AttributeName **Status**, click **Select Attribute** and click **OK**.

You can use the filter function to easily find the category and Attributename you would like.

- f. Click **Select Operator** and set a value of **Is** and click **OK**.

- g. Click **Select Value**, select the **Static Value** radio button, select the value of **failed** from the dropdown and click **OK**.

- h. Click **OK** in the **Select Attribute** window and then click **OK** in the **Edit Filter** window.

The scenario for which you are configuring the alert affects how you configure and how, if at all, you further filter rule alert trigger.

6. Configure the alert:

- a. In **Alert**, click **Select**.

The **Edit Alert** window opens.

- b. In the **Alert Fields** tab, select the severity from the dropdown.

- c. In the **Description & Resolution** tab, configure any description and resolution information you want to be sent with the alert.

- d. In the Associated Reports tab, select a system template report or create a custom report that you would like to be generated upon the alert.

7. Click one of the save options.

Creating event-based rules for configuration change

Event-based rules work in response to data that is streaming into the DPA server in real-time and triggers alerts for any type of configuration changes. For example, changes for active or inactive, version, OS type, specific fields, increase or decrease by a certain percentage of any metric that DPA monitors.

About this task

This procedure focuses on a change from client active to inactive.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to configuration change you are triggering the alert for.

For example, `client active status changed`

You can enter a condition description as well, if you like. This is optional.

3. In the **Category** field, select **Change Management** from the dropdown.

4. Configure the object type:
 - a. In **Object Type**, click **Select**.
The **Select Object Types** window opens.
 - b. Expand Backup Applications, and select **Backup Client** from the **Select Object Type** list and click **Select Object Type**.
5. Configure the alert trigger. For this example, we want to look any client that changed from active to inactive, so we select the appropriate trigger.
 - a. In **Alert Trigger**, click **Select**.
The **Select Alert Trigger** window opens.
 - b. Select **Change Control** radio button and then click **Select and Edit Filter**.
The **Edit Alert Trigger - Change Control** window opens.
 - c. Select **ClientConfig** from the dropdown.
 - d. Select the box next to **Active** and click **OK**.
Note that this rule configuration alerts on any changes in this field, not just active to inactive.

No conditions filters are needed because we want to see the configuration change on all clients.
6. Configure the alert:
 - a. In **Alert**, click **Select**.
The **Edit Alert** window opens.
 - b. In the **Alert Fields** tab, select the severity from the dropdown.
 - c. In the **Description & Resolution** tab, configure any description and resolution information you want to be sent with the alert.
 - d. In the Associated Reports tab, select a system template report or create a custom report that you would like to be generated upon the alert.
7. Click one of the save options.

Creating event-based rules for lack of event

Event-based rules work in response to data that is streaming into the DPA server in real-time and triggers an alert if an event does not occur for defined period of time; for example, Agent is down.

About this task

The procedure focuses on creating a rule to alert for production Agent is down, and to keep generating the alert every hour that the production Agent is down.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to the lack of event for which you are setting for the rule.
For example, `DPA Agent down`
You can enter a condition description as well, if you like. This is optional.
3. In the **Category** field, select the relevant category from the dropdown that best fits the rule you are setting.

For example, **Administrative**.

4. Configure the object type. For this example, we want to alert on Agents that have gone down, so we select the AgentStatus.
 - a. In **Object Type**, click **Select**.
The **Select Object Types** window opens.
 - b. Expand the Host object, and then select **AgentStatus** from the **Select Object Type** list and click **Select Object Type**.
You can use the filter function to easily find the object you would like to monitor on.
The object type you select depends on the scenario on which you want to trigger the alert.
5. Configure the alert trigger. For this example, we want to look only at production Agents that have gone down, so we select the trigger and set conditions filters to find only Agents that are down:
 - a. In **Alert Trigger**, click **Select**.
The **Select Alert Trigger** window opens.
 - b. Select **Event/Data Collection Did Not Occurr** radio button and then click **Select and Edit Alert Trigger**.
The **Edit Alert Trigger** window opens.
 - c. For option 1, Select what you want to monitor, select the radio buttons for **Event did not occur** and **AgentStatus**.
 - d. For option 2, select the radio button next to **Keep Generating**.
 - e. For option 3, if you want to specify a type of hostname with a naming convention, for example, *prod* for production, select **Edit Conditions Filter** radio button and then click **Select Attribute**.
 - f. Ensure that **Attribute** radio button is selected for **Value Type** field and click **Browse** for the **Attribute** field.
 - g. In **Browse Attributes**, select the **name** attribute and then click **OK**.
 - h. Click **Select Operator** and set a value of **Contains** and click **OK**.
 - i. Click **Select Value**, select the **Static Value** radio button, in the **Value** field type `prod` and click **OK**.
 - j. Click **OK** in the **Edit Filter** window.
 - k. For option 4, select the radio button next to **Time Period** and select **Static Value** from the drop down and select **1** from the number dropdown and **hours** from the time period dropdown, and then click **OK**.

The scenario for which you are configuring the alert affects how you configure and how, if at all, you further filter rule alert trigger.
6. Configure the alert:
 - a. In **Alert**, click **Select**.
The **Edit Alert** window opens.
 - b. In the **Alert Fields** tab, select the severity from the dropdown.
 - c. In the **Description & Resolution** tab, configure any description and resolution information you want to be sent with the alert.

- d. In the Associated Reports tab, select a system template report or create a custom report that you would like to be generated upon the alert.
7. Click one of the save options.

Creating event-based rules for inventory change

Event-based rules work in response to data that is streaming into the DPA server in real-time and triggers alerts if there is new type of node is auto-created.

About this task

The procedure focuses on creating a rule to alert when an RMAN backup client instance is auto-created.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to the condition you are setting for the rule.

For example, `new RMAN database backed up to central recovery catalog`

You can enter a condition description as well, if you like. This is optional.
3. In the **Category** field, select the relevant category from the dropdown that best fits the rule you are setting.

For example, **Configuration**.
4. Configure the object type. For this example, we want to alert on new RMAN backup client instances, so we select the OracleRMANBackupclient object.
 - a. In **Object Type**, click **Select**.

The **Select Object Types** window opens.
 - b. Expand **Host**, expand the **Applications and Databases**, expand the **Oracle Application**, and then select **OracleRMANBackupclient** from the **Select Object Type** list and click **Select Object Type**.

You can use the filter function to easily find the object you would like to monitor on.

The object type you select depends on the scenario on which you want to trigger the alert.
5. Configure the alert trigger. For this example, we want to look only at newly created objects, so we select the trigger and set conditions filters to find only inventory changes:
 - a. In **Alert Trigger**, click **Select**.

The **Select Alert Trigger** window opens.
 - b. Select **Inventory changes** radio button and then click **Select and Edit Filter**.

The **Edit Alert Trigger - inventory Change** window opens.
 - c. In option 1 **Select operations to monitor**, ensure that **Created** is selected, and then click **OK**.

The scenario for which you are configuring the alert affects how you configure and how, if at all, you further filter rule alert trigger.
6. Configure the alert:
 - a. In **Alert**, click **Select**.

The **Edit Alert** window opens.

- b. In the **Alert Fields** tab, select the severity from the dropdown.
 - c. In the **Description & Resolution** tab, configure any description and resolution information you would like to be sent with the alert.
 - d. In the Associated Reports tab, select a system template report or create a custom report that you want to be generated upon the alert.
7. Click one of the save options.

Creating event-based rules for prediction

Event-based rules work in response to data that is streaming into the DPA server in real-time and triggers alerts of an event occurs in a defined period of time.

About this task

The procedure focuses on creating a rule to alert when an Avamar server is predicted to reach 90% utilized within the next 24 hours.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to the condition you are setting for the rule.

For example, *Avamar server predicted to reach 90% in next 24 hours*

You can enter a condition description as well, if you like. This is optional.

There is an existing system template rule called Backup Failed, which you can edit if you like. This example shows you how to create it from scratch.
3. In the **Category** field, select the relevant category from the dropdown that best fits the rule you are setting.

For example, **Resource Utilization**.
4. Configure the object type. For this example, we want to alert Avamar backup servers, so we select the Backup Application object.
 - a. In **Object Type**, click **Select**.

The **Select Object Types** window opens.
 - b. Expand Backup Applications, expand the Backup Server, and then select **Backup Application** from the **Select Object Type** list and click **Select Object Type**.

You can use the filter function to easily find the object you would like to monitor on.

The object type you select depends on the scenario on which you want to trigger the alert.
5. Configure the alert trigger. For this example, we want to look only at particular backup servers reaching a target utilization within a certain period, so we select the trigger and set conditions filters to find only predictive behaviour:
 - a. In **Alert Trigger**, click **Select**.

The **Select Alert Trigger** window opens.
 - b. Select **Predictive Time** radio button and then click **Select and Edit Filter**.

The **Edit Filter** window opens.
 - c. For option 1, Select attribute to predict, click **Browse**.

The **Select Attribute** window opens.

- d. From the BackupApplication Object Type select the row with the AttributeName **Utilisation**, click **Select Attribute** and click **OK**.

You can use the filter function to easily find the category and Attributename you would like.

- e. For option 2, Set threshold, select **Static Value** and type or scroll up to 90 .
- f. For option 3, Specify when to send alert, select **Static Value** and select **1** and **Days** from the dropdowns.
- g. Skip option 4; there are no conditions filters for this example.
- h. For option 5, Select prediction method, leave the default selection.
- i. Click **OK** .

The scenario for which you are configuring the alert affects how you configure and how, if at all, you further filter rule alert trigger.

6. Configure the alert:

- a. In **Alert**, click **Select**.

The **Edit Alert** window opens.

- b. In the **Alert Fields** tab, select the severity from the dropdown.
- c. In the **Description & Resolution** tab, configure any description and resolution information you want to be sent with the alert.
- d. In the Associated Reports tab, select a system template report or create a custom report that you would like to be generated upon the alert.

7. Click one of the save options.

Creating schedule-based rules

Schedule-based rules periodically compare data in the DPA Datastore against rules to detect a match against a specific problem that you want to track. It uses a report to do this. You can use a System Template report or a custom report.

About this task

The procedure focuses on creating a rule to alert for three strikes failed backup clients.

Procedure

1. Go to **Policies > Analysis Policies > Custom Rule Templates**, and then click **Create Custom Rule Template**.
2. Populate the **Name/Alert Message** field with a rule name relevant to the condition you are setting for the rule.

For example, `schedule based three strikes failed backup`

You can enter a condition description as well, if you like. This is optional.

There is an existing system template rule called Backup Failed, which you can edit if you like. This example shows you how to create it from scratch.

3. In the **Type** field, select **Scheduled** from the dropdown.
4. In the **Category** field, select the relevant category from the dropdown that best fits the rule you are setting.

For example, **Data Protection**.

5. Select the report. For this example, we want to alert on three strikes failed backup clients, so we select the Three Strike Failed Client report.

- a. In **Select Report Template**, click **System Report Templates**.
 - b. Select **Three Strike Failed Client** from the **System Template Name** list and click **Select Template and Edit Options**.
 You can use the filter function to easily find the object you would like to monitor on.
 The object type you select depends on the scenario on which you want to trigger the alert.
6. Configure the options.
 - a. In **Number of Alerts**, select the options that best suits your needs.
 If you select **Generate a separate alert for each row**, DPA sends a different alert for each client. This information is useful because it is granular. However, if you are alerting on a lot of clients you may receive a lot of alerts.
 If you select **Generate one alert for all rows**, DPA sends an alert for top-level nodes. This is useful if you want fewer alerts because you have a lot of clients; however, the information is less granular.
 - b. In **Default settings**, click **Select Schedule** and select one of the **Manage Schedule** options or click **Create Schedule** to create your own schedule that defines when the rule will run.
 We do not recommend selecting **Always** from among the **Manage Schedule** options because this option overloads the server.
 - c. Ensure that you review the time period selection and either leave the default selection or change the selection.
 - d. Click **OK**.
 7. Configure the alert:
 - a. In **Alert**, click **Select**.
 The **Edit Alert** window opens.
 - b. In the **Alert Fields** tab, select the severity from the dropdown.
 - c. In the **Description & Resolution** tab, configure any description and resolution information you want to be sent with the alert.
 - d. In the **Associated Reports** tab, select a system template report or create a custom report that you want to be generated upon the alert.
 - e. In the **Rule Objects** tab, ensure that you select the **Object Type** and select **Name Field** and **Sub Name Field** from the dropdowns.
 8. Click one of the save options.

Adding an analysis rule to an Analysis Policy

After a rule template is added to an Analysis Policy, the Analysis Engine carries out certain operations and then displays the resulting events in the **Alerts** section of the web console.

About this task

The Analysis Policies can contain multiple analysis rules that apply to different types of objects. DPA automatically applies the appropriate rules from the applied Analysis Policy to an object. For example, DPA applies rules for switches to switches only, not to backup servers.

Analysis Engine actions log file

The `actions.log` contains one record for each successful Analysis Engine action notification.

The Analysis Engine actions can be:

- email
- SNMP
- script
- Windows event log

The `actions.log` contains only the information about successful actions. It does not contain failure information or warnings of failing actions. The default location for the `actions.log` is `$instalationDir\services\logs`. This location is not user-configurable.

Analysis policy rule categories

Capacity planning

Capacity planning analysis policies create alerts about events that indicate that resources might soon run out. The following table describes these jobs.

Assigning alerts for pools and storage array analysis policies

When assigning the following analysis policies to objects, the recommended severity levels are:

- Storage pool is filling Up - Severity 3
- Storage pool is filled Up - Severity 2
- Storage Array is filling Up - Severity 1

Table 38 Capacity planning

Rule	Description	Parameters
File system filling up	Generates alerts if a file system utilization will exceed 90% in the next 2 weeks.	Max Predicted Utilization - 100% Number of hours to forecast 336
Running out of backup client licenses	Generates alerts if the license only permits you to monitor less than an additional 25 computers.	Maximum client licenses - 25
Storage pool is filling Up	Alerts when according to the growing trend there will not be space left on the pool for the selected time period.	Minimum Free Space Allowed - 0 Days to Forecast - 90
Storage pool is filled up	Alerts when there is no space on the pool to physically allocate a new LUN.	Initial Consumed Capacity - 3
Storage Array is Filling Up	Alerts when there is no space left to allocate a new LUN on the pool and there are no free disks available on the storage array.	Initial Consumed Capacity - 2

Table 38 Capacity planning (continued)

Rule	Description	Parameters
Empty tapes running low	Generates alerts if there will be no empty tapes available in a tape pool within 6 weeks.	Maximum Predicted Count - 0 Number of hours to forecast - 1008
TSM Database filling up	Generates an alert if the TSM Database is predicted to reach 100% usage within 2 weeks.	Number of Hours to Forecast - 336 Maximum Predicted Utilization - 100
TSM Database utilization high	Generates an alert if the TSM Recovery log is predicted to reach 100% usage within 2 weeks.	Number of Hours to Forecast - 336 Maximum Predicted Utilization - 100

Change management

Change management analysis policies alert about changes in the environment. The following table describes these jobs.

Table 39 Change management

Rule	Description	Parameters
Backup client configuration changed	Generates alerts if the configuration of a backup client has been modified.	N/A
Backup device configuration changed	Generates alerts if the configuration of a backup device has been modified.	N/A
Backup group configuration changed	Generates alerts if the configuration of a backup group has been modified.	N/A
Disk firmware level changed	Generates alerts if the firmware level of a disk has changed.	N/A
Disk serial number changed	Generates alerts if a disk serial number has changed.	N/A
Object operating system changed	Generates alerts if the operating system of a object has changed.	N/A
RecoverPoint Active RPA changed	Generates an alert if the active RPA has changed since the last analysis run.	N/A
RecoverPoint for VMs Consistency Group Copy is disabled	Alert if a RecoverPoint for VMs Consistency Group Copy is disabled	N/A

Table 39 Change management (continued)

Rule	Description	Parameters
RecoverPoint RPA Link Status Changed	Generates an alert if the status of the RPA link has changed since the last analysis run.	N/A
Tape drive firmware level changed	Generates alerts if the firmware level on a tape drive has changed.	N/A
Tape drive serial number changed	Generates alerts if the serial number of a tape drive has changed.	N/A

Configuration

The configuration analysis policies monitor the environment for device or application configuration issues. The following table describes these jobs.

Table 40 Configuration

Rule	Description	Parameters
Backup client inactive	Generates alerts if a backup client is not scheduled to run.	N/A
Fileserver export and LUN on same volume	Generates alerts if a fileserver export is on the same volume as a LUN.	N/A
LUN on given volume	Generates alerts if a LUN has been configured on vol0.	Volume - vol0
IP autonegotiation mismatch	Generates alerts if there is an autonegotiation mismatch between a host and its switch port.	N/A
IP duplex mismatch	Generates alerts if there is a duplex mismatch between object and switch.	N/A
Not enough virtual memory	Generates alerts if the amount of virtual memory on a computer is less than 1.5 times the amount of physical memory.	N/A
Volume priority not normal	Generates alerts when volume priority is set to something other than normal.	N/A

Data protection

The data protection analysis policies monitor the environment for exceptions related to backup and recovery issues. The following table describes the monitored jobs.

Table 41 Data protection

Rule	Description	Parameters
Application restore time estimate too high	Generates alerts if it is estimated that it will take more than 12 hours to restore an application.	Recovery time objective - 12 hours
Application recovery point objective missed	Alert if an application has not had a successful backup in more than 72 hours.	Recovery point objective - 72 hours
Backup failed	Alert generated if a backup fails.	N/A
No Successful backups in one minute	Alert generated if a backup fails two consecutive times.	Maximum failures - 2
Backup larger than average	Generates an Alert if a backup Job is double its size of its average size over the last 14 days.	Days of history - 14 days Deviation - 100%
Backup not occurred for many days	Alert is generated if a host has not had a backup in the last 3 days.	Maximum days not backed up - 3
Backup Running at Same Time as Server Operation	Generates an alert if there were any backups completed over a period that overlapped with any of the following operations on the backup server: <ul style="list-style-type: none"> • Delete volumes • Expirations • Storage pool copies • Moves • Database backup • Migrations • Reclamations 	None.
Backup spans multiple tapes	Alert is generated if a backup spans more than 3 tapes.	Maximum number of tapes - 3
Full backup smaller than average	Generates alerts if a Full backup is less than 50% of its usual size.	Days of History - 14 days Deviation - 50%
Full backup not occurred for many days	Generates alerts if a host has not had a successful full backup in the last 14 days.	Maximum Days Not Backed Up - 14

Table 41 Data protection (continued)

Rule	Description	Parameters
Mirror not updated for a number of hours	Generates alerts if a Remote Disk Mirror has not been updated in at least 2 days.	Maximum Exposure - 48 hours
Too many backups without a full	Generates alerts if there have been more than seven runs of a backup Job since the last Full backup.	Maximum Non Fulls - 7
No NetWorker bootstrap generated	Generates an alert if there has not been a NetWorker bootstrap ran in the last 48 hours.	Maximum hours without bootstrap - defaults to 48 hours
TSM Database Backup Running at Same Time as Server Operation	Generates an alert if a database backup process completed while there was other activity on the backup server, including other backups	None.
TSM Database Backup Occurred	Alerts if there was a TSM database backup in the last 24 hours, or returns the last TSM backup time if there was no backup.	Time - 24 Hours

Licensing

The licensing analysis policies monitor the environment and generate alerts about licensing issues. The following table describes these policies in more detail.

Table 42 Licensing

Rule	Description	Parameters
License expired	Generates an alert if a license in DPA has expired.	N/A
License nearing expiration	Generates an alert if a license will expire in the next week.	Minimum days before expiry - defaults to 7 days

Performance

The performance analysis policies monitor the environment and generate performance problem alerts. The following table describes these jobs in detail.

Table 43 Performance

Rule	Description	Parameters
Backup slower than average	Generates an alert if the performance of a backup job is 50% less than its average over the last 2 weeks.	Days of history - 14 Deviation - 50%

Table 43 Performance (continued)

Rule	Description	Parameters
Backup Job overrunning	Generates an alert if a backup has been running for more than 18 hours.	Max Runtime - 18 hours
Fileserver cache hit rate low	Generates alerts if the cache hit rate of a fileserver drops below 80%.	Minimum cache hit rate - 80%
Full backup succeeded but slow	Generates an alert if a full backup ran at less than 300 KB/sec.	Minimum expected speed - 300 KB/sec

Provisioning

The provisioning analysis policies generate alerts about events that might require provisioning operations. The following table describes the jobs.

Table 44 Provisioning

Rule	Description	Parameters
File system snapshot space under utilized	Generates alerts if the peak snapshot usage over the last 14 days is less than 80%.	Days to examine usage - 14 Minimum peak snapshot usage - 80%

Recoverability

Recoverability analysis policies alert about Recoverability. The following table describes these jobs.

Table 45 Recoverability

Rule	Description	Parameters
DR Host Visibility Check for TF/Snap	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for MirrorView/A	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for MirrorView/S	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for RecoverPoint/A	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for RecoverPoint/S	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A

Table 45 Recoverability (continued)

Rule	Description	Parameters
DR Host Visibility Check for SanCopy	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for Continuous SRDF/S	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
DR Host Visibility Check for Point in Time SRDF/S	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Recoverability Exposure	Recoverability Exposure	N/A
Consistency Group Check for Continuous SRDF/A	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Consistency Group Check for Point in Time SRDF/A	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Consistency Group Check for Continuous SRDF-S/EDP	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Consistency Group Check for Point in Time SRDF-S/EDP	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
DR Host Visibility Check for SRDF/A	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for SRDF/S	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for SV/Clone	that the devices of the recovery-point are mapped,	N/A

Table 45 Recoverability (continued)

Rule	Description	Parameters
	masked and visible by the DR-host	
DR Host Visibility Check for SV/Snap	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for TF/Mirror	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
DR Host Visibility Check for TF/Clone	that the devices of the recovery-point are mapped, masked and visible by the DR-host	N/A
Consistency Group Check for Continuous SRDF-A/EDP	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Consistency Group Check for Point in Time SRDF-A/EDP	Check that the devices of the recovery-point are configured in the same consistency group and that the consistency group is enabled	N/A
Consistent Device Replication Check for Point in Time SV/Clone	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time SV/Snap	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time TF/Mirror	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time TF/Clone	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A

Table 45 Recoverability (continued)

Rule	Description	Parameters
Consistent Device Replication Check for Point in Time TF/ Snap	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time MirrorView/A	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time MirrorView/S	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time RecoverPoint/A	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time RecoverPoint/S	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time SanCopy	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Consistent Device Replication Check for Point in Time SNAPVX	Best practice check. An alert will be generated in case the consistency operation was not issued; but it is recommended by the vendor.	N/A
Application Consistency Violation	Inconsistent Replication: Application was not in backup mode during replication process	N/A
Application Not in Backup mode	Application Not in Backup mode during replication creation	N/A
Consistency Group is disabled	Consistency Group is disabled.	N/A
Invalid Replication	Images for object have failed or missed schedule	N/A

Table 45 Recoverability (continued)

Rule	Description	Parameters
Logs not on Disk	Application log file is not found on disk	N/A
Not all the devices are part of a replication group	The device is not part of a Replication group	N/A
Not Protected Logs	Inconsistent Replication: The file is required for recovery but was not protected	N/A
Partially Replicated	Object is partially replicated	N/A
The continuous replication is halted	The continuous replication is halted	N/A
Storage Object Not Protected	Application Storage Object not protected	N/A
The link status for a continuous replication is down	The link status for a continuous application is down	N/A

Resource utilization

Resource utilization analysis policies generate alerts about events that have occurred because of resource utilization problems within the environment. The following table describes these jobs in detail.

Table 46 Resource utilization

Rule	Description	Parameters
Aggregate snapshot utilization high	Generates an alert if an aggregate snapshot utilization is higher than a specified threshold.	Maximum aggregate snapshot utilization - default is 90%
CPU pegged	Generates an alert if the CPU Utilization on a host is greater than 90% for last 30 minutes.	Maximum CPU utilization - defaults to 90% Number of minutes - 30 minutes
Disk pegged	Generates an alert if a disk on a host is greater than 90% busy for over 30 minutes.	Maximum Disk Busy Percentage - 90% Number of minutes - defaults to 30 minutes
Fibre Channel port utilization high	Generates an alert if a Fibre Channel port exceeds 70% of its max throughput.	Maximum utilization - 70%
Fibre Channel port no BB credits	Generates an alert if a Fibre Channel port has ran out of buffer to buffer credits.	N/A
File system file utilization high	Generates an alert if the number of files on a file	Maximum file system file utilization - 90%

Table 46 Resource utilization (continued)

Rule	Description	Parameters
	system is greater than 90% of the max number allowed.	
File system snapshot utilization high	Generates an alert if a file systems snapshot utilization is above 90%.	Maximum file system snapshot utilization - defaults to 90%
File system utilization high and increasing	Generates alerts if a file system utilization is above 90% and is increasing.	Maximum file system utilization - defaults to 90%
Memory utilization high	Generates an alert if memory utilization on a host is greater than 90%.	Maximum memory utilization - defaults to 90%
Network utilization high	Generates an alert if a network interface exceeds 70% of its rated throughput.	Maximum utilization - defaults to 70%
RecoverPoint Journal Utilization High	Generates an alert if the journal utilization for an RPA is above a specified warning or critical threshold.	Warning threshold Critical Threshold
RecoverPoint Journal Utilization High	Generates an alert if the SAN utilization for an RPA is above a specified warning or critical threshold.	Warning threshold Critical Threshold
RecoverPoint RPA WAN Usage High	Generates an alert if the WAN utilization for an RPA is above a specified warning or critical threshold.	Warning threshold Critical Threshold
RecoverPoint Replication Lag High	Generates an alert if the replication time or data lag is above a specified warning or critical level.	Time Lag Warning threshold Time Lag Critical Threshold Data Lag Warning threshold Data Lag Critical Threshold
TSM Database Utilization High	Generates an alert if the TSM Database utilization exceeds 90%.	Maximum Database Utilization - 90%
Expiration Process Duration Exceeds Expectation	Generates an alert if the TSM Expiration process take longer than an hour to run, or more than 25% longer that the average expiration process time over the last seven days.	% Increase - 25% Period - 7 Max Duration - 1
TSM Recovery Log Utilization High	Generates an alert if the TSM Database utilization exceeds 90%	Maximum Recovery Log Utilization - 90%

Service Level Agreements

Service Level Agreement (SLA) analysis policies generate alerts about SLA violations. The following table describes the SLA jobs.

Table 47 Service Level Agreement

Rule	Description	Parameters
Backup succeed but failed SLA requirements	Generates an alert if a backup was successful but outside of its backup window.	N/A

Status

Status category analysis policies generate alerts when there is concern of the current status of a monitored device or application match. The following table describes status jobs.

Table 48 Status

Name	Description	Rule	Parameters
Backup Server Errors	Generates an alert if a backup server error is logged (TSM only).	Backup server errors	N/A
CG Copy for VMs link down	CG Copy of RecoverPoint for VMs is enabled and the link is down.	CG Copy for VMs link down	Entity; CgCopyStatus Condition "enabled is false (or 0- please check)" fields: data transfers not "Active"
CPU Offline	Generates an alert if a CPU is offline.	CPU offline	N/A
Agent Heartbeat Failed	Generates an alert if an agent fails to send in its heartbeat.	Agent heartbeat failed	N/A
Agent Log File Message	Alerts on any message that appears in the agent log files.	Agent Log Messages	N/A
Disk Failed	Generates an alert if a disk has failed.	Disk failed	N/A
EDL Failover occurred	Generates an alert if one EDL appliance fails over to another.	EDL Failover Occurred	N/A
Fan Inactive	Generates an alert if a fan on a device is inactive.	Fan inactive	N/A
Fibre Channel Port Changed State	Generates an alert if a Fibre Channel port has changed state.	Fibre Channel port changed state	N/A

Table 48 Status (continued)

Name	Description	Rule	Parameters
Less than 75% of Backup Devices Available	Generates an alert if less than 75% of the backup devices on a backup server are Up.	Less than x% of backup devices available	Lowest backup device availability - defaults to 75%
More Than 3 Backup Devices Unavailable	Generates an alert if there are more than 3 backup devices on a backup server Down.	Many backup devices unavailable	Maximum number of downed devices - 3
Network Interface Changed State	Generates an alert if network interface gets a link up or link down event.	Network interface changed state	N/A
Object Restarted	Generates an alert if a host has been rebooted.	Object restarted	N/A
Object Status not Up	Generates an alert if a object's status changes to anything except active.	Object Status not Up	N/A
PSU Inactive	Generate an alert if a Power Supply Unit is not active.	PSU inactive	N/A
Publisher Hung	Generates an alert if the Publisher queue hasn't changed since the last poll.	Publisher Queue Hung	N/A
Server Log File Message	Alerts on any messages appearing in server log files.	Server Log Messages	N/A
Tape Drive Needs Cleaning	Generates an alert if a tape drive needs cleaning.	Tape drive needs cleaning	N/A
Tape Drive Not Okay	Generates an alert if a tape drive is reporting a status other than OK.	Tape drive not okay	N/A
Tape Library Not Okay	Generates an alert if a tape library is reporting a status other than OK.	Tape library not okay	N/A
Thermometer Inactive	Generates an alert if a thermometer becomes inactive.	Thermometer Inactive	N/A
Thermometer Overheating	Generates an alert if a thermometer on a	Thermometer overheating	N/A

Table 48 Status (continued)

Name	Description	Rule	Parameters
	device indicates that it is overheating.		
Waiting For Writable Tapes For More Than 30 Minutes	Generates an alert if a backup server has been waiting more than 30 minutes for a writable tape.	Waiting for writable devices	Maximum outstanding devices - defaults to 0 Minutes before alerting - defaults to 30 minutes
Xsigo Fan Less Than 90% of Normal Speed	Generates an alert if the speed of a fan on a Xsigo Director falls below 90% of the normal speed.	Xsigo Fan Speed Less than Expected	Percentage to Check - defaults to 90%.

Troubleshooting

The troubleshooting analysis policies provide help for troubleshooting problems with the environment. The following table describes these jobs.

Table 49 Troubleshooting

Rule	Description	Parameters
Backup failed due to client network errors	Generate an alert if a backup failed on a client while it experienced an increase in network errors.	N/A
Backup job failed due to high client CPU utilization	Generate an alert if a backup failed on a client, while the CPU utilization on the computer was greater than 90%.	Maximum processor utilization - defaults to 90%
Backup job failed due to high client memory utilization	Generates an alert if a backup failed on a client whilst the memory utilization on that client was greater than 90%.	Maximum memory utilization - defaults to 90
Backup failed due to high server CPU utilization	Generates an alert if a backup failed on a client whilst the CPU utilization on the backup server was greater than 90%.	Maximum processor utilization - defaults to 90%
Backup failed due to high server memory utilization	Generates an alert if a backup fails whilst the memory utilization on the backup server is greater than 90%.	Maximum memory utilization - defaults to 90%
Backup failed due to server network errors	Generates an alert if a backup failed while there was an increase in the number of network errors on the backup server.	N/A

Table 49 Troubleshooting (continued)

Rule	Description	Parameters
Disk failed for a number of hours	Generates an alert if a disk is in a failed state for more than 48 hours. Applicable to Linux and Solaris.	Maximum failure time - defaults to 48 hours
Fibre Channel port reporting errors	Generates an alert if a Fibre Channel port is reporting errors.	N/A
Fibre Channel port reporting more than x% errors	Generates an alert if more than 1% of all frames going through a Fibre Channel port have errors.	Maximum percentage errors - defaults to 1%
Network interface reporting errors	Generates an alert if errors are being seen on a network interface.	N/A
Network interface reporting more than x% errors	Generates an alert if more than 1% of the packets travelling through a network interface have errors.	Maximum percentage errors - defaults to 1%
Tape drive reporting errors.	Generates an alert if there is an increase in the number of errors seen on a tape drive.	Include Recoverable Errors - defaults to False

Protection policies

Protection policies are used to define service level agreements and exposure reporting to calculate whether a backup ran in its backup window and to calculate whether an application or host is meeting its recovery time objective (RTO) and recovery point objective (RPO). Protection policies also determine how an application, host, or device should be replicated or backed up. Policies are assigned to objects and consist of a set of rules that dictate:

- For replication: the type of copy, the replication level, and the schedule.
- For backups: the level of backup and the schedule.

DPA reports then compare the protection policy for an object to the actual replication or backup taking place to display the level of compliance with policy.

Recoverability checks

Recoverability checks are additional consistency checks that DPA performs on an environment, if you configure recoverability analysis. A recoverability check verifies that the storage and recoverability environment is configured to a user's particular requirement; for example, disaster recovery.

If you enable a recoverability check and DPA detects an inconsistency, a recoverability check generates an exposure just like an exposure generated by a Protection Policy breach or a Recoverability request. Recoverability check exposures are displayed in the Exposure reports.

There are three system recoverability checks that identify gaps, as described in the following table.

Table 50 Recoverability checks

Recoverability check	Description
Consistency Group Check	Checks whether the devices of the recovery point are configured in the same consistency group and the consistency group is enabled. If no consistency group exists, a consistency violation gap is generated for the recovery point.
Consistent Device Replication Check	Checks whether the consistency option was used when the images were created, when applicable. This is a best practice check. If the consistency option was not used, then a Consistency Violation gap is generated for the recovery point.
DR Host Visibility Check	Checks whether the devices of a recovery point are mapped, masked, and visible to the Disaster Recovery host. Otherwise, a Consistency Violation gap is generated.

Chargeback policies

Chargeback reports provide the ability to perform a financial cost analysis for backups, restores, and data protection replication operations in a customer's environment. DPA calculates a cost for each backup client and can be charged back to the business unit that is responsible for that client or set of clients.

DPA calculates chargeback using two models: one for data backup and restore, and one for the protection and replication of storage data by RecoverPoint. DPA calculates chargeback for clients based on the inputs for each type.

Backup chargeback

DPA breaks out backup chargeback by cost per GB backed up and other backup costs.

Cost Per GB Backed Up uses the following inputs:

- Base Size - Baseline backup size in GB for base costing.
- Base Cost - Total cost for backup up to the base size specified.
- Cost of Each Additional GB - Additional cost per GB for backups greater than the base size.

DPA derives other Backup Costs from the Chargeback Policy and uses the following inputs:

- Cost Per Backup - the cost per backup (derived from the chargeback policy).
- Cost per GB Retained - the cost per gigabyte stored (derived from the chargeback policy).
- Cost Per Restores - the cost per restore (derived from the chargeback policy).
- Cost per GB Restored - the cost per gigabyte restored (derived from the chargeback policy).
- Cost Per Tape - the cost per tape used for backup (derived from the chargeback policy).

Storage chargeback

DPA breaks out storage chargeback by cost per GB stored, cost per GB replicated, and snaps.

Cost Per GB Stored uses the following inputs:

- Cost Based On - chargeback calculated on either storage used or storage allocated.
- Base Size - Amount of base storage space allocated in GB.
- Base Cost - A one-off price for the base size.
- Cost of Each Additional GB - the price per GB after base size is exceeded.

Cost Per GB Replicated uses the following inputs:

- Base Size - Amount of base storage space allocated in GB.
- Base Cost - A one-off price for the base size.
- Cost of Each Additional GB - the price per GB after base size is exceeded.

Snaps uses the following inputs:

- Cost Per GB - the price per GB.

A Chargeback Policy allows you to specify a value for each of these parameters. DPA calculates the total cost for a client by adding each of the different cost elements. For example, if you want to implement a chargeback model where you charge \$5 for each backup that took place and \$0.20 for each GB that was backed up, then you can specify values for these fields in the chargeback policy but not specify values for the other parameters.

You assign a backup client objects a cost center, which allows DPA to calculate Chargeback costs by cost center. A default cost center exists for objects that have not been assigned a cost center.

You can create multiple chargeback policies, and different clients or groups of clients can have different policies assigned to them. For example, if you wanted to calculate the chargeback cost for one group of backup clients based on the number of backups performed and another group based on the number of tapes used during the backup process, you can create two chargeback policies and associate them with each group of clients.

Policies and generating events

When an analysis policy finds a matching condition, DPA generates an event. All events are automatically logged in to the DPA Datastore. You can view all events in the **Alerts** section of the web console.

You can edit policies to:

- generate an email
- run a script
- send an SNMP trap
- write an event to a Windows Event Log

Editing rules in policies

To edit all the rules in the policy, go to **Policies > Analysis Policies > Edit > Edit Policy-based actions**.

About this task

Alternatively, edit actions on a per-rule basis. To edit actions on a per-rule basis:

Procedure

1. Go to **Policies > Analysis Policies > [select a policy] and click Edit**.
2. Under Analysis Rules, highlight the rule name to edit, and click **Edit Actions**.
3. In the **Edit Actions** window, ensure that the Rule-based actions radio button is selected.

Alternatively, edit or overrule all the rules in a policy or on a per-rule basis from the Inventory area. This is applicable only to the roles that have permissions to edit the policy.

4. Go to **Inventory** and select the object.
5. Select **Properties**.
6. Within the object **Details** window, click the **Policies** tab.
7. Click **Edit Override Settings** .

Edit Override Settings is available only if the role has privileges to do so. Otherwise, the option is **View Settings**

8. Within the object **Override Policy Settings** window, make applicable changes, either on a per-rule level or at a policy level; and click **OK** when finished making changes.

Results

The *Data Protection Advisor online help system* provides additional information on create, edit, or copy an Analysis Rule template.

Parameters for generating alerts from scripts

You can place scripts in any directory. However, we recommend that you use the `<install-dir>/services/shared/` directory because in a clustered environment, you will need to put the scripts just once. If you choose a different location, in a clustered environment you will need manually to copy the scripts to every DPA Application server.

The following table describes the parameters to the script to use to perform actions.

Table 51 Script field parameters

Parameter	Description
Node	Name of the node to which the alert applies.
Text	Textual error message as defined in the ruleset.
Severity	Severity of the alert (Critical, Error, Warning, Informational).
Name	Name of the analysis that triggered this alert.
Alert ID/Event ID	ID that uniquely describes this alert.
First occurrence	Timestamp that details the time that this alert first occurred.
Last occurrence	Timestamp that details the time that this alert last occurred.
Count	Number of times this alert has been issued.
View	Name of the view to which the analysis is assigned.
Node	Name of the node to which the analysis is assigned.
Category	Category of the rule (possible values: Administrative, AssetManagement, CapacityPlanning, ChangeManagement, Compliance, Configuration, DataProtection, Execution, Performance, Provisioning,


Table 51 Script field parameters (continued)

Parameter	Description
	Recoverability, ResourceUtilization, SLA, Status, System, Troubleshooting).

The following table describes the arguments that are passed to a script in an alert action.

Table 52 Script alert arguments

Argument	Description
\$1	Event node.
\$2	Event message.
\$3	Event severity (as set in the analysis properties).
\$4	Name of analysis that caused the event.
\$5	Alert ID (unique for this run of the script).
\$6	Event ID (unique for this alert).
\$7	First occurrence (timestamp).
\$8	Last occurrence.
\$9	Count.
\$10	Category.
\$11	Description of the alert.

 **Note:** If you are running a script in a UNIX environment, you must enclose parameters with 2 digits in curly brackets: {xx}. For example, \$ {11}.

Rule Template

A rule is the set of instructions that the DPA Analysis Engine uses to determine whether a condition has been met and if an alert is generated. For example, the file system filling up rule contains the set of rules to determine if any file systems will exceed the threshold at a certain point in the future.

An Analysis job uses a rule to perform analysis and alerting based on information within the DPA database. When DPA is installed, a number of pre-defined rules are installed that can monitor for common problems that might occur in the environment. You can use these rules as the basis for implementing an analysis policy. DPA provides a rules editor that you can use to create entirely new rules.

The term *rule template* is used to differentiate the rule definition from the rule instance. The rule template defines the rule's logic. When a rule template is added to an analysis policy, it becomes a rule instance (or a rule) that the Analysis Engine will run. Also, when rule templates are added to a policy, users can specify the values for any parameters. This allows rules to be reused by different policies.

For example

A Tier 1 policy might generate an alert when disk space is 80% utilized. A Tier 2 policy can generate an alert when disk space is 90% utilized. This can be handled with the same rule template that uses a parameter for utilization.

Policy application

You can apply policies directly to a group or an object. Policies that are applied directly to an object always take precedence. When you set a policy at the group level, objects in the group that do not have their own policies, they inherit the group's policy. The best practice is to apply the policy at the highest group level. Policies cannot be applied to Smart Groups.

If an object is moved from one group to another group, the most recently applied policy is implemented. For example, if you move an object from Group A to Group B, the object inherits the policy of Group B.

An administrator or any user with the Edit Node privileges can apply a policy to a group or object.

CHAPTER 5

Uninstalling DPA

This chapter includes the following sections:

- [Uninstalling the software](#).....236
- [Agent-only uninstallation](#).....236

Uninstalling the software

This section describes how to uninstall DPA in both UNIX/Linux and Windows environments.

Procedure

1. Run the following command:

```
<DPA_install_directory>/_uninstall/ Uninstall_Data_Protection_Advisor
```

Add `-i silent` to the command if you want silent uninstallation. The uninstaller will not ask you for input.

Results

When uninstalling the DPA Datastore, a warning indicating that the uninstaller will remove the features that were installed during product installation appears indicating that the database will be removed.

Uninstalling by using silent command line

About this task

- On UNIX/Linux machines, start a command shell, navigate to the `_uninstall` directory and type the following command: `./Uninstall_Data_Protection_Advisor -i silent`
- On Windows machines, type the following command through the command line:
`Uninstall_Data_Protection_Advisor.exe -i silent`

Uninstalling through user interface on Windows

About this task

Procedure

1. Select **Start > Control Panel > Programs and Features**.
2. Uninstall **Data Protection Advisor** from the list of installed applications.

Agent-only uninstallation

You cannot uninstall only the Agent from the DPA Application server or Datastore server installation.

If you would like to upgrade the DPA Agent, upgrade the Agent only on the existing DPA Application server or Datastore server installation. [Upgrades](#) on page 64 provides information on carrying out upgrades.

CHAPTER 6

Troubleshooting

This chapter includes the following sections:

- [Installation troubleshooting](#)..... 238
- [Log files](#).....239
- [Data collection troubleshooting](#)..... 242
- [Client/Storage discovery for replication analysis troubleshooting](#)..... 243
- [Troubleshooting report output failure](#).....248
- [Troubleshooting report generation or publishing problems](#)..... 248
- [System clock synchronization](#)..... 249

Installation troubleshooting

DPA Agent does not restart or register after DPA Server password change

If the DPA Agent does not restart or register after the DPA Server password was changed during installation, it could be because the Agent password on the DPA Server has been changed and the password on the DPA Agent has not been changed to match it.

To get the DPA Agent to restart or register, set the password on the Agent to the same value as is set on the DPA Server. [Installing the DPA Agent](#) on page 51 provides information.

DPA Datastore on Linux failure to start after installation

In certain circumstances the Kernel settings of the system running the DPA Datastore may need to be tuned for the Datastore to start up correctly.

If the Datastore fails to start and errors in the DPA log file reference shared memory segments, then the values specified in the following file may need to be tuned according to your system specifications.

- Linux: Investigate tuning values for SHMMAX and SHMMIN in the `/etc/sysctl.conf`

DPA web console launch failure on Windows Server 2012

If the DPA web console fails to launch on Windows Server 2012, check the following items:

- The Internet Explorer Enhanced Security Configuration (IE ESC) stops the DPA web console from launching. Do not stop the notification of the block by clearing the Continue to prompt when website content is blocked option because DPA never comes past Starting services. Please wait.
The workaround for this is to disable the IE ESC.
- Internet Explorer in Windows server 2012 doesn't support Flash. The workaround for this is to enable Desktop Experience in Windows server 2012.

Postinstallation memory adjustment

When the DPA Application and Datastore services are originally installed, they automatically tune memory parameters based on your system RAM. If at a later stage you either increase or decrease the amount of RAM installed on the host you must run the `tune` command so that the DPA memory parameters are adjusted correctly.

When you run the `tune` command, you must specify the amount of RAM installed on the host. For example, if the Application server memory is changed to 64GB and the Datastore memory is changed to 32GB, you would run the following commands:

- On the application server: `dpa app tune 64GB`
- On the datastore server: `dpa ds tune 32GB`

DPA automatically configures itself to use a portion of the memory amount specified in the command.

Error messages during upgrades

If there is an error during the upgrade process, the DPA server stops. This could occur under the following circumstances:

- Errors in SQL upgrade scripts
 - Result: The server stops and does not continue.
 - Suggested action: Contact EMC Technical Support.
- Errors in system metadata upgrade; for example, system reports, rule templates
 - Result: The server stops, but you have the option to continue the upgrade.
 - Suggested action: You can disregard this message and continue with the DPA server upgrade. However, the DPA system might be unstable. If you do stop the server upgrade, Contact EMC Technical Support
- Errors in the custom data upgrade; for example, custom analysis rules
 - Result: An error message is thrown indicating the problem.

Suggested action: You can disregard this message and continue with the DPA server upgrade. However, you should expect the custom rule that failed to upgrade not to work. An error is recorded in the log file.

Log files

Log files provide important information when troubleshooting problems.

Note: The following section describes the log file locations for a standard DPA installation. If the default installation directory was changed during installation, the location of the log directory will be different.

By default, logs contain warnings and error and informational messages. These may not provide enough information when troubleshooting complex problems.

Changing default log detail level

Go to **Admin > System > Configure System Settings**.

Viewing install log file

The `Data_Protection_Advisor_Install_[two-digit date]_[two-digit month]__[year]_[two-digit hour]_[two-digit minute]_[two-digit seconds].log` file is generated during installation and contains all log messages. For successful installations, you can find this file in the install directory (for example, `/opt/emc/dpa/_install`). For unsuccessful installations on UNIX platforms, you can find the file in the root of the system drive. On Windows platforms, you can find the file on the desktop.

Viewing server log files

DPA generates the server log files in the following locations:

About this task

- UNIX: `/opt/emc/dpa/services/logs`

- **Windows:** `C:\Program Files\EMC\Data Protection Advisor\services\logs`

Server log files

The default location for following log files is `<install_dir>\services\logs\ .`

- `Server.log`—Contains all log comments generated from the DPA Application Server
- `actions.log`—Contains successful Analysis Engine actions
- `reportengine.log`—Contains all log comments generated from the DPA Report Engine
- `listener.log`— Contains all log comments generated from the DPA Listener related to the server receiving agent data and processing it

Viewing agent log files

The agent log files are generated in the following locations:

About this task

- **UNIX:** `/opt/emc/dpa/agent/logs`
- **Windows:** `C:\Program Files\EMC\Data Protection Advisor\agent\log\agent.log`

Managing log files

When a log file reaches its maximum size, and the maximum number of log files exist in the log file directory, DPA deletes the oldest log file for that process and creates a new log file. You can modify the maximum log file size and maximum number of log files. You can also change the location of log files, if required.

About this task

Enabling alternative log rotation on VMs running Windows

There is a known issue on VMs running Windows that causes the logs not to rotate due to the file being locked. To fix this, enable the alternative log rotation method. This will change the way the logs are being used, where the highest numbered log is the latest and not the `agent.log` file. This is pertaining to DPA-24288.

Procedure

1. Create the following string registry:
`HKLM\SOFTWARE\EMC\DPA\Agent\ALTLOGROTATE`
2. Set the value to *true*.

Erroneous memory data in installer log file

The Free Memory and Total Memory data indicated at the top of the installation log files is erroneous. The correct Free Memory and Total Memory data is located further down in the log file, under `STDERR ENTRIES`.

The Corrected Total Memory data indicated under `Executing IAUpdatePostgesconfFile: [INFO]` refers to data being used for the DPA Datastore service.

Running a DPA Agent request in debug mode using DPA web console

The DPA Agent request in debug mode, also sometimes called a *modtest*, is a support tool. If you are encountering problems with a Data Collection Defaults, an EMC Technical Support Engineer

may ask you to run the Agent request debug mode from the DPA web console. You can run DPA Agent request in debug mode, download the zip file directly from the DPA web console with no need of going to DPAServer to retrieve the zip file, and send the zip file for analysis. The Agent request debug mode runs the selected request and retrieves the output and the log messages, in debug log level, and by default stores that report xml as a zip file to the following location:

<DPA_HOME>\services\shared\modtests, where <DPA_HOME> is the location of the DPA installation.

About this task

Consider the following when running DPA Agent request in debug mode using DPA web console:

- The test cannot be run if the Collection Request is disabled.
- The test cannot be run if the Collection Request isn't applicable on the object.
- If you are running Google Chrome: you should change the default security setting for the URL to low:
Go to **Trusted Sites**, add the URL to Trusted Sites list, and set security to **low**.

Procedure

1. In the web console, select **Inventory > Object Library**.
2. In the Object Library, select the DPA server under **All hosts**.
3. In the host details window, select the **Data Collection > tab**.
4. In **Data Collection**, select the Request.
5. Right-click **Run** and select **Run in Debug**.
6. In the **Run in Debug - host/status** window, select credentials and data options.
7. Click **Close** to the a dialog box that appears confirming that the test is running.
8. Click **History** to view collected tests. The rows highlighted in orange indicate results from a DPA Agent request in debug mode.
9. Click the test result. If a Windows Security Login appears, enter your DPA server credentials and click **OK**.
10. To access the successfully collected tests, go to <DPA_HOME>\services\shared\modtests.

If you are on a remote web browser, you can download a link which allows you to transfer the zip to your machine (where the browser is) if you look at the history for the request and click on the orange modtest line.

Default modtest deletion schedule

DPA deletes modtest files from the DPA server weekly on Sunday at 4:00 a.m. DPA removes all test results files older than seven days. This schedule is not configurable.

Generate Support Bundle

The Generate Support Bundle option is a support tool. [Generate Support Bundle](#) on page 91 provides information.

Data collection troubleshooting

This section describes the steps that you can take to diagnose problems when trying to gather data. We assume the following scenario:

- DPA was successfully installed.
- The Discovery Wizard was successfully run to create the object to monitor.
- Requests have been assigned to the object and the agent has been reloaded.
- Sufficient time (fifteen minutes) has passed to allow the agent to gather data.
- An appropriate report has been run that returns no data when data should exist for the object.

Troubleshooting data collection: first actions

Review any errors returned by the Agent Errors report and take corrective action if possible; for example, resolve an authentication problem.

Procedure

1. Verify that the time period selected for the report is correct.
2. Check that the correct requests have been assigned to the object.
Select **Inventory > Object Library > [select node] > Data Collection**. Verify that the requests are configured correctly.
3. Rerun the request.

Troubleshooting data collection: second actions

Procedure

1. If no resolvable agent errors are reported, select **Admin > System**, click **Configure System Settings**, and verify the data collection agent settings.
2. If the status shows that the agent is active, verify that the process is active on the operating system on which the agent is installed.
3. Run the Agent log reports in the web console followed by the Agent Status, and then the Data Collection History report.
4. Rerun the report. If the report continues to show no data, open the agent log and look for any problems. For example, was an incorrect value entered during agent installation. [Log files](#) on page 239 describes how to view the log files.

Preparing a log file for submission to EMC Support

Procedure

1. Set the Log Level of the process to **Debug** in System Settings, as described in [Log files](#) on page 239.
2. Stop the agent process.
3. Navigate to the directory in which the log file is stored. Rename or remove all existing log files for the process.
4. Restart the process.

Restarting an agent reloads all the requests assigned to that agent and starts the data gathering routine. This ensures that all requests have been attempted. Starting a new log file removes the need to search through unnecessarily long log files for a problem.

5. Select **Inventory > Object Library > [select node] > Data Collection** and then select **History**.
Alternatively, run a Agent History report.
6. Rerun the request to confirm that data is not being gathered.
7. Select **System Settings > Log Level** and set to **Info**.
8. Make a copy of the log for submission to EMC Support.

Client/Storage discovery for replication analysis troubleshooting

This section describes the steps that you can take to diagnose problems when trying to configure VNX Block/CLARiiON or Symmetrix storage arrays for replication analysis. We assume following scenario:

- DPA was successfully installed.
- The DPA server and the storage array host meet the requirements as specified in the *Data Protection Advisor Software Compatibility Guide*.
- Solutions Enabler has been successfully installed.

Client/storage discovery using remote execution

The following table describes possible problems and solutions encountered when attempting to discover clients or storage remotely; that is, without a DPA agent.

Table 53 Client/storage discovery problems and solutions

Problem	Solutions
<p>Client discovery fails: No authentication defined, or unable to log in.</p>	<ul style="list-style-type: none"> • Create a credential in DPA (Admin > System > Manage Credentials) and assign it to the client. • Check that the username and password supplied with the credential is able to connect the client. • Ensure su or sudo are not needed to connect, and if so that the correct parameters are supplied in the credential.
<p>Client Discovery fails: Failed to connect to client using RPC, or a specified logon session does not exist.</p>	<ul style="list-style-type: none"> • Check that the username and password supplied with the credential is able to connect the client. • Ensure you supplied the username along with the domain name: <code><domain> \<username ></code> for remote computers, <code><computer name>\<username></code> for local computers. In most cases <code>localhost \<username></code> can be used. • Check if the host is accessible from the DPA server using the admin share: <code>\ \hostname\Admin\$</code>

Table 53 Client/storage discovery problems and solutions (continued)

Problem	Solutions
	<ul style="list-style-type: none"> • If the error exists after you tried all the preceding actions, change the Log on as value for the DPA server service from local system to any other user with administrator privileges. A local administrator also can be set.
<p>Client Discovery failed: Failed to connect to client using RPC. The network path was not found.</p>	<ul style="list-style-type: none"> • Check if the client's Name, IP, or Alias is defined correctly and reachable from the DPA server. • Check if the host is accessible from the DPA server using the admin share: <code>\hostname\Admin\$</code> If the share is not accessible, check that it is not blocked by a firewall.
<p>Client Discovery failed: User does not have enough privileges to get the device mapping information.</p>	<ul style="list-style-type: none"> • Follow the system requirements for remote execution permissions. • Assign a user with administrator capabilities to the credential. • Check that the user that connects to the client has write and execute privileges to the path <code>/var/tmp</code>. (Unix)
<p>Client Discovery failed: Failed to send file for discovery to client using SCP or Failed to send file for discovery to client using FTP.</p>	<p>Check for free disk space in <code>/var/tmp</code>.</p>
<p>Client Discovery fails: Error (977). Overlapped IO operation in progress.</p>	<p>Verify that there is no antivirus software installed on the host; antivirus software might be blocking the <code>irxsvs.exe</code> operation. Disable antivirus blocking by authorizing the <code>irxsvs.exe</code> file in the anti-virus software.</p>
<p>Client discovery fails with the following error: <client_name> irx errMsg: Unable to connect host:<client_name> with user:<domain> \<username> using RPC irx output: Error (1203): No network provider accepted the given network path.</p>	<p>Ensure that the following services are running: Server, Computer Browser, and Workstation.</p>
<p>When using sudo, the Host Config request might fail to return volume group information on AIX hosts with the following message: SymMapVgShow exited with code 161 (SYMAPI_C_VG_NOT_AVAILABLE) SessionId: 0 - for VG:<vg_name> with type:</p>	<p>This occurs only when the credentials are configured to use sudo. Add the following line to the sudoers file: Defaults env_keep += "ODMDIR"</p>

Table 53 Client/storage discovery problems and solutions (continued)

Problem	Solutions
2(AIX LVM) VolumeGroup information will not be parsed.	

Client/storage discovery with agent

The following table describes possible problems and solutions encountered when attempting to discover clients or storage using the DPA agent.

Table 54 Client/storage discovery with agent problems and solutions

Problem	Solution
Client Discovery request is using remote execution instead of using the installed agent.	<p>Ensure that the agent is installed on the host.</p> <p>Ensure that the DPA server is defined as the controller for the agent.</p> <p>Restart the agent service.</p>

General client/storage discovery

The following table describes general possible problems and solutions encountered when attempting to discover clients or storage from DPA.

Table 55 General client/storage discovery problems and solutions

Problem	Solution
Client discovery finished with warnings: Failed to discover application storage objects for application <application_name> on client <client_name>.	<ul style="list-style-type: none"> Check if the application is running and available for connection. Check if the user configured in the DPA credential has enough privileges to query the application's system data.
Client discovery failed: can't connect to any IP.	<ul style="list-style-type: none"> Check that port 25011 or port 135 from the Windows Proxy Collector is not blocked by a firewall.
Client discovery finished with warnings: Home directory was not found for application.	<p>Select Admin > System > Manage Credentials.</p> <p>Click Edit to edit the Credential.</p>
Unsupported File System Type encountered: <filesystem_name>.	<p>DPA does not support this file system type.</p> <p>To avoid this warning in the next client discovery, you can ignore discovery for this file system.</p> <p>DPA will not display recovery data for this file system.</p>

Table 55 General client/storage discovery problems and solutions (continued)

Problem	Solution
<p>Client discovery failed with the error:</p> <p>Please verify that you have enough disk space and write permission.</p> <p>or</p> <p>Failed to unpack file on client <client_name>.</p>	<p>Ensure that there is enough disk space in the host's root file system according to the system requirements.</p>
<p>Client discovery finished with warnings:</p> <p>Can't find or no permission to execute file <home_dir>.</p>	<ul style="list-style-type: none"> • Check if <home_dir> that was discovered exists on the client. • Check if the file sqlplus in the home directory that was discovered has enough permissions for being executed by DPA.
<p>Client discovery finished with error:</p> <p>Timeout waiting for agent response on client <client_name>.</p>	<ol style="list-style-type: none"> 1. In DPA, select Admin > System > Configure System Settings and click Select Server. 2. Change the parameter Timeout(s) from the default 120 to a larger value. <p>Alternatively:</p> <p>Check if the DPA server has two network cards enabled and the client is able to reach both of them. If the client is not able to reach one of the cards, disable the network card that the client cannot reach.</p>
<p>When connecting to ECC 6.1, client discovery finished with error:</p> <p>Error Import Clients for w2k3-96-52.dm1nprlab.com finished with errors.</p> <p>Check previous error messages for further information.</p> <p>Unable to logon (Connection refused).</p>	<p>Run a batch file containing the following command:</p> <pre style="background-color: #f0f0f0; padding: 5px;"> %ECC_INSTALL_ROOT%\tools\JRE\Nt \latest\bin\java -cp %ECC_INSTALL_ROOT%\ECCAPIServer \class;%ECC_INSTALL_ROOT% \ECCAPIServer\ecc_inf\exec \eccapiclient.jar; com.emc.ecc.eccapi.client.util.EccA piPopulateRandomPassword ApiClient </pre> <p>The additional <code>classpath</code> parameter is only required if the command is not used from within the ECC classes directory.</p>
<p>Client discovery is not correctly resolving LUNs on VFMS on ESX 4.1.</p> <p>When trying to correlate the virtual devices to the remote storage they reside on, and the hosting ESX of the virtual machine cannot resolve the name of the VM (DNS configuration), the correlation fails and the virtual devices are displayed as local devices.</p>	<p>Configure the DNS on the ESX properly or add the VM name and IP to the ESX hosts file.</p>

Table 55 General client/storage discovery problems and solutions (continued)

Problem	Solution
Importing CLARiiON information request fails with the following error message: "An error occurred while data was being loaded from a Clariion ClarEventGet exited with code 3593 (SYMAPI_C_CLARIION_LOAD_ERROR) "	Run the <code>SYMCLI</code> command on the SE host for this CLARiiON: <code>symcfg sync -clar</code>
Host configuration request exceeds 60 minutes	Edit the <code>TIMEOUT</code> value to greater than 3600, for example to 7200, in the <code>services/remotex/deploy/<platform>/apolloreagent.ini</code> file, where <i><platform></i> the host on which the request exceeded, not the DPA server: <pre> <REARGS> <LOGFILE>apolloreagent.log</ LOGFILE> <LOGLEVEL>Info</LOGLEVEL> <WORKINGDIR>.</WORKINGDIR> <TIMEOUT>7200</TIMEOUT> </REARGS> </pre>

Incorrect Recovery Point times Synchronization

If there is a time difference between the DPA server and the storage array being monitored, recovery points might be displayed with times that do not match the expected times. For example, a system administrator initiates a recovery point at 2:00 but the recovery point is displayed at 4:00 in DPA.

The discovery requests have a Time Offset option that accounts for time differences and allows recovery points to be displayed with consistent times. You must calculate the exact offset between the DPA server and the storage array host.

In the following instructions, Solutions Enabler host refers to the DPA host which has the SYMAPI/CLARAPI Engine Discovery request assigned.

The time offset is calculated in seconds.

Synchronizing incorrect Recovery Point times on VNX/CLARiiON

To calculate the time offset between the VNX/CLARiiON and the DPA server:

Procedure

1. Query the VNX/CLARiiON's time using the `navicli` command `getsptime`.
2. At the same time, query the Solutions Enabler host's time.
3. If the Solutions Enabler host time and the DPA server time are equal (no Time Zone difference):

$$\text{TimeOffset} = \text{Solutions Enabler host time} - \text{VNX/CLARiiON time.}$$
4. Otherwise, if there is a time difference also between the Solutions Enabler host time and the DPA server time:

$$\text{TimeOffset} = (\text{DPA server time} - \text{Solutions Enabler host time}) - \text{VNX/Clariion's time.}$$

5. Set the time offset for the request. [Configuring the time offset](#) on page 248 provides information.

Synchronizing incorrect Recovery Point times on Symmetrix

To calculate the time offset between the Symmetrix and the DPA server:

About this task

Procedure

1. Query the DPA server time.
2. At the same time, query the Solutions Enabler host's time.
3. If the Solutions Enabler host time and DPA server time are not equal:
TimeOffset = DPA server time - Solutions Enabler host Time.
4. Otherwise, no time offset for Symmetrix needs to be set.
5. Set the time offset for the request. [Configuring the time offset](#) on page 248 provides information on this.

Configuring the time offset

After you have calculated the time offset, set the value for the request. To set the time offset value:

Procedure

1. Select **Inventory > Object Library > [select the Solutions Enabler host] > Data Collection**.
2. Select the relevant request and click **Edit**.
3. Set the time offset you calculated **Client-server Time Difference** in seconds or minutes.
4. Click **Apply**.

Troubleshooting report output failure

If reports are hanging after you save them with the message `Please wait while generating report`, and you are using Internet Explorer, it could be because you do not have the XMLHTTP option enabled. To enable the XMLHTTP option:

About this task

This is in relation to DCE-1546.

Procedure

1. Go to **Internet Options > Advanced**
2. Scroll to **Security** and select **Enable Native XMLHTTP Support**, then click **OK**.

Troubleshooting report generation or publishing problems

If scheduled reports fail to generate, or if they generate properly but fail to publish, perform the following actions:

About this task

- If a custom report, check that report template has been designed correctly in **Run Reports** area.

- Check that report template runs properly in **Run Reports** area.
- Check that report template properly saved (exported) in desired format.
- Check errors/warnings in `server.log` regarding scheduled reports.

If these actions do not resolve the issue, contact EMC Technical Support.

System clock synchronization

As part of the User Authentication process, DPA relies on the system clock times on the client machine and the server differing by less than one minute. In the event that clock times are unsynchronized, the following error message is displayed:

```
User Authentication failed due to the times on the client and server not matching. Ensure that the times are synchronized.
```

To resolve this issue, ensure that the system clock times on the client and server are synchronized.

You should use NTP to synchronize the DPA Server and all the DPA Agent hosts as well. This is imperative for accurate data collection.

