

INTRODUCTION TO XTREMIO METADATA-AWARE REPLICATION

Abstract

This white paper introduces XtremIO replication on X2 platforms. XtremIO replication leverages XtremIO's unique content addressable storage, in-memory metadata and copy data management implementation while supporting all data protection operations. We present XtremIO replication improvements over conventional replication offerings. We describe XtremIO asynchronous and synchronous replication and the capabilities of XtremIO Metadata-Aware Replication, including best RPO, performance, efficiency, change only replication, many Points-in-Time (PITs) at the Target, failover, full DR support, fan-in and fan-out support, and retention policy management.

February 2020

Contents

Abstract.....	1
Executive Summary.....	3
Conventional Replication Products.....	3
Understanding XtremIO Asynchronous Replication	4
XtremIO Replication Efficiency and Benefits	5
Deduplication and Compression	5
Changes Only.....	5
Write Folding	5
Asynchronous Replication Benefit Summary.....	5
Replication Flow	6
Understanding XtremIO Synchronous Replication	8
Replication Retention Policy	9
Accessing a Snapshot-Set at the Target	10
Refresh Capability	11
Testing a Snapshot-Set.....	12
Failover.....	12
Cleanup Flows.....	14
DR Host Volumes Accessibility	15
Capabilities	15
Comparison	16
iCDM Integration	17
Conclusion	18
XtremIO X2 Technical Whitepapers	18
How to Learn More	19

Executive Summary

Critical business processes require data replication for various purposes such as data protection or production environment duplication for development, testing, analytics or operations.

Data replication over the wire requires many customers to compromise either on RPO (Recovery Point Objective) or bandwidth cost for transferring the high throughput generated by the applications. This is even more acute with AFA (All Flash Array) high throughput.

While there are several approaches for replication, they all fundamentally struggle with RPO, limited performance and complex operational processes, forcing customers to make tradeoffs between RPO and cost.

XtremIO leverages its unique content addressable storage, in-memory metadata and copy data management implementation to offer the most efficient solution by replicating unique data only. This unique implementation significantly reduces bandwidth and storage array utilization. Furthermore, XtremIO simplifies the replication management for all use cases.

As with all other XtremIO data services, all data protection operations, such as Failover, Failback, etc., are supported and their management is very intuitive.

The integration of replication and XtremIO iCDM capabilities simplifies the repurposing processes and decreases the time for creating and refreshing environments from hours and days to just minutes.

XtremIO also enables instantaneous test and failover operations.

Conventional Replication Products

Conventional replication products use different mechanisms that are streaming or snapshot based. When sending the data over the wire, the bandwidth requirements for such solutions require the customer to either compromise on RPO or increase the bandwidth cost to allow copying all the changes.

The reason for this is that with conventional solutions, all changes need to be replicated. For example, in [Figure 1](#) data is replicated from Array 1 in primary site to Array 2 in DR (Disaster Recovery) site. When a new, non-unique block "A" is written to address "4" in Array 1, it will be fully replicated to the remote array even though the same data already exists in address location "1" as shown in [Figure 2](#).



Figure 1. Conventional Replication: Non-Unique Block "A" Added to Array 1




Figure 2. Conventional Replication: Non-Unique Block "A" Replicated in Array 2

This means that all data changes are fully replicated on the wire, causing the following:

- The bandwidth needs to be sized according to all changed data.
- The performance of the arrays on both sides are impacted by replicating all data changes. The source array reads the data and transfers it to the remote array, and the remote array receives all the changed data and writes it to the drives.

Understanding XtremIO Asynchronous Replication

XtremIO replication leverages a unique content-based storage architecture to reduce bandwidth utilization. XtremIO stores only unique data at the array physical layer, and manages volume information at the logical layer, also called in-memory metadata. For more information, see the [Introduction to Dell EMC XtremIO X2 Storage Array](#) white paper.

Every data block that is written in XtremIO is identified by a fingerprint that is kept in the data block's metadata information. When the fingerprint is unique, the data block is physically written and the metadata points to the physical block. When the fingerprint is not unique, it is kept in the metadata and points to an existing physical block.

A non-unique (deduplicated) data block that already exists on the target array is not sent again. Instead, only the block metadata is replicated and updated at the target array.

For example, in [Figure 3](#) a new block "D" is added to address 6. Block "D" already exists at address "5" and was already replicated to the array at the DR site.




Figure 3. Non-Unique Block "D" Written to Address "6"

When a non-unique block is written to the source array, XtremIO replication will not replicate the data, but rather will update the metadata at the target array to point to the physical block that already exists on the target array.

In [Figure 4](#), the metadata at the DR array is modified to point to the unique block "D" that is already stored at the physical layer of the DR array. Instead of needing to replicate the full data block, XtremIO efficiently replicates only the metadata.




Figure 4. Address "6" Modified to Point to Block "D"

XtremIO Replication Efficiency and Benefits

The following sections describe unique XtremIO capabilities which provide replication benefits and improved efficiency.

Deduplication and Compression

For every changed data block, XtremIO replication checks if its fingerprint exists at the target array. If the fingerprint already exists, only the metadata is updated at the target and no data is sent.

In case of a new unique block with no previous fingerprint, the source replicates the full compressed data block.

This efficient replication is not limited per volume, per replication session or per single source array, but is a global deduplication technology across all volumes and all source arrays.

XtremIO deduplication is inline, always-on, and is not sensitive to the cluster's utilization, thus making it the most efficient replication technology in the market.

Changes Only

XtremIO Asynchronous replication is based on a snapshot shipping method. Snapshots at the source are created at a frequency derived from the RPO setting of the protection session (called Cycle). Snapshots are efficiently transferred (shipped) to the target site. In each cycle, a new snapshot is created and XtremIO calculates the changes between the last 2 snapshot-sets at the source, and transfers the changes to the target array where it is merged with previous data. This mechanism essentially means that only changes between cycles are transferred to target, thus providing additional operational efficiency.

Write Folding

In many cases, applications repeat writes to specific addresses of the volumes. These addresses are called "hot spots". In streaming based replication technology, when the replication granularity is per a single IO, all write I/Os need to be replicated individually and in the same order even if they are written to the same address. However, in XtremIO snap-based replication, for every snapshot, only the last write I/O to every address determines the data that needs to be replicated. In this way, write I/Os that were overwritten are not replicated. This is called "Write folding", and it reduces the amount of data that needs to be replicated.

The greater the time between replication snapshots, the larger the savings which can be obtained from Write Folding.

Asynchronous Replication Benefit Summary

Based on all the above capabilities, XtremIO replication provides the following benefits:

- The bandwidth required on the wire is based only on the unique changed data.
- Data is sent compressed over the wire.
- Only unique data blocks, at either the source or target array, need to be replicated.
- Only data block changes between subsequent cycles, and not the complete block, need to be replicated.

For a typical DRR of 4:1, only quarter of the bandwidth needs to be sent over the wire, thus the bandwidth savings is 75%.

Replication Flow

XtremIO Asynchronous replication uses snapshot-shipping method to replicate crash-consistent copies to the remote array.

On the source cluster, XtremIO capabilities are used to create crash-consistent virtual copies, and to calculate the difference between two consecutive virtual copies. The copy creation and diff commands are performed fully in-memory and therefore are done at memory speed.

The replication flow is as follows:

1. When replication starts, a full copy is performed to ensure that the production and the target are identical. To perform the full copy, XtremIO replicates the first snapshot-set to the target cluster. The data is sent efficiently, and only unique blocks are sent (compressed) over the wire. This initialization phase is shown in [Figure 5](#):

Step 1 – First snapshot created.

Step 2 – Only unique blocks sent compressed.

Step 3 – Snapshot stored at target.



Figure 5. Copy Initialization Phase

2. For every subsequent cycle, a new snapshot-set is created at the source and only the differences between the previous snapshot-set and the current snapshot-set are sent to the target cluster. The difference between the snapshot-sets reflects all the blocks that were changed between the 2 cycles and that need to be replicated. The unique data is sent compressed.

Each new snapshot-set is consistent across all volumes in the consistency group. XtremIO compares the metadata of the current snapshot-set to the metadata of the previous snapshot-set to determine the differences. It then gets the fingerprints of the changed blocks and checks if the same fingerprint exists at the target array. If the fingerprint exists, the metadata is updated at the target and in case not, the data is read and sent compressed to the target array.

3. On the XtremIO target array, a new snapshot-set is created to which the incoming data is written. When replication of the snapshot-set delta is completed, the snapshot-set that is received from source becomes the latest snapshot at the target array and is available for the DR host. When the snapshot-set becomes available to the DR host, the previous snapshot-set at the source array can be deleted. Figure 6 shows the steps in these subsequent phases:

Step 1 – A new snapshot is created.

Step 2 – Changes between the current and previous snapshot are calculated.

Step 3 – Only unique blocks are sent compressed.

Step 4 – The new snapshot is merged and stored at the target.




Figure 6. Subsequent Copy Phase

4. The trigger for starting a new cycle depends on the RPO settings specified by the user for the protection session. The RPO in XtremIO can be as low as 30 seconds and up to 24 hours. XtremIO will attempt to meet the required RPO and will start a new cycle half way through the desired RPO. For instance, if the RPO is set to one minute, XtremIO will start a cycle every 30 seconds to ensure compliance with the replication RPO.

In case there are many changes or packet losses on the wire, which cause a new time interval to pass before the previous cycle has finished, the new cycle will start on completion of the previous one. A longer interval results in more efficient data transfer, because of Write-folding, where a single, final write replaces multiple writes to the same location in the same cycle at the source volume.

5. XtremIO automates the cycle of creating a new snapshot-set at the source, sending the changes to the target array, creating a new snapshot-set at the target, and deleting the snapshot-set from the source. The process is repeated indefinitely.

Understanding XtremIO Synchronous Replication

Synchronous Replication is suitable for applications that require zero data loss in the event of a disaster.

To achieve zero data loss, Synchronous Replication ensures that each block of data written by the host is saved to both the source volume locally and to the target volume remotely, before the write is acknowledged to the host.

Figure 7 describes the flow of Synchronous Replication.

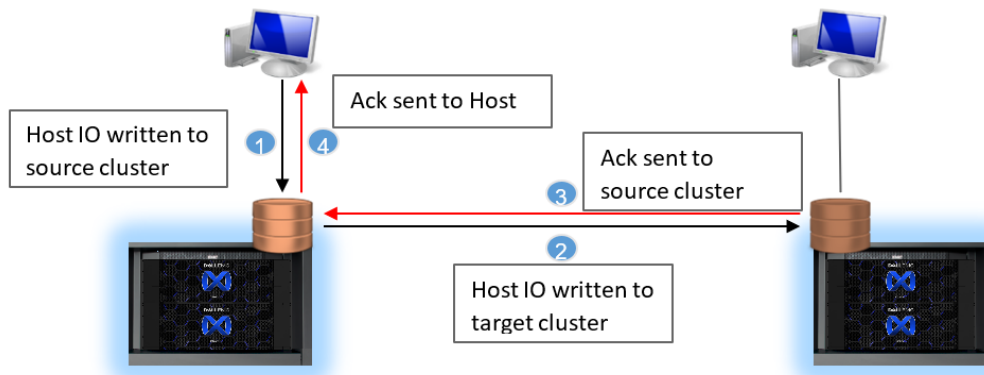


Figure 7. I/O Flow with Synchronous Replication

1. The host initiates a write I/O to the source cluster.
2. The source cluster sends the I/O to the target cluster.
3. After the I/O is written at the target, the target cluster sends an acknowledgement to the source cluster.
4. When the I/O is written to both clusters (source and target), the source cluster sends an acknowledgement to the host.

As each write needs to be saved both locally and remotely, the trade-off of synchronous replication is the added latency caused by the remote write as well as by the link latency. This response time increases as distance between the remote clusters increases. Synchronous Replication has a distance limitation based on latency between systems. This limitation is generally 60 miles or 100 km between sites.

Synchronous Replication presents a challenge when the source and target get out of sync. This can happen during the initial sync phase, as well as when disconnection occurs, due to link failure or a user operation (e.g. pausing the replication or performing failover).

The unique capabilities of XtremIO's Synchronous Replication high efficiency include:

- **Metadata Aware Replication** – During the initial synchronization and when the target becomes out of sync, the replication uses the metadata aware replication to efficiently and quickly replicate the data to the target. The replication uses multiple cycles until the gap is minimal and then switches to synchronous replication, thus reducing the impact on the production to minimum, and accelerating the sync time.
- **Recovery Snapshots** – To avoid a full-copy or even a full-metadata-copy, XtremIO leverages the in-memory snapshots capabilities. Every few minutes recovery-snapshots are created on both sides, which can be used as a base-line in case a disconnection occurs. When connection is restored, the system only needs to replicate the changes that occurred from the point the recovery snapshots were taken.
- **Prioritization** – To ensure the best performance for the applications using Sync Replication, XtremIO automatically prioritizes the I/O of Sync Replication over Async Replication. This prioritization occurs automatically, and no tuning or special definition is required.
- **Auto-Recovery from Link Disconnection** – Replication resumes automatically when the link is back to normal.

XtremIO Synchronous Replication is managed at the same location as Asynchronous Replication and supports all Disaster Recovery operations, as is supported for Asynchronous Replication.

Switching between Async to Sync is as simple as a single command.

XtremIO Synchronous Replication is fully integrated with Asynchronous Replication, in-memory snapshots and iCDM capabilities, making the replication, recovery and repurposing very efficient.

Replication Retention Policy

The protection window and the maximum number of snapshot-sets that XtremIO keeps are determined by the Retention policy. The maximum snapshot-sets per protection session is 500 (for the initial version, please check the RN for the latest scalability numbers).

The retention policies are defined once and can then be used by multiple protection sessions.

The retention policy is applicable for both Synchronous and Asynchronous replication.

The retention policy specifies:

- The Required Protection Window
- The number of snapshot-sets that will be kept within the protection window

For example, for a Production Consistency Group, a user can create a "Gold" policy, with a short protection window of 30 snapshots for 60 minutes, a medium protection period of 23 snapshots for 23 hours, and a long protection period of 2 snapshots for 2 days. This essentially means your production consistency group has a PIT every 2 minutes for the first 60-minute period, a PIT every 1 hour for the next 23 hours, and a PIT per day after that.

A Silver policy, for Test/Dev, may use a single long period window of 1 snapshot per day for a week.

The screenshot shows a web interface for creating a new retention policy. The title is "New Retention Policy". There is a text input field for "Retention Policy Name" containing the word "Gold". Below this is a section for "Retention Policy Periods" with a "+ ADD PERIOD" button. Three periods are listed: "Short" with "KEEP 30" and "SNAPSHOTS FOR 60 Minutes", "Medium" with "KEEP 23" and "SNAPSHOTS FOR 23 Hours", and "Long" with "KEEP 2" and "SNAPSHOTS FOR 2 Days". Each period has a small summary below it: "1 Snapshot every 2 Minutes", "1 Snapshot every 1 Hour", and "1 Snapshot every 1 Day". At the bottom right are "CANCEL" and "APPLY" buttons.

Figure 8. Retention Policy Configuration

The Required Protection Window is specified by the user; from a short period of minutes to a long period of up to one year.

In addition, with XtremIO the protection window can be split to 1 to 3 time-periods with different number of snapshot-sets and different granularity.

XtremIO automatically manages the retention of the snapshot-sets according to the retention policy settings. In case the protection window settings or the number of snapshot-sets is changed, XtremIO will automatically adhere to the new retention policy settings.

The retention of snapshot-sets runs:

- As Retention policy is executed every 5 seconds.
- At the end of a replication cycle.
- Whenever the retention policy settings are updated.

When the protection session is suspended, the retention policy is suspended as well, and the system will not delete any snapshot-sets. This will allow the user to test the snapshot-sets at the target when recovery is required, and prevents any snapshot-set deletion until a decision is made regarding the use of the snapshot-sets for recovery.

Accessing a Snapshot-Set at the Target

Accessing the snapshot-sets at the target is required for the following use cases:

- Permanent:
 - When near-instantaneous disaster recovery or testing are required
 - Repurposing copies and environments
- Temporary:
 - Technology refreshes
 - System upgrades
 - Data center moves or expansion
 - Production data migration to a target array

When replicating to the remote array, the volumes at the remote site must be defined as read-only or no- access mode depending on user preference for the host access on that site. In the event of a failover, the volumes at the remote site are changed to read-write.

Refresh Capability

XtremIO provides the option to refresh the data of a host volume from any PIT (Point in Time) that exists in the system. The refresh can be done from any PIT snapshot-set that was created for protection. The refresh operation is instantaneous, as the refresh operation is a metadata operation that is done completely in-memory. The refresh operation doesn't require any metadata copy or roll-forward/backward of any data or metadata. The refresh doesn't change the SCSI personality of the volume and therefore is transparent to the host. Thus, there is no need to perform any mapping or scanning operation for the data to be accessible to the host.

The refresh operation performs the following steps:

1. Create a virtual copy from the snapshot-set used for the refresh.
2. Change the host volume to point to the newly created virtual copy, as shown in [Figure 9](#).




Figure 9. Host Volume Points to Data Container Including Metadata Information

When the data of the host volume needs to be restored or refreshed from a virtual copy, the refresh operation will just update the host volume to point to the new virtual copy data, as shown in [Figure 10](#).




Figure 10. Refresh Operation

Testing a Snapshot-Set

Testing a remote replica is required to validate that all data is fully and correctly replicated. Testing a snapshot-set at the target arrays is instantaneous, using a single command. All that is required is that the snapshot-set which contains the desired point in time be promoted to be the volume of the DR-host. While the host accesses the new PIT replica, snapshot-sets continues as usual, adhering to the RPO.

Testing the snapshot-set at the remote host is not limited by time and can therefore be used indefinitely. This is very useful when extensive testing of a replication snapshot-set is required.

During a test scenario, the volumes at the remote site are switched from No Access/Read-only to Read-Write. However, when testing of a snapshot-set is completed, the access-mode of the target volumes changes back to read-only, and all the data written to the target volumes are discarded. To keep the data of a tested copy, a protection-copy of the target needs to be taken prior to the finish-test-copy command.

Any target snapshot-set can be used for testing. When the user selects a snapshot-set, XtremIO provides access to a copy of the selected snapshot-set, thus the original protection copy is not affected.

The test-copy operation is performed by using the "Refresh" capability as described above.

Before finishing the test copy, host applications accessing the copy must be shut down and the file-system must be unmounted. The test-copy-finish command removes write-access from target volumes to prevent data corruption on the next test-copy or failover operations.

Failover

When performing failover, the replication direction is reversed. The target Consistency Group and target volumes become the production and the production volumes of the consistency group become the target volumes. As part of the failover command there is an option to start the replication in the reverse direction immediately after switching sides. When the replication to the new target (original production) is started, a full check is performed which will check the whole volume. But replications will only be made of the differences between the source and the target volumes. These differences are calculated by matching the fingerprints between the source and target snapshot-sets.

After failover, the previous snapshot-sets are deleted. To keep the snapshot-sets, create a protection copy where needed.

Failover to the target site can be done on demand with or without first using a "Test snapshot-set" operation. When a snapshot-set is mounted using the Test snapshot-set option, the user can run some recovery procedures on the host. Upon completion, the user can decide whether or not to use the current mounted data for the failover.

XtremIO supports all failover scenarios including "Planned Migration" (sync-and-failover) and "Disaster Recovery" to any PIT at the target. The failover operation for "Disaster Recovery" scenarios with XtremIO is instantaneous and requires no metadata copy, log roll-forward or log roll-backward.

The following diagrams represent the failover process.

1. [Figure 11](#) shows the initial state at the DR site.
 - Volumes are mapped and visible to the DR host.
 - Write-access is disabled to the DR host volumes.
 - Multiple Point-in-Time Copies exist on the target.




Figure 11. Initial State at DR Site

2. When a failover needs to be performed, select the PIT copy for failover, and execute the failover command, as shown in [Figure 12](#).




Figure 12. PIT Selection

- The failover command creates a copy from the selected PIT copy and uses the "Refresh" capability to change the host volume to point to the new created copy as shown in [Figure 13](#).




Figure 13. Failover to Selected Copy

Because promoting a snapshot-set to the target host is essentially instantaneous, near-zero recovery time objective (RTO) is possible when failing over, regardless of the selected point-in-time.

Cleanup Flows

At the end of a recovery incident, a cleanup process should be performed to delete unnecessary snapshots. Key aspects of the cleanup are as follows:

- Removing a protection session removes all protection session snapshot-sets from both the source and target clusters.
- Terminating a protection session removes the corresponding auto-created snapshot-set from the source cluster. User-created snapshot-sets (protection copies that were taken on the replication snapshot-set) will not be removed. Terminating the replication session will perform a metadata-aware full-copy when the protection session will be restarted. After the first replication, when a full sweep is needed, there is a good chance that the copy at the target is similar to the production so most of the data does not need to be transferred.
- Removing snapshot-sets from the target allows clean-up of the protection window at the target cluster.
- On Failover, when the source cluster is accessible, failing over to the target cluster triggers a swap. The source cluster volumes change to read-only/no access mode, and the cluster switches identity from source cluster to target cluster. Similarly, the volumes at the target cluster change to Read-Write, and change identity from target cluster to source. When the source cluster is inaccessible, it is possible to failover all applications to the target cluster (thus making it the source cluster) but the original source cluster cannot be changed to target. This creates a "split brain" situation, in which there are two source clusters. In such a situation, some applications may still write to the original source cluster even after the target cluster has become the new source. When connection between the clusters is restored, the user needs to determine which of the clusters will function as the source (i.e. which cluster data should be used) before initiating replication. The data of the target volumes will be discarded. To preserve the data on the new target, use the create-protection-copy command on the target prior to executing the failover-cleanup command.

DR Host Volumes Accessibility

With XtremIO, the target volumes are inaccessible by default. The access mode for the target volumes can be defined in the protection session. There are 2 access modes for the target volumes: "read-only" and "no-access". The default value is "read-only".

The target-volume access mode is changed to read-write in the following scenarios:

- When testing a snapshot set at the target host
- When a failover is performed
- When the volumes are removed from the protection session

Capabilities

XtremIO Metadata-Aware Replication provides the following capabilities:

- **Best RPO** – As low as 30 seconds.
- **Performance** – Support for All Flash Array high performance workloads.
- **Efficiency** – Only unique data is sent over the wire. The data is sent compressed.
- **Changes Only** – Only changes between cycles are replicated. Resuming the replication is always incremental and is not impacted by communication failures.
- **Many PITs at the Target** – Hundreds of PITs can be kept at the target and can be used for failover or repurposing.
- **Best RTO** – Failover or test a PIT at the target is instantaneous and does not require any roll forward or metadata copy. DR Host can be easily recovered, while preserving all SCSI information (thus eliminating the need for SCSI-BUS rescan on the host side).
- **Full Support of DR Operations** – Such as test copy to any PIT, failover, and failback.
- **Bi-Directional** – XtremIO supports bi-directional replication, replicating one CG from Array "A" to Array "B" and a different CG from Array "B" to Array "A".
- **Fan-In Support** – Replicating from multiple clusters to a single target cluster benefits from global dedupe with XtremIO replication.
- **Fan-Out Support** – Replicating different CGs to different remote clusters.
- **Retention Policy Management** – The system automatically manages the PITs at the target according to the Retention Policy.

Comparison

Table 1 compares some of the key differences between the various replication technologies and the XtremIO Metadata-Aware Replication.

Table 1. Existing Replication Technologies vs. XtremIO Metadata-Aware Replication

	Other Snapshot-based Replication Products	Streaming Based Replication	XtremIO Metadata-Aware Replication
Space-Efficient Data on Source and Target Arrays	No, or limited	No, or limited	Yes + Inline and Always On
Transfer Efficiency – Compression	May be supported, usually limited and has performance impact	When supported, usually limited and has performance impact	Yes + Inline and always On
Transfer Efficiency – Deduplication	No, or limited	No, or limited	Yes + Inline and always On
Transfer Efficiency – IO Folding	Yes	No, or limited	Yes
Sensitive to communication failures – full sweep	No	Usually, a resync is needed when communication fails for long time	No
Test Copy Limitation	No	Usually has impact when used for long time	No
Data Services Limitation	May limit data services	May limit data services	No
Instant Failover/Test to any copy at the target	Yes, but usually requires data or metadata copy	No or limited	Yes
Addresses DR scenarios	Yes	Yes	Yes
Addresses Logical Corruption Scenarios	Yes	Very Limited	Yes
Repurposing 1 or many environments from remote copies	Difficult and time consuming	Difficult and time consuming	Yes
Performance	Yes, but usually not dedupe aware	Not optimized and usually not dedupe aware	Yes
Instant RTO	Requires roll forward or metadata copy	Yes	Yes

iCDM Integration

XtremIO replication is fully integrated with XtremIO iCDM capabilities, which allow quick setup and refresh of development & test environments. The ability to quickly refresh an environment from a remote production cluster with built-in virtual copy capabilities, allows setup of many environments with small capacity impact.

Creating a virtual copy with XtremIO is instantaneous and doesn't have any impact on the capacity. Only the changes that will be made on the copy will consume space.

Multiple repurposing environments are supported and can be easily setup, for instance:




Figure 14. Replicating Production Environment to DEV Environment

In this example, different clusters exist for the production and DEV environments. The production environment is replicated to the DEV cluster (cluster2) with a single command. When created, the DEV environment has the following benefits:

- It does not consume any additional metadata or physical space. Only changes will consume space thus providing huge capacity savings.
- It will benefit from the same performance as a regular volume.
- It will benefit from the same data services as a regular volume.

Refreshing an environment with updated data is also instantaneous. The user can refresh from any replication copy newer or older than the data currently used on his environment.

In cases where sensitive data from the production system must be masked before replicating to the DEV environment, using XtremIO iCDM and replication capabilities will speed up the overall process and make it very easy to manage.

With XtremIO the master copy will be refreshed using XtremIO iCDM "Refresh" option. The replication will replicate the masked image to the remote cluster. And with the iCDM repurpose and refresh option the remote environments will be easily refreshed.




Figure 15. Repurposing Remote DEV Environment from Master Image

Conclusion

The XtremIO Replication data service offers the most efficient replication technology and supports all Flash Array high performance workloads and enterprise level protection requirements.

XtremIO supports hundreds of PITs that can be used for failover, testing and repurposing. Accessing the PITs at the target for testing and failover is instantaneous and does not require any metadata copy or log roll-forward operations.

XtremIO replication is fully integrated with XtremIO iCDM capabilities and creating new Development/Test environment from the replication copies is very easy and efficient. The time needed to create and refresh remote environments will drop from days or hours with existing technologies to minutes with XtremIO replication and iCDM capabilities.

XtremIO X2 Technical Whitepapers

- XtremIO X2 Technical Architecture – <http://www.emc.com/collateral/white-papers/h16444-introduction-xtremio-x2-storage-array-wp.pdf>
- XtremIO X2 High Availability Architecture – <http://www.emc.com/collateral/white-papers/h16445-high%20availability-x2-architecture-wp.pdf>
- XtremIO X2 XDP Data Protection – <http://www.emc.com/collateral/white-papers/h16446-xtremio-x2-xdp-wp.pdf>
- XtremIO X2 Data at Rest Encryption – <http://www.emc.com/collateral/white-papers/h16447-xtremio-x2-dar-encryption-wp.pdf>

How to Learn More

For a detailed presentation explaining XtremIO X2 Storage Array's capabilities and how XtremIO X2 substantially improves performance, operational efficiency, ease-of-use and total cost of ownership, please contact XtremIO X2 at XtremIO@emc.com. We will schedule a private briefing in person or via a web meeting. XtremIO X2 provides benefits in many environments and mixed workload consolidations, including virtual server, cloud, virtual desktop, database, analytics and business applications.



[Learn more](#) about Dell EMC XtremIO



[Contact](#) a Dell EMC Expert



[View more](#) resources



Join the conversation
[@DellEMCStorage](#) and
[#XtremIO](#)