Datrium

Datrium DVX
System Management

Version 5.3.1.0

VMware, Inc.

3401 Hillview Ave.

Palo Alto, CA 94304

www.vmware.com

Contents

DVX System Management

The Automatrix platform includes the **DVX System** a self-protecting, scaleable data storage and protection infrastructure, available either as software -defined converged infrastructure or as the Cloud DVX service on public clouds. And ControlShift, a global data management SaaS application providing workflow automation for mobility and disaster recovery.

Automatrix Platform

Automatrix integrates compute, storage, virtualization, backup, and recovery. Datrium software runs on the following platforms:

- Compute Nodes for I/O processing, using local compute and flash resources.
- On-premises network-attached Data Node to provide a durable storage Data Pool.
- Remote Compute Nodes and Data Pools for on-premises private cloud storage.
- Cloud DVX storage as an AWS service for off-site cloud storage.
- ControlShift DR orchestration of workload mobility between on-premises systems and between on-premises systems and Cloud DVX.

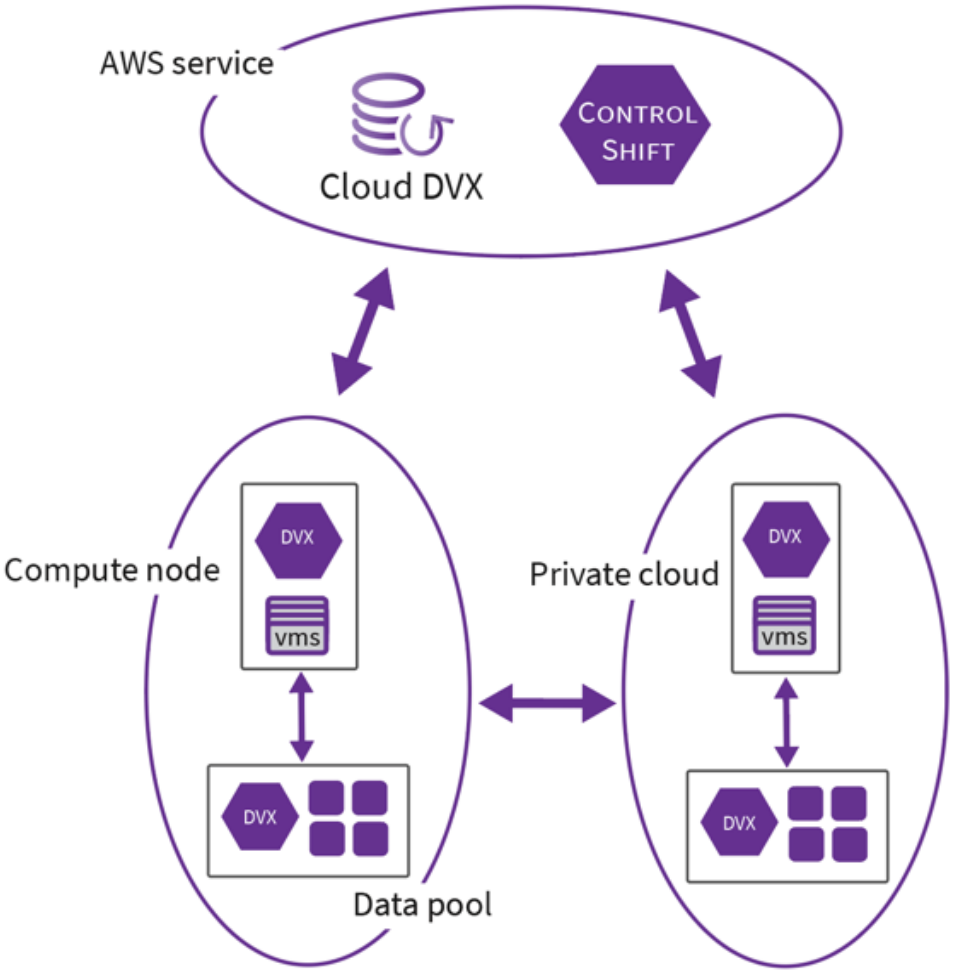

The DVX System provides advanced copy data management operations that applies the automatic data reduction features of deduplication and compression across unified platforms for always-on efficiency. To maintain data protection, the copy data management suite of features also includes always-on erasure coding, continual integrity checks with Blanket Verification, and optional end-to-end Blanket Encryption. A DVX System converges compute, primary storage, backup, and cloud-based disaster recovery to provide the following storage capabilities:

- **High performance** – Maximum speed is achieved by putting all active data in flash on Compute Nodes to support mixed workloads that require low latency. The DVX System uses local flash as primary storage supporting read requests. Performance scales linearly as you add Compute Nodes.

- **Scale-out backup** – Durable storage capacity provides a built-in backup and it scales separately from performance by adding Data Nodes to the Data Pool. The expanded Data Pool increases write bandwidth and disk rebuild speed.
- **Cloud-based DR** – Cloud-based disaster recovery maintains the DVX copy data management activities across on-premises sites for private cloud storage and data storage in an off-site AWS cloud. Economy of scale is achieved by applying global deduplication to data in on-premises and offsite Data Pools. Disaster recovery operations support the transfer of workloads from on-premises DVX protected sites to on-premises DVX recovery sites, and can use the Datrium cloud storage in an AWS Cloud as the source of data for recovery.

High Performance

The DVX System uses CPU, memory, and local flash on a Compute Node to achieve high performance I/O processing. It can service reads from data in local RAM, from the data stored in local flash for very high speed, or from the Data Pool if the request cannot be resolved locally. Writes are synchronously replicated to mirrored NVRAM on a Data Node for low latency and data durability.

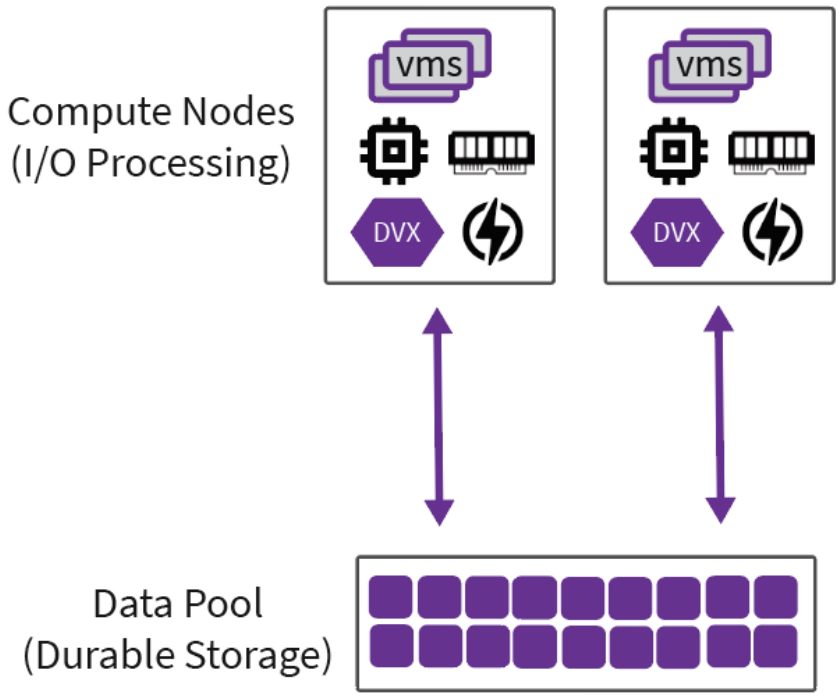

A Compute Node is stateless – the DVX System ensures that data is saved in the Data Pool before additional I/O requests are processed and before it performs any further processing on the data. Failure events on the Compute Nodes do not affect the data stored in the Data Pool and do not require data to be rebuilt.

For more information, see DVX Architecture | Datrium Technical Report.

Compute Nodes

A Compute Node is an x86 server that has the DVX Hyperdriver installed on it and it has flash memory for DVX use. The Compute Node uses the Hyperdriver for local I/O and for access to the virtual machine files on the Data Node.

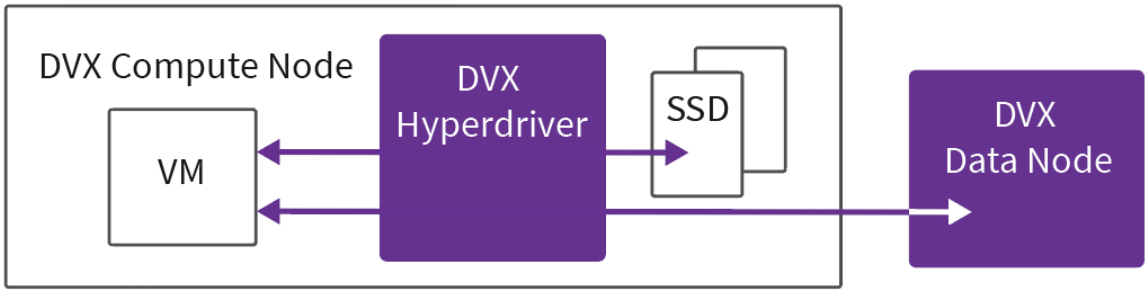

A DVX Compute Node has a minimum of two SSDs that are dedicated for DVX System use.

The DVX System can manage a heterogeneous environment that includes DVX Compute Nodes and third-party x86 servers running DVX Software. Datrium provides DVX Compute Nodes with VMware ESXi. You can also create a Compute Node with KVM (Kernel-based Virtual Machine) installed on CentOS-7 or with Red Hat Virtualization.

The DVX Hyperdriver

The DVX Hyperdriver is software that is installed on the Compute Node or third-party server in all cases. It provides optimized DVX file system access on the local Compute Node and through a tight coupling with the DVX Data Pool. The Hyperdriver uses flash, RAM and compute resources on the Compute Node to accelerate storage operations and provide data services such as compression, deduplication, erasure coding, encryption, snapshots and clones, replication, RAID, and space reclamation for data in the local host and in the Data Pool. Hyperdriver software on different Compute Nodes operates independently for the most part, with some transient exceptions. The Compute Nodes are coordinated by software on the Data Node(s).

The Hyperdriver provides the following capabilities:

- Presents a DVX Datastore's NFS server interface to the host system. This NFS interface is virtual; the protocol terminates within the host at the Hyperdriver. The Hyperdriver uses a proprietary network protocol to communicate with the Data Node.
- Uses host compute resources to process I/O requests. For most operations, this is for local processing only. In some cases, for example in disk rebuilds and space

reclamation, the storage compute load is distributed across the hosts to minimize impact to a particular host and to accelerate the results. As more Compute Nodes and Data Nodes are added to the same DVX System, load per host decreases and the speed of these global operations increases.

- Uses local host SSD space to fulfill the vast majority of virtual machine I/O requests.
  - Dispatches write data immediately to Data Node NVRAM for durability. It also writes the data to the Compute Node SSD.
  - Uses the Compute Node for fast local reads. Data will stay in the SSD until the SSD fills up. When SSDs are sized to fully accommodate all powered-on virtual machines,
    read requests are satisfied from local flash without ever issuing I/O requests to the network.
  - Uses the Compute Node RAM and SSD to coalesce data for a specific virtual machine into a larger, compressed whole stripe that is erasure coded and dispatched to the Data Pool after the whole stripe is collected.
  - Because of the DVX System end-to-end data reduction capability and the size of modern SSDs, Compute Node reads from the Data Node are typically limited to requests during cold snapshot boots. If the Compute Node SSD is undersized, its behavior will degenerate in caching, evicting on the basis of least-recent use.
- Provides VM fault tolerance through Peer Cache Mode. In the case of the failure of all host SSD drives or the host's RAID controller, all virtual machines on the affected host stay online. Peer Cache allows virtual machines on the affected host to use the SSDs on a neighboring host until the SSD or RAID controller failure can be addressed.
- Provides snapshot, clone, and replication operations for copy data management at a virtual machine, vDisk, or Docker Persistent Volume granularity, visible in the Snapstore.
- Supports Microsoft Windows VSS (Volume Shadow Copy Service) for application quiescing during DVX Snapshot operations.
- Provides Blanket Encryption for data in use, at rest, and in transit. The DVX System supports FIPS 140-2 approved or FIPS 140-2 validated modes of operation. The DVX System software uses a cryptographic module that is FIPS 140-2 certified.
- Supports virtual machine backup and recovery operations through vSphere's VADP (vStorage APIs for Data Protection) and its support of ESXi NFS-attached storage.
- Supports VAAI (vSphere Storage APIs - Array Integration) storage operations.

Scale-Out Backup

A Data Pool is a set of DVX Data Nodes that provide a durable data store for Compute Nodes. DVX durable storage capacity provides a built-in backup for the primary storage on Compute Node flash and Data Node NVRAM that supports virtual machine I/O.

In the DVX System, the separation of active data from durable copies supports Split Provisioning, in which resources are provisioned and scaled independently. The durable storage scales separately from performance by adding Data Nodes to the Data Pool. The expanded Data Pool supports additional Compute Nodes and it increases write bandwidth and disk rebuild speed.

To increase speed:

- Add Compute Nodes to the DVX System to increase compute resources and balance the virtual machine load.
- Add flash to a Compute Node to improve I/O performance. With sufficient flash, all read requests can be satisfied locally.

To increase capacity:

- Add a Data Node to the DVX System to increase total storage capacity, aggregate write bandwidth to the Data Pool, and disk rebuild speed.
- A single Data Node can support up to 1500 powered-on virtual machines distributed across 32 hosts.
- A DVX Data Pool consists of a set of 1 to 10 Data Nodes.

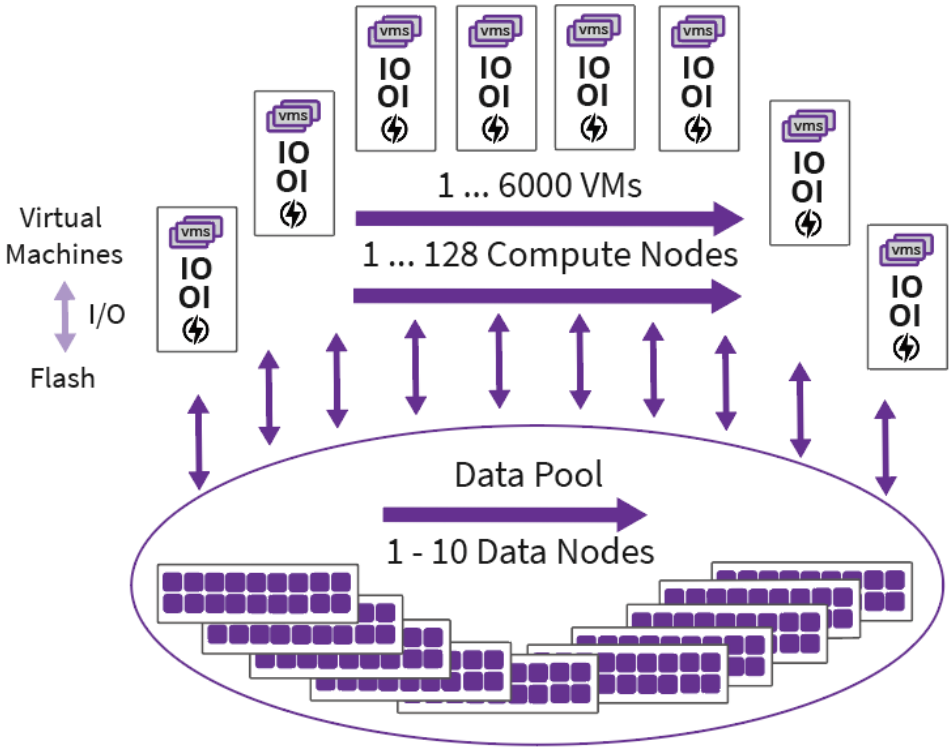

Depending on the Data Node model, different levels or usable capacity are supported:

- Data Node models D12x4, D12x4B, D12x4C provide a usable capacity of **29TB**(**26.4TiB**).
- Data Node models D12x10D provide a usable capacity of **73TB** (**66TiB**).
- F24x2B and F24x2D provide a usable capacity of **28TB** (**25.4TiB**

Each additional Data Node, up to the maximum total of 10, increases the usable raw storage capacity by 26TiB. Using a ten-node Data Pool, the DVX System can store over a petabyte of deduplicated and compressed data.

The Compute Node and virtual machine maximums can be achieved in a 4-node Data Pool, which can support up to 128 Compute Nodes and up to 6000 powered-on virtual machines.

DVX Datastores

The virtual machine files on a DVX Data Pool are visible as a hypervisor-mountable namespace called a Datastore. A DVX Datastore provides an NFS mount point for a host. A

DVX Data Pool can contain one to a maximum of 32 Datastores. Any of the Datastores can use the total available Data Pool capacity. You can limit the set of hosts that can access a specific Datastore.

DVX Snapstore

The DVX Snapstore is a separate namespace that contains snapshots and Protection Groups for policies regarding virtual machines and their snapshot scheduling, retention, deletion, and replication. This does not increase the size of data in the Data Pool; the Snapstore is a namespace of references and relationship metadata. It is not visible to hypervisors.

The DVX System creates snapshots of all objects within a Protection Group at the same point in time to ensure state consistency. The DVX System supports two types of snapshots:

- A crash-consistent view of the virtual machine data at a single point in time. When you take a protection group snapshot, the DVX System pauses all I/O to all virtual machines and files in the protection group at the same I/O instant, even if the virtual machines are on different hosts. Protection group contents are snapped simultaneously. A DVX snapshot does not include the contents of memory or transactions in progress.
- An application-consistent view of the virtual machine data at a single point in time. With the Datrium VSS Agent installed on a virtual machine running on Windows Server 2008 or 2012, you can take app-consistent snapshots. The Datrium VSS agent uses Microsoft VSS (Volume Shadow Copy Service) to quiesce applications. When activity has ceased, the DVX System pauses I/O to the virtual machine and takes the snapshot.

In addition, each snapshot can serve as a standalone backup, since there is no dependency of between snapshots or upon the snapped object. This allows you to specify different retention periods for snapshots across local or target site (if replicated). You can retain snapshots for as long as necessary and you can easily access snapshots from the Snapstore, even if the original object no longer exists. The Snapstore is effectively a backup-class catalog that provides a search capability for finding and selecting the specific object on which to operate.

Data Pool High Availability

The DVX System provides High Availability (HA) access to data. The system can handle simultaneous failures of two drives in a Data Pool. After a failure, drive rebuilds use all Compute Nodes and Data Nodes, so as the system gets bigger, rebuilds get proportionately faster.

To provide internal HA, each Data Node contains redundant controller and power components to support continuous operation.

Ways that Data Nodes use system redundancy to protect your data:

- Controller module redundancy – If the active controller module fails, the DVX System uses the second controller module for storage operations.
- Battery backup for NVRAM. Each Data Node controller has a pair of dual redundant batteries for sufficient backup power to transfer its local NVRAM contents to durable drives, ensuring proper shutdown in the event of a power failure. Synchronous writes to a Data Node include internal mirroring of NVRAM contents to the second controller.
- Network path redundancy – Each controller supports adaptive path data connections through active/active associated data network interfaces. The associated data interfaces provide aggregate bandwidth. If one of the data paths fails, the DVX System will continue to use the remaining active interface. The controllers also support active/passive bonded pair management interfaces. If the active management path fails, the other management interface provides access to the Data Node. DVX network path redundancy imposes no requirements on switch configuration.
- Power and cooling module redundancy – A Data Node contains two power and cooling modules. If a power and cooling module fails, the other module is capable of handling the entire power and cooling load.
- Erasure coding - Data within DVX systems is erasure-coded to allow for simultaneous failure of i) 2 disks on one Data Node, ii) plus a third disk in a second Data Node, and iii) a latent sector read error on a fourth disk that may be encountered when attempting rebuilds.
- Data Pool Redundancy – In DVX systems with 3 or more Data Nodes, you can activate *Data Node Fault Tolerance* to guarantee data and performance high availability for

unplanned failovers due to network or power loss, maintenance operations, or disasters.

Cloud-based Backup for DR

Cloud-based backup maintains the DVX copy data management activities across on-premises sites for private cloud storage and data storage in an off-site AWS cloud. Economy of scale is achieved by applying global deduplication and incremental forever backups to data in on-premises and offsite Data Pools.

Datrium BaaS with Cloud DVX

The DVX System supports replication to and retrieval from a Datrium BaaS with Cloud DVX. Cloud DVX is a replication destination that runs as an AWS service for archiving snapshots.

- The DVX System encrypts all data in-transit to and from Cloud DVX. Cloud DVX encrypts the data at rest in the cloud.
- The on-premises DVX System sends unique compressed and de-duplicated data to Cloud DVX. The replication operations from multiple on-premises sites and subsequent Cloud DVX operation produce globally deduplicated data on Cloud DVX. This approach significantly reduces the amount of data transferred between the Cloud DVX and the on-premises DVX System and thus reduces the cost per gigabyte of data transferred from AWS.

> **Tip:** When using Cloud DVX, if you are using if blanket encryption with your on-prem DVX, we strongly recommend you enable replication encryption as well.

DVX User Interface

The DVX System supports a Graphical User Interface and a Command Line Interface.

The figure below shows a single DVX Data Node in a VMware vSphere environment. The vCenter Server manages three Compute Nodes. The Hyperdriver software is installed on each Compute Node.

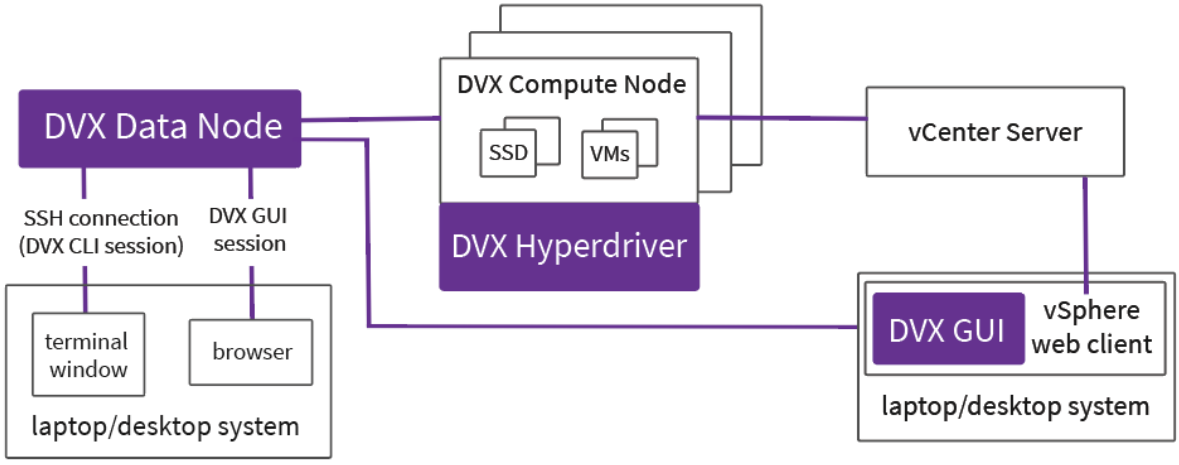

The figure also shows three methods of access to monitor and manage the DVX System.

- You can use the DVX GUI embedded in the vSphere Web Client.
- You can use a browser to create a DVX GUI session that runs on the Data Node.
- You can create an SSH connection to the Data Node to use the DVX CLI.

DVX Support for vSphere Capabilities

The DVX System supports all vSphere data operations except for the following:

- vSphere Host Profiles
- vSphere Content Libraries
- vSphere Auto Deploy
- vSphere Virtual Machine Encryption
- The DVX System supports vSphere Fault Tolerance with the exception of DVX snapshot and replication operations.

DVX Data Pool: Monitoring and Management

A DVX Data Pool consists of a set of 1 to 10 Data Nodes. A DVX Data Pool can contain one to a maximum of 32 Datastores.

This section contains the following topics:

Storage Capacity

A Data Pool with a single Data Node can support up to 1500 powered-on virtual machines distributed across 32 hosts. You can increase the Data Pool capacity by adding Data Nodes. See Expanding the Data Pool.

The physical capacity of a single Data Node is as follows:

- 29TB (26.4TiB) for a D12x4 series Data Node
- 28TB (25.4TiB) for a F24x2 series Data Node
- 72TB (65.4TiB) for a D12x10D Data Node

The DVX Real Time display and the `pool capacity show` CLI command display the available storage capacity.

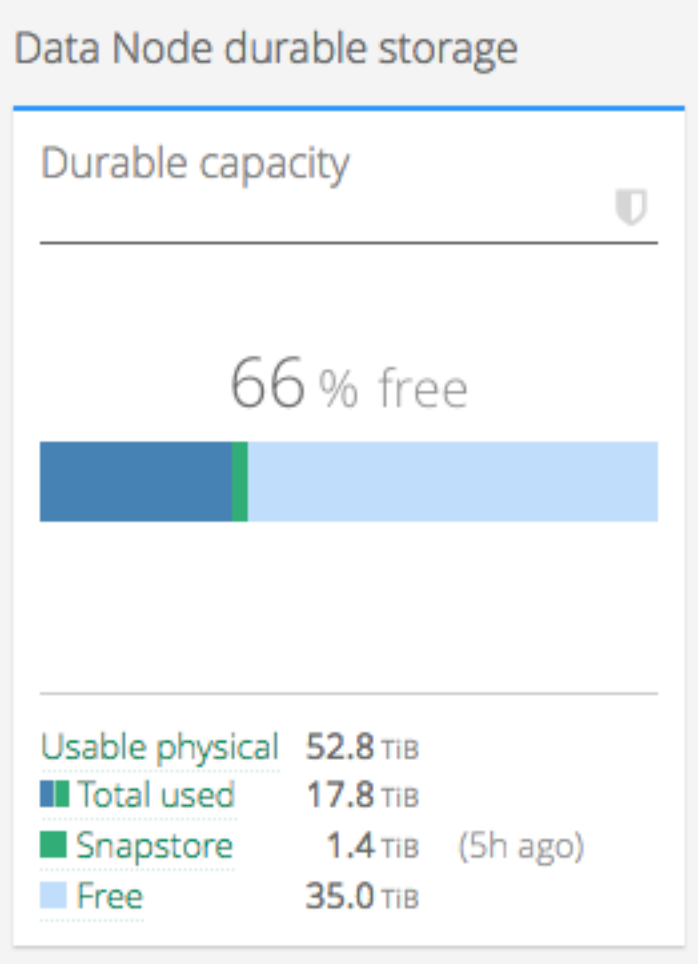


```
b5.controller1>> pool capacity show

-------------------------------------------------------------------
Name                Usable (TiB)   Used (TiB)   Available (TiB)
-----------       -------------- ------------ -----------------
Datastore1  [...]         26.42         7.68             18.74
-----------       -------------- ------------ -----------------
```

The DVX GUI shows real time display of your Date Nodes' durable capacity.

If the DVX System runs out of space, or if it's running low on space and Space Reclamation (SR) is taking a long time, use the The `pool show` command to check the pool state,

which indicates the state of the individual drives. You can also see the drive state(s) in the DVX GUI.

If the pool state is DEGRADED, and there are drives marked MISSING or FAILED, replace the drives to restore normal operation.

```
pool show
State: DEGRADED

---------------------------------------
Slot          State           Health
------------ ---------- ----------
node1.disk01  ACTIVE             OK
node1.disk02  MISSING        UNKNOWN
node1.disk03  ACTIVE             OK
node1.disk04  ACTIVE             OK
node1.disk05  ACTIVE             OK
node1.disk06  ACTIVE             OK
node1.disk07  MISSING        UNKNOWN
node1.disk08  ACTIVE             OK
node1.disk09  ACTIVE             OK
node1.disk10  ACTIVE             OK
node1.disk11  ACTIVE             OK
node1.disk12  ACTIVE             OK
------------ ---------- ----------
```

sc122

Storage Pool States

A DVX System with a full complement of storage drives operates as a fully redundant system that can sustain drive failures, or the failure of an entire Data Node with Storage Capacity activated.

To ensure that the DVX system is redundant:

- On a DVX System with a single Data Node, the DVX System reserves storage space equivalent to one drive's worth of storage capacity to rebuild data.
- On a DVX System with an expanded Data Pool consisting of two or more Data Nodes, and Data Node Fault Tolerance is not activated, the DVX System reserves storage space equivalent to a maximum of two drives' worth of storage capacity to rebuild data. This reserve capacity is not available for datastore filesystems.

- On a DVX System with an expanded Data Pool with 3+ nodes, and Fault Tolerance activated, one entire Data Node is reserved for HA and is not available for datastore filesystems.

The DVX System defines a set of storage pool states to represent storage redundancy and reserve capacity.

Pool State	Description
OK	The DVX System is capable of sustaining two simultaneous drive failures.
DEGRADED	The DVX System has already sustained one drive failure and is capable of sustaining one more drive failure.
CRITICAL	The DVX System can no longer sustain additional drive failures.
DOWN	The DVX System cannot service data requests.

The DVX System response to drive failure depends on the number and sequence of drive failures. Any time a drive fails, you should contact Datrium Support and replace the drive as soon as possible. After you replace the failed drive(s), the DVX System will rebalance the data and the storage pool returns to normal operation.

Recovery from drive failure might delay other system operations, such as Space Reclamation (SR). If the system is already close to full, this delay might cause the system to become full.

Space Reclamation (SR)

During normal operations in a virtual environment, a file delete operation will free up storage space. This occurs when users delete files in the guest operating system environment or when you delete files during virtual machine management operations. Recently freed space will not be available immediately if deduplication operations produced one or more references to the deleted content. When these references are deleted, it produces “dead” storage that you cannot access.

The Datrium Space Reclamation (SR) component analyzes storage to determine space that is no longer referenced, and it restores access to the unused storage. SR is designed to run in the background. There is no conflict between SR and normal operation of the DVX

System. Virtual machines can continue to access the Data Pool without affecting SR processing.

There are two modes of SR execution:

- Automatic operation – The DVX system will automatically run SR as needed. The DVX System will also run SR if 6 hours have elapsed since the last run.
- Manual operation – You can use the `datastore sr start Datastore1` command to start an SR operation manually. If SR is currently running, the command will have no effect.

Space reclamation can be a lengthy process, depending on how much of the DVX System storage capacity is being used. If there is available DVX system space, SR will provide immediate access to the system space for virtual machine storage and it will restore the system reserve space during reclamation.

To obtain information about SR activity, use the DVX CLI command `datastores sr show`. See the description of the`datastores sr` commands in the *DVX Command Line Interface* manual.

Managing Storage Drive Space

Freeing Up Storage Drive Space – The DVX system will not perform reads or writes when full. To free up drive space you can use the vSphere Web Client to delete snapshots or virtual machines, and you can use the DVX UI to delete DVX snapshots. Then you can issue the DVX CLI command `datastorse sr start`. The system might take more than an hour to recover from the system full condition.

Availability of freed storage drive space – There is a time lag between when data is deleted and when the space freed up becomes available for reuse. This lag increases as the system is filled up and might stretch into several hours when there is more than 10TB of data in the system.

Data Pool Resilience for (HA)

DVX Data Pool resilience prevents data loss and ensures high availability of DVX data and management services in the event of disk, controller, power, and entire Data Node failures. Data Pool resiliency ensures that your DVX system will continue serving and storing data normally when:

- Multiple drive failures in a scaled-out Data Pool.
- A controller fails on one of the Data Nodes.
- A Data Node loses network connection or power, or fails completely (for example, for some reason both controllers fail and need to be replaced).

Standard Data Node High Availability

By default, the DVX System provides redundant Data Node controller components to support continuous access to data in the following situations:

- Controller failure – The Data Node has a pair of controllers that operate in active/passive configuration for high availability. If the active controller fails, the DVX System promotes the passive controller to active status and uses the newly active controller to continue to serve data requests.
- Data Node loss of network connectivity – With a Data Pool of more than one Data Node, a DVX System supports dynamic controller failover based on connectivity between the Data Nodes. The DVX System will fail over controllers on Data Nodes to maintain connectivity between all nodes in the Data Pool.
- Single Data Node controller loss of connectivity – If the following conditions are satisfied, the DVX System will failover to the passive controller:
    - The standby controller must have connectivity to more Data Nodes than the active controller.
    - The standby controller set of Data Nodes must include the same set of Compute Nodes that are connected to the active controller.
    - Connectivity between Data Nodes in the Data Pool is not compromised.

- Power loss – In the event of a total power outage, the Battery Backup supports orderly shutdown of the Data Node, including vaulting NVRAM contents to persistent storage.

Optional: Data Node Fault Tolerance (DNFT)

In DVX systems with 3 or more Data Nodes, you can activate Data Node Fault Tolerance (DNFT) to guarantee even higher availability in larger Data Pools for unplanned network or power loss, maintenance operations, or disasters, such as when:

- An entire Data Node fails or loses power.
- A Data Node is unreachable due to power loss or network outage, or due to hardware or software failure on both controllers of a Data Node.

> **Tip:** You can also enable DNFT when adding an extra Data node.

By default, DNFT is disabled. In order to activate DNFT you must have a minimum of 3 Data Nodes in your Data Pool.

Once DNFT is activated:

- One Data Node worth of storage in your Data Pool will be set aside for HA.
- Once enabled, DNFT cannot be disabled.

Data Pool Rebalancing After Enabling DNFT

After enabling DNFT, it can take up to a day before all data in the Data Pool is fully rebalanced. This process does not interfere with normal Data Pool operation.

DNFT and “Claimed” Storage

By design, the DVX storage pool claims a certain amount of storage capacity to support DNFT (when enabled), roughly equivalent in capacity to one entire Data Node, plus a single drive. This space is not available for datastore file systems.

For example:

- In a DVX deployment with 2 or more Data Nodes, *without* DNFT activated, a storage pool can recover from two drive failures, because DVX always retains 2 drives worth of storage. It will use this claimed space to rebuild the contents of the failed drives.
- In the DVX deployment *with* DNFT activated, a DVX storage pool can recover from an entire Data Node failure plus a single drive, by claiming the capacity of an entire Data Node plus a drive for resilience.

Before You Enable DNFT

Before you enabling DNFT in your cluster, you should understand the two general contexts for enabling the feature:

Enabling DNFT when adding a new node to the DVX cluster

You can ONLY enable DNFT when adding a new Data Node IF the current datastore space usage is **less than 85%**. That is, on the DVX dashboard, the Free capacity metric of the Capacity panel must show more than 15% free capacity. This allows for enough free storage during the DNFT enabling process. While DNFT is being enabled, the DVX space reclamation task will run slower at almost half of its original speed.

If your free capacity is **15% or less**, DVX will not allow you to enable DNFT. If your free capacity is **more than 15%** (generally the more free capacity the better), then it is OK to enable DNFT.

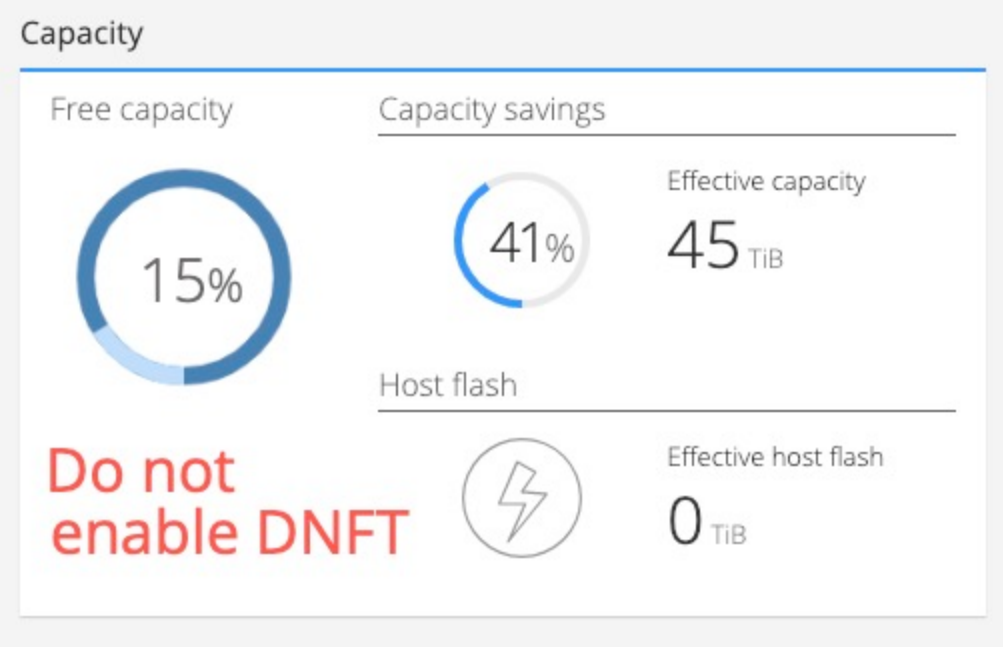

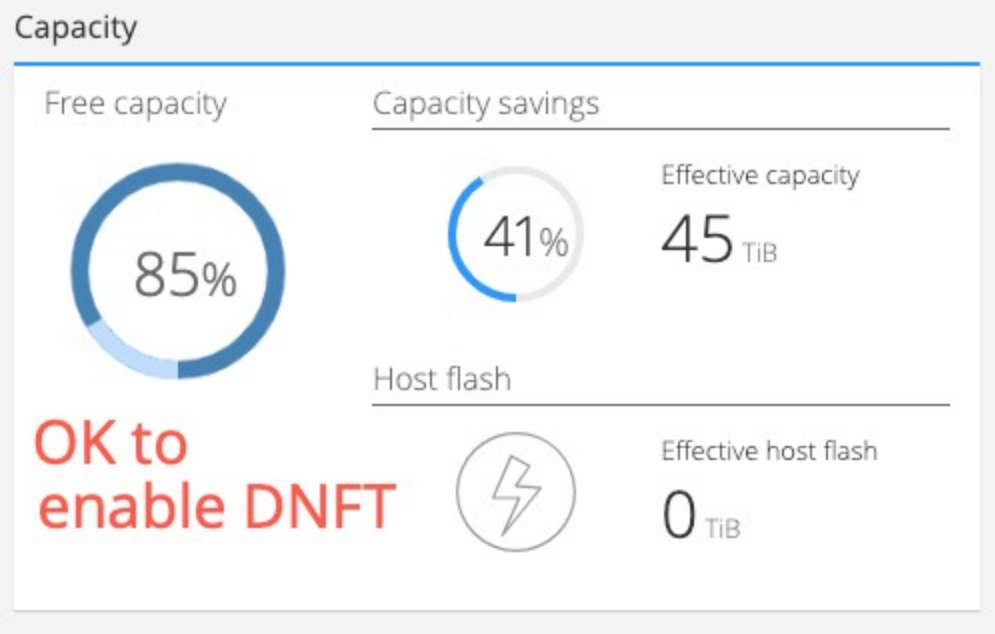

Enabling DNFT but *not* adding a new node to the cluster

If you are enabling DNFT, but you are NOT adding a new Data Node to your cluster, your cluster needs enough free space to accommodate extra parity data for enabling DNFT. This means that you should only enable DNFT in this context if the "projected" free capacity space usage after enabling DNFT is more than 15%.

For example, in a 4 node cluster you need to have current datastore space usage below 63% (this translates to free capacity more than 37%) in order to have the projected space usage after enabling DNFT to be less than 85%.

Because enabling DNFT can take up to 1 day to finish, we recommend that you observe your DVX space usage pattern, plan ahead, and ensure that enough free space available to absorb the extra data written during that time frame.

For DVX cluster with a different number of nodes, the threshold for current datastore space usage varies. Use the following table to project how much free space will be available when enabling DNFT, based on the size of your cluster:

# of nodes in DVX cluster	Minimum free capacity (%) needed to enable DNFT without adding a new node
3	45
4	37
5	33
6	30
7	27
8	26
9	25
10	24

Check if DNFT is enabled

If you aren't sure if DNFT is enabled, you can check the DVX GUI, or use the DVX CLI.

To check if DNFT is enabled - DVX GUI:

1. In the top navigation bar of the DVX GUI, click on the Data Nodes view. The Data Nodes view shows the Data Node(s) in the Data Pool.

2. Look at the bottom of the list of Data Nodes and check the value for DNFT. It will show as either Enabled or Disabled.

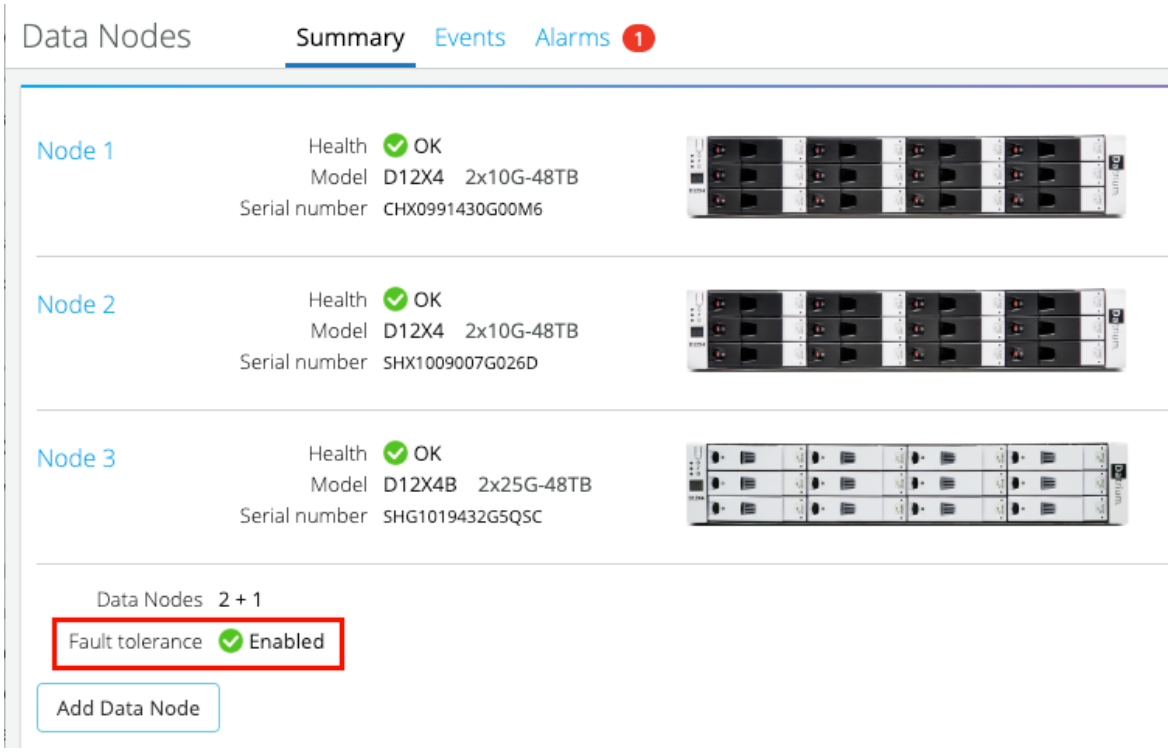

To check if DNFT is enabled - DVX CLI

1. Log in to a Data Node and enter the following command:

   ```
   nodes fault-tolerance show
   ```

2. The output will show the following information:

   ```
   Data Node fault-tolerance:  Enabled
   ```

   If DNFT is not enabled, the value shown will be `Disabled`. Also, if there has been a Data Node failure, the output of this command will indicate the current fault tolerance status for the Data Pool. For example, if there is a Data Node failure and DNFT is enabled, running this command will indicate the failure status as degraded, and then recovering:

   ```
   nodes fault-tolerance show
   Data Node fault-tolerance:  Enabled - degraded
   ```

And as the DVX system recovers:

```
nodes fault-tolerance show
Data Node fault-tolerance:  Enabled - degraded (recov-
ering...)
```

Enable DNFT

You can enable DNFT from the DVX GUI, either by clicking the Enable Data Node Fault Tolerance button, or when you are Adding a Data Node.

You can also enable DNFT using the DVX CLI, either by enabling it and/or when you add a Data Node.

> **Important:** Once DNFT is enabled, you cannot disable it.

When you enable DNFT, the DVX System starts a background process to rebalance data across the entire Pool. Rebalancing DVX System data will take several hours or longer. This process does not interfere with normal Data Pool operation.

If your Data Pool has less than 3 nodes (3 minimum required), then DNFT will be disabled and will not be available until you are Adding a Data Node:

Less than three Data Nodes:

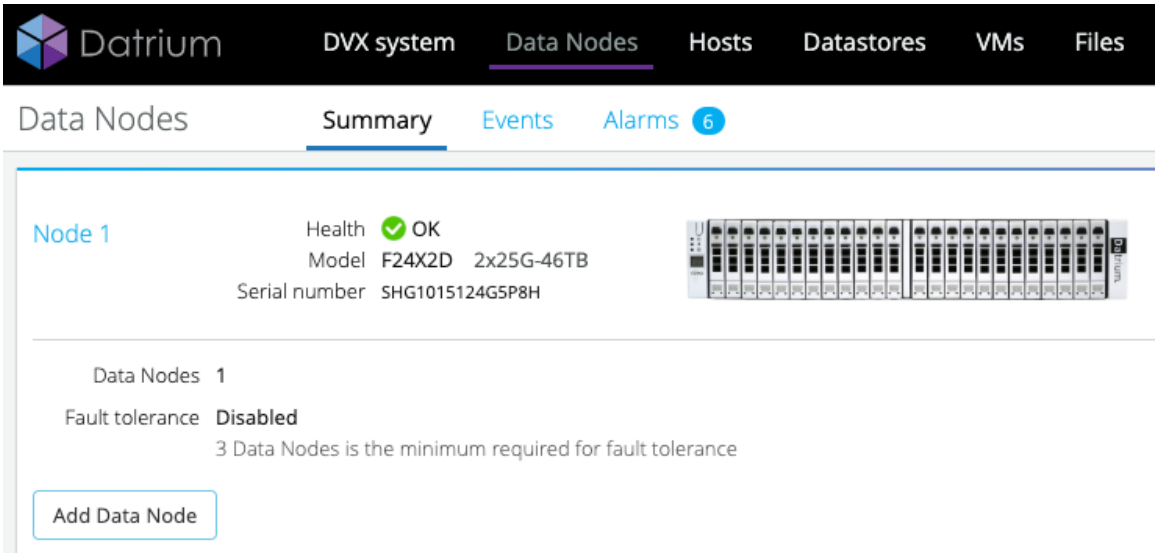

Three or more Data Nodes:

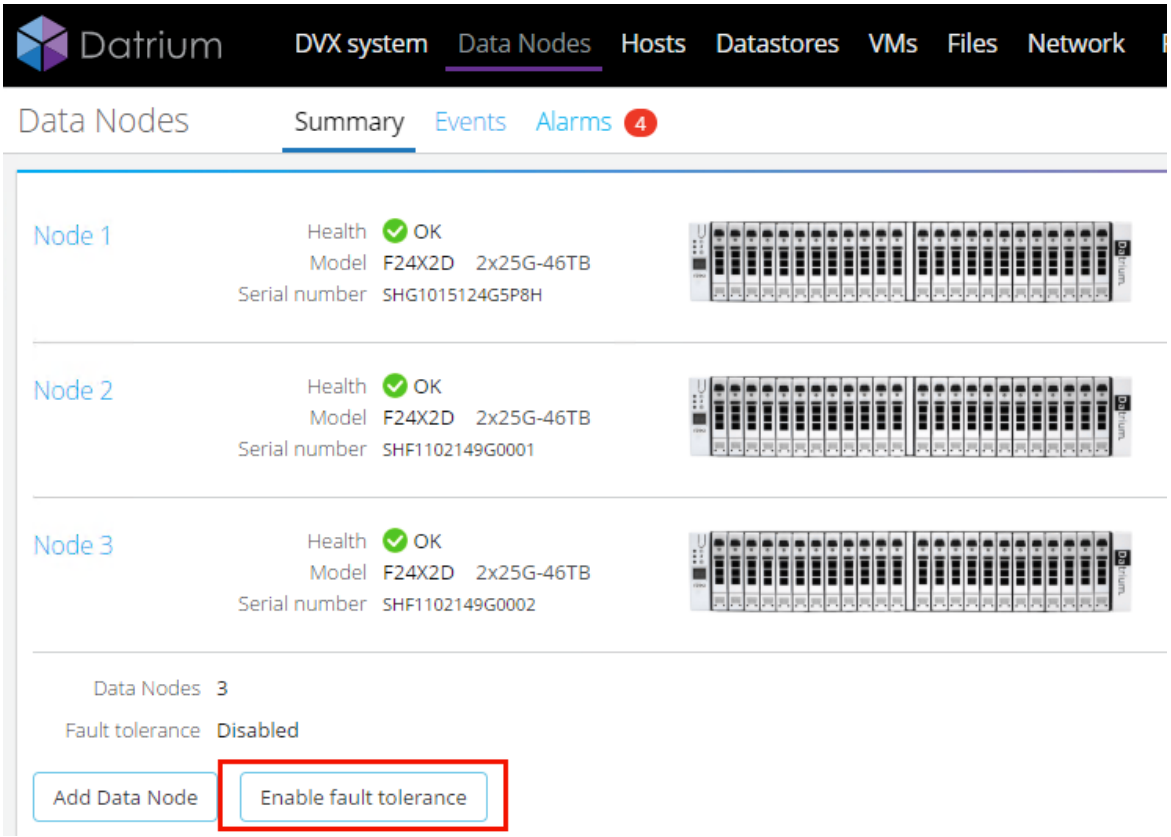

Enable DNFT – DVX UI

> **Important:** If your Data Pool has only 15% or less 'free capacity', you cannot enable DNFT. If you attempt to enable DNFT without sufficient free capacity, the DVX UI will display an error stating: "Failed to enable Fault Tolerance. Try again later."

1. In the top navigation bar of the DVX GUI, click on the Data Nodes view. The Data Nodes view shows the Data Node(s) in the Data Pool.
2. At the bottom of the list of Data Nodes, click the Enable button next to Data Node Fault Tolerance.
3. In the Enable Data Node fault tolerance dialog, type the words ENABLE FAULT TOLERANCE and then when you are ready to enable it, click OK.

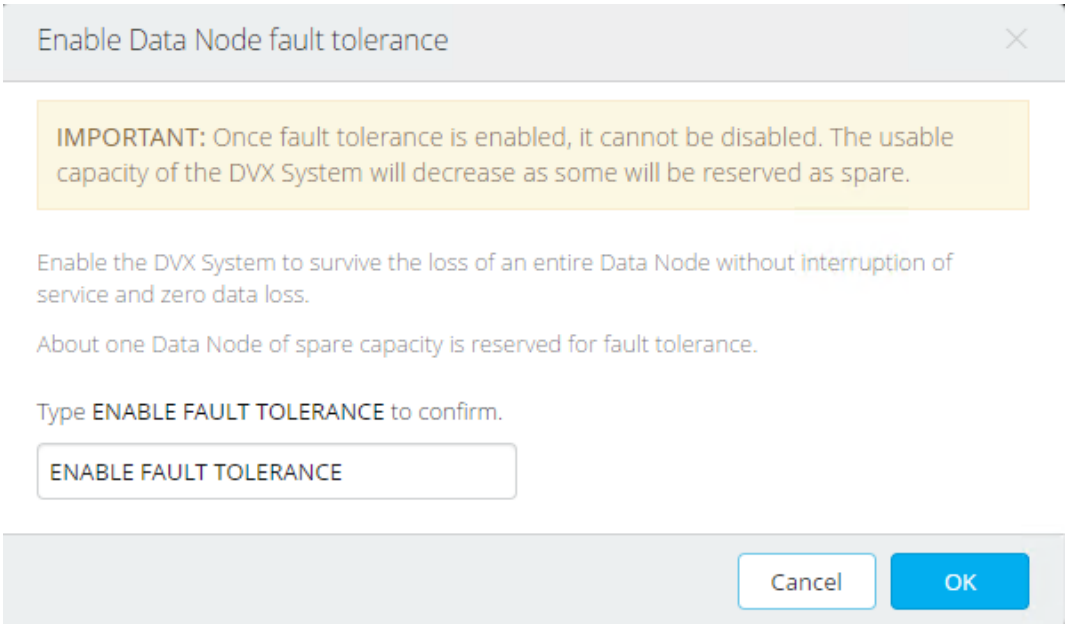

Enable DNFT – DVX CLI

> **Important:** If your Data Pool has only 15% or less 'free capacity', you cannot enable DNFT. If you attempt to enable DNFT without sufficient free capacity, the DVX CLI will display the error: "Data Node fault tolerance requires at least x TiB available capacity."

1. Log in to a Data Node and enter the following command:

   ```
   nodes fault-tolerance enable
   ```

   If DNFT cannot be enabled (you do not have at least 3 Data Nodes in your Data Pool), or it is already enabled, or you do not have enough free capacity (at least 15% or more), the CLI will indicate so in the output.

Expanding the Data Pool

You can add Data Nodes to the DVX System to increase the number of virtual machines, hosts, and storage capacity supported by the System. You can create the following sizes of DVX Data Pools:

Data Node Type	Maximum Nodes
D12x4 series	10
D12x10D	10
F24x2 series	5

If you have activated Optional: Data Node Fault Tolerance (DNFT), the capacity of one of the Data Nodes is claimed for HA. This means if you expand the Data Pool with DNFT turned on, the overall Data Pool Capacity increases by less than the sum of the raw capacity of the additional data nodes you're adding. For example:

Data Node Type	Maximum Nodes	Claimed Data Node Capacity for DNFT	Usable Data Node Capacity
D12x4 or D12x10 D series	10	1	9
F24x2 series	5	1	4

1 to 10 Data Nodes

Each Data Node supports 1500 powered-on virtual machines distributed across 32 hosts.

Each additional node, up to the maximum number of 10 nodes, increases the usable storage capacity. The amount of additional capacity depends on the type of Data Pool, and total scaled out capacity depends on whether or not you have Optional: Data Node Fault Tolerance (DNFT) enabled or not.

- In a **D12x4 series Data Pool**, each additional Data Node increases the usable storage capacity by **29TB**(**26.4TiB**).
  - Using a ten-node Data Pool *without* DNFT, the DVX System can store 580 - 1740 TB per 10-node DVX.
  - Using a ten-node Data Pool *with* DNFT on, the DVX System can store roughly 820 TB of deduplicated and compressed data.
- In a **D12x10 Data Pool**, each additional Data Node increases the usable storage capacity by **73TB** (**66TiB**).

    - Using a ten-node Data Pool *without* DNFT, the DVX System can store up to 1460 - 4380 TB per 10-node DVX.
    - Using a ten-node Data Pool *with* DNFT on, the DVX System can store up to 2 petabytes of deduplicated and compressed data.
- In a **F24x2 series Data Pool**, each additional Data Node increases the usable storage capacity by **29TB** (**25.4TiB**).
    - Using a five-node Data Pool *without* DNFT, the DVX System can store 577 TB - 1730 TB per 10-node DVX.
    - Using a five-node Data Pool *with* DNFT on, the DVX System can store up to 330 TB of deduplicated and compressed data.

> **Important:** The Data Node must be in the factory initialized state. If you want to add a Data Node that is already in use, contact Datrium Support for information about how to proceed.

You must connect both management and data interfaces to your network. Both the management and data interfaces for each Data Node in a Data Pool must be in the same subnet.

> **Important:** Datrium *strongly* recommends that you use separate subnets for data and management traffic.

Adding a Data Node

When you add a Data Node to a Data Pool, you select the network interfaces that you will use for management and data traffic, and you supply IP addresses for the corresponding ports on both controllers of the new Data Node. The DVX System will integrate the new Data Node into the data pool and start using the additional capacity immediately.

You only need to power on the new Data Node and connect it to the physical (or virtual) subnets that the existing Nodes in your Data Pool use.

> **Caution:** With DVX Software 5.1.1.0 , do not attempt to replace the Power Control Module (PCM) Field Replacement Unit (FRU) on any existing Data Nodes in the Data Pool while a new Data Node is being added.

When you use the DVX GUI or CLI, the DVX System will use zeroconf to discover the new Data Node. You can then select the network interfaces for data and management traffic and assign IP addresses to the corresponding controller ports.

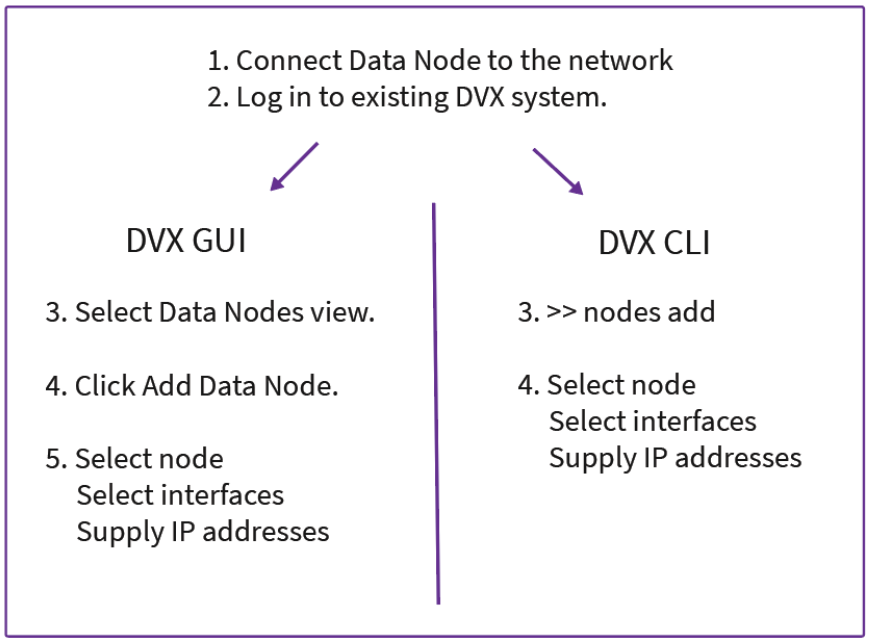

If zeroconf discovery does not identify the new Data Node, see Adding a Data Node Without zeroconf Discovery.

You can add only one Data Node to a Data Pool at a time. If you are adding more than one Data Node to the Pool, you must wait for the rebalancing task to complete before adding the next Data Node. See Data Pool Rebalancing After Adding a Data Node.

Optional: If you have deployed two or more Data Nodes, you have the option to enable Enable DNFT (DNFT) after deploying the third node.

Data Pool Compatibility:

- Data Pools are homogeneous.
  - You can create D12x4 series Data Pools, D12x10D Data Pools, and F24x2 series Data Pools.
  - You cannot mix D12x4 series Data Nodes with D12x10 Data Nodes in the same Data Pool.

    - You cannot mix hard-drive D12x4 series or D12x10D Data Nodes and all-flash F24x2 series Data Nodes in the same Data Pool.
- To add a D12x4C model Data Node to an existing D12x4 series Data Pool, all of the Data Nodes in the Data Pool must be running DVX version 4.1.x or later. Upgrade the existing Data Pool, before you add the D12x4C Data Node to the Pool.
- To add an F24x2D model Data Node to an existing F24x2 series Data Pool, all of the Data Nodes in the Data Pool must be running DVX version 5.x. Upgrade the existing Data Pool, before you add the F24x2D Data Node to the Pool.

You can use the `nodes discover` CLI command to display relevant information. If there is an incompatibility, the command indicates the Node with the message “The node model is not compatible”.

If you use the `nodes add` command to add an incompatible F24x2D, you will get “DaExceptionIncompatibleModel” and the message will be “Model F24x2D is incompatible.

Important: Before you add a D12x4C Data Node to a D12x4 Pool, or before you add an F24x2D to an F24x2 Pool, you must upgrade the existing DVX Systems. The upgrade process will upgrade all DVX Nodes in the Data Pool.

Adding a Data Node – DVX GUI

Caution: With DVX Software 5.1.1.0 , do not attempt to replace the Power Control Module (PCM) Field Replacement Unit (FRU) on any existing Data Nodes in the Data Pool while a new Data Node is being added.

Important: The Data Node you add must be in the factory initialized state. If you want to add a Data Node that is already in use, contact Datrium Support for information about how to proceed.

Use the following procedure to add the Node.

Note: You can only add one Data Node at a time.

1. In the top navigation bar of the DVX GUI, click on the Data Nodes view, which displays all Data Node(s) in the Data Pool. If you are using the embedded GUI in the vSphere

Web Client, and you have more than one DVX System registered with the vCenter Server, make sure that the correct DVX System is selected.

2. In the Data Nodes View, at the bottom of the list of Data Nodes, click Add Data Node. (Or, you can click the Add Data Node button on the top right of the DVX GUI.) The Add Node dialog displays and begins displaying only those Data Nodes that can be added to this Data Pool. If you select Show incompatible Data Nodes, the GUI expands the list to include any Data Nodes that are currently in use in other DVX Systems.
3. In the Add Node dialog, select the Data Node for addition to the Data Pool.

    a. Optional: Data Node Fault Tolerance (DNFT) (Optional). Select this option if you already have deployed two or more Data Nodes and you want to enable your DVX system to be able to survive the loss of an entire Data Node without interruption of service or data loss. Once DNFT is enabled, it cannot be turned off. Also, with this option enabled, the equivalent of one entire Data Node worth of storage is reserved for HA.

    b. Type ENABLE DATA NODE FAULT TOLERANCE to confirm. If you do not want to enable this option, go to the next step.

4. Click Next. The GUI displays fields for network interface selection and IP addresses.
5. Select the management and data interfaces that the DVX System will use on the new Data Node.
6. Enter IP addresses for the corresponding management and data ports on both controllers of the new Data Node. The management interfaces for each Data Node in a Data Pool must be in the same subnet. Likewise, the data interfaces for each Data Node in a Data Pool must be in the same subnet.
7. Click Add Data Node. You might see a temporary increase in latency while the DVX System adds the Data Node to the Data Pool.

Adding a Data Node – DVX CLI

Caution: With DVX Software 5.1.1.0 , do not attempt to replace the Power Control Module (PCM) Field Replacement Unit (FRU) on any existing Data Nodes in the Data Pool while a new Data Node is being added.

> **Important:** The Data Node you are adding must be in the factory initialized state. If you want to add a Data Node that is already in use, contact Datrium Support for information about how to proceed.

To use the DVX CLI to add a Data Node to an existing Data Pool, log in to the existing DVX System and use the CLI `nodes add` command. This command performs discovery and presents a wizard that guides you through the steps:

To add a Data Node – DVX CLI:

> **Note:** You can only add one Data Node at a time.

1. SSH in to the DVX system and run the following command:

   ```
   nodes add
   ```

2. The output shows you a list of Data Nodes to choose from (if available in your environment). Select the Data Node to add.
3. Select Network interfaces for both management and data connections.
4. IP address assignment for the corresponding ports on both controllers of the Data Node.
5. If you have at least 2 Data Nodes deployed and want to enable Optional: Data Node Fault Tolerance (DNFT), use the `--enable-data-node-fault-tolerance` argument to enable it.

   For example:

   ```
   nodes add -enable-data-node-fault-tolerance
   ```

Replacing a Failed Data Node

If you have Optional: Data Node Fault Tolerance (DNFT) enabled, you can replace a failed Data Node with a new one. Before you replace a Data Node with a new one, make sure the new Data Node is connected directly to both the data and management networks.

Note: If the Data Node is functioning properly, it cannot be replaced; you can only add new Data Nodes if all nodes are healthy, and any malfunctioning Data Nodes have been shutdown and disconnected.

Caution: With DVX Software 5.1.1.0, do not attempt to replace the Power Control Module (PCM) Field Replacement Unit (FRU) on any existing Data Nodes in the Data Pool while a new Data Node is being added.

To replace a failed Data Node — GUI:

1. In the top navigation bar of the DVX GUI, click on the Data Nodes view, which displays all Data Node(s) in the Data Pool.
2. In the Data Nodes View, at the bottom of the list of Data Nodes, click Replace Data Node. The Replace Data Node dialog displays only those Data Nodes that can be used to replace the existing Data Node. If you select Show incompatible Data Nodes, the GUI expands the list to include any Data Nodes that are currently in use in other DVX Systems.
3. The GUI displays fields for network interface selection and IP addresses. By default, the DVX system will default to the IP addresses used by the old Data Node. However, you can reassign a new IP addressed, if desired.
4. Select the management and data interfaces that the DVX System will use on the new Data Node.
5. You can reuse the existing IP addresses from the failed for the corresponding management and data ports on both controllers of the new Data Node (or you can also use new IP addresses, if you wish). The management interfaces for each Data Node in a Data Pool must be in the same subnet. Likewise, the data interfaces for each Data Node in a Data Pool must be in the same subnet.
6. Click Replace Data Node. You might see a temporary increase in latency while the DVX System replaces the Data Node to the Data Pool.

To replace a Data Node — CLI:

1. SSH in to the DVX system and enter the following command:

   nodes add

2. The output will indicate that there is a failed node, and after discovery, will provide a list of available Data Nodes to choose as the replacement.
3. Select Network interfaces for both management and data connections.
4. Choose the IP address assignment for the corresponding ports on both controllers of the Data Node.

Adding a Data Node Without zeroconf Discovery

If zeroconf discovery does not identify the new Data Node, you must use the DVX CLI to add the Data Node to the Data Pool.

> **Caution:** With DVX Software 5.1.1.0 , do not attempt to replace the Power Control Module (PCM) Field Replacement Unit (FRU) on any existing Data Nodes in the Data Pool while a new Data Node is being added.

1. Log in to the new Data Node and assign a temporary IP address to the management interface. Then, log in to the management interface of the existing DVX System and use the “nodes add” CLI wizard to complete the process.
2. **DVX CLI** – Log in to new Node (zeroconf software required on client system).
3. Set up a Temporary Ethernet Connection to the New Data Node.
4. Use the Data Node Zero Configuration URL to create the CLI session with the Data Node.
5. Assign a Temporary IP Address to the New Node.
6. **DVX CLI** – Log in to existing DVX System.
7. Add the New Data Node to the Data Pool (`nodes add --ip`).

Temporary Ethernet Connection to the New Data Node

The laptop or other system that you use to connect to the new Data Node must have a multicast DNS responder daemon, and your system must allow access over UDP port 5353.

For initial configuration, use an Ethernet cable to connect your laptop or other client system directly to the eth1 management port on the Data Node. Make sure to connect to

the management port on the active controller. The temporary connection provides access for the CLI session.

The picture below shows a Data Node that does not have any other network connections. It is not necessary to connect the Data Node to your network to perform initial configuration, but you can do so if desired.

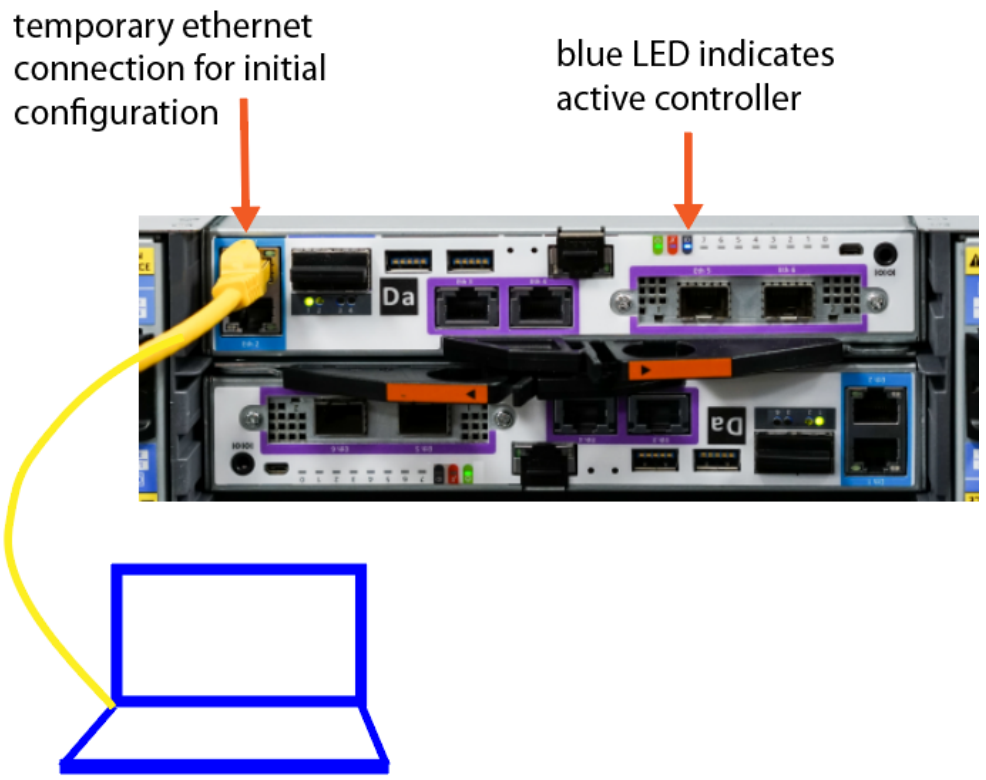

When you save the network configuration, the DVX System will disable the zeroconf connection. At that time, replace the temporary connection with a connection between the configured management port and the management subnet. Then you can use the management floating IP address to log in over the management subnet. If you have not already connected the other ports that you have configured, do so at this time.

If you do not have multicast DNS responder software, you must use the Data Node serial port for the connection. For information about obtaining multicast DNS responder software or about using serial port access, see “Data Node and DVX Configuration (CLI Setup Wizard)” in the *DVX Software Configuration* manual.

Data Node Zero Configuration URL

The Data Node zeroconf URL contains the Data Node host name. The host name has three components:

- Host name prefix – “datrium-NS”.
- Data Node serial number – The serial number is on the back of the chassis (see below).
- Host name suffix – Identifies the active controller on the Data Node.

Controller	Host name suffix
1	-C1.local:7443
2	-C2.local:7443

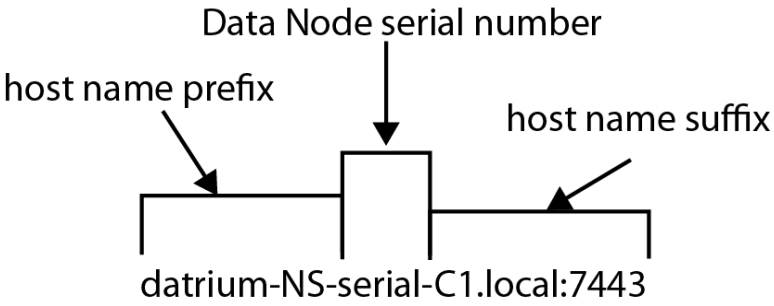

The Data Node serial number is located on the inside of the left “ear” of the PCM on the left side of the Data Node.

The following string is an example of a zeroconf URL with a Data Node serial suffix of `0000000000000000`:

```
https://datrium-NS-0000000000000000-C1.local:7443
```

It may take a few minutes to establish the zeroconf connection. Ping the Data Node zeroconf URL to determine connectivity.

Assign a Temporary IP Address to the New Node

1. Using the DVX CLI, log in to the new Data Node and provide a temporary management IP address.
2. When you log in, the CLI presents the option of creating a new DVX System or of joining an existing DVX System. In this case, specify “join”.
3. Indicate that you are not using zeroconf to add the Data Node to the DVX System.
4. Enter a temporary IP address and subnet mask for the management interface.

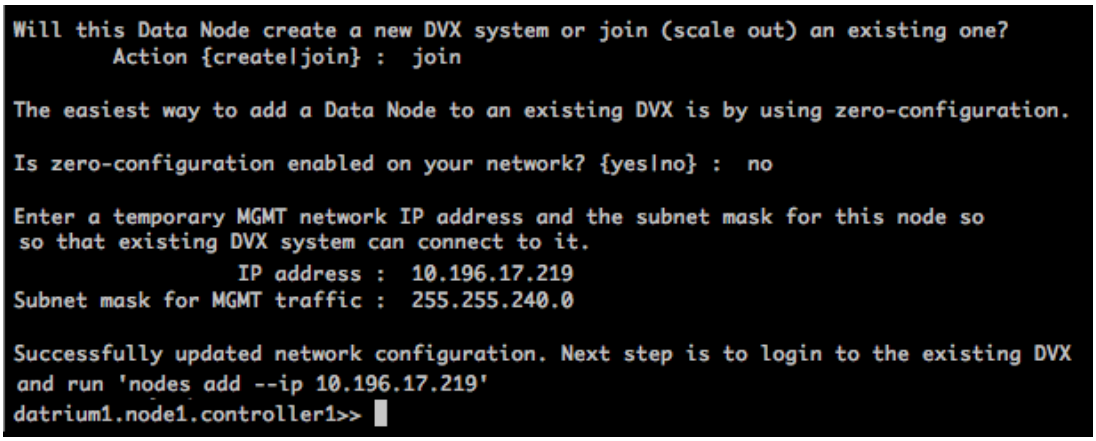

Add the New Data Node to the Data Pool

1. Log in to the management interface for the existing DVX System and use the following CLI command:

   ```
   nodes add --ip tmp-ip
   ```

2. Specify the temporary IP address that was assigned to the new Data Node.
3. When you run the command, the DVX System runs a wizard that locates the specified Data Node and presents the choice of network interfaces for data traffic:
4. Enter the letter for the teamed data interfaces that the DVX System will use for data traffic.
5. Enter the IP addresses for the corresponding ports on both controllers

   This example shows active links for the eth5 and eth6 10Gb SFP+ interfaces.

```
dvx93.node1.controller1>> nodes add --ip 10.196.17.219

ADD NEW NODE

Finding node with IP 10.196.17.219 in local network...

SHG1015124G5T2S


Choose the teamed interfaces for node DATA traffic.
    (a)  eth3 eth4              BaseT         No Link
    (b)  eth5 eth6              SFP+          Link

           Data ports {a|b} :  b

eth5 eth6

Enter the node1 DATA IP addresses
      node1.controller1.eth5 :  10.196.112.219
      node1.controller1.eth6 :  10.196.112.220
      node1.controller2.eth5 :  10.196.112.223
      node1.controller2.eth6 :  10.196.112.224
```

6. Select the interface(s) for management traffic and provide IP addresses for the corresponding ports on both controllers of the new Data Node.

   This example shows an active link for the eth1 1Gb BaseT interface.

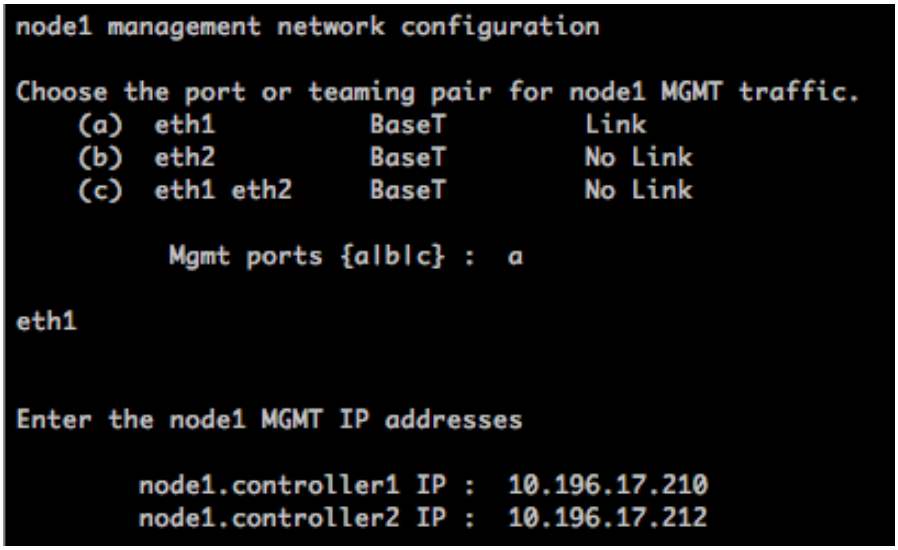

You might see a temporary increase in latency while the DVX System adds the Data Node to the Data Pool.

Data Pool Rebalancing After Adding a Data Node

When you add or replace a Data Node to a Data Pool, or when you enable DNFT, the DVX System uses the expanded pool immediately.

- The DVX System writes new data to Data Node drives across the entire Pool.
- The DVX System starts a background process to rebalance data across the entire Pool. This process does not interfere with normal Data Pool operation.

  The amount of time required for rebalancing across the expanded pool depends on the amount of data in the original pool and the number of hosts in the DVX System. Rebalancing DVX System data of any reasonable size will take several hours or longer. Additional hosts will contribute to the rebalancing effort and can reduce the amount of time it takes to rebalance the data. If you have additional hosts that you can add to the DVX System, add them before you add the new Data Node to the Data Pool.

The DVX UI does not show the expanded capacity until rebalancing is complete.

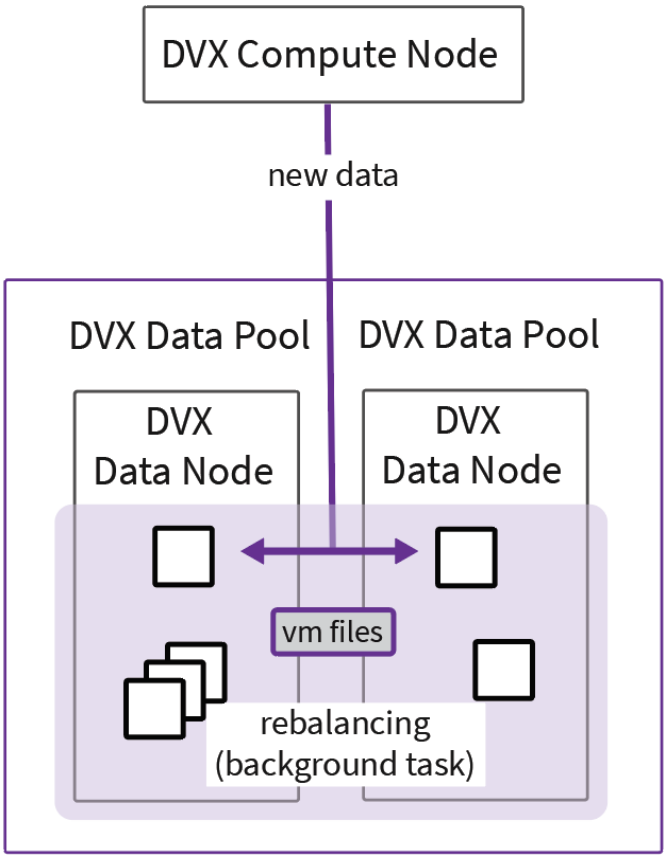

You can add up to two Data Nodes to a Data Pool at a time. If you are adding more than two Data Node to the Pool, you must wait for the rebalancing task to complete before adding the next Data Node. Use the `pool show` command to monitor the rebalancing task. During the task, the UI will show the state as "POOL_REBALANCE".

Monitoring the Data Pool

The DVX GUI Real Time display shows information about Data Node durable storage, including durable capacity, data reduction, scale-up potential, and total network usage. The DVX CLI command nodes show displays overall health for the Data Pool nodes.

When there is an error on a Data Node, the DVX UI reports the error. The DVX GUI shows an event banner on the Dashboard and Real Time views, in addition to the alarms list in the right-hand frame of the GUI page. The DVX CLI events show command displays the list of

events. In both cases, GUI and CLI, the event information identifies the Data Node component on which the error occurred. To get more information:

- DVX GUI – Navigate to the Data Nodes page (select the Data Nodes view) and then click on the appropriate node link. The GUI will display the node page, with health icons indicating the status. Place the cursor over the health icon and the icon for the component to display tooltips with more information.
- DVX CLI – Use the “nodes show” command and specify the appropriate node. The command output indicates the overall node health, and it shows the status for the node components.

Managing Datrium Datastores

A DVX Data Pool can contain one to a maximum of 32 Datastores. When you install a DVX System, the System creates a Data Pool that contains a single Datastore named “Datastore1”. Each datastore has access to the entire available capacity in the Data Pool.

The following sections provide information about managing Datrium Datastores:

Host Access to Datastores

When you install the first Data Node in a Data Pool, the DVX System creates the first datastore. You can create additional datastores, up to a total of 32 datastores in a single Data Pool.

If there is only one datastore in the Data Pool, DVX host configuration automatically mounts the datastore on the host. If there is more than one datastore, the host configuration process presents the list of datastores from which you can select a datastore for mounting. You can mount a maximum of 8 DVX datastores on a single host.

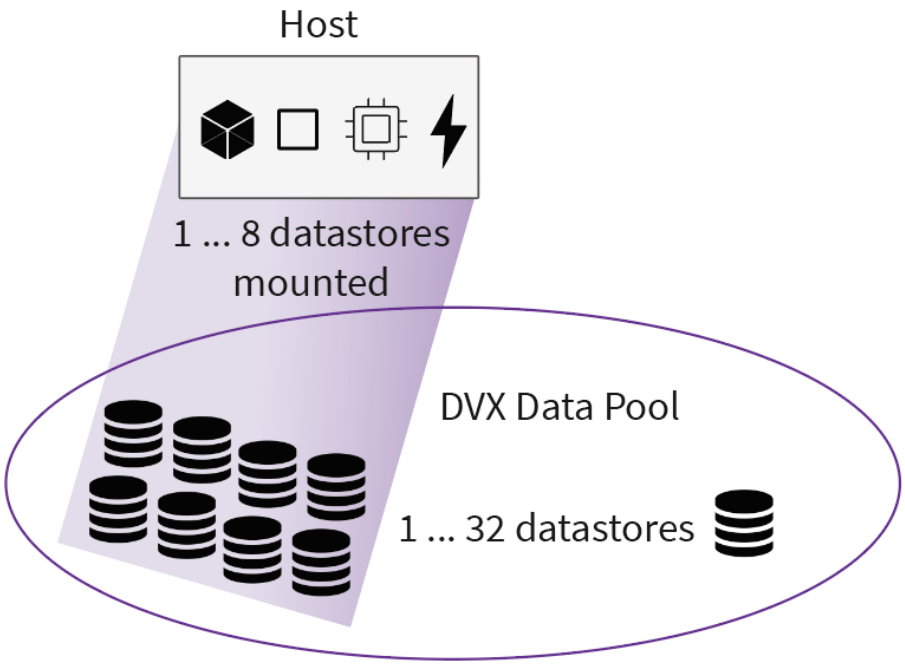

In certain situations, for example Changing the Data Floating IP Address, it is necessary to unmount and later re-mount the datastore.

The following section describes how to unmount and mount a Datrium Datastore. To perform these actions, use the vSphere Web Client. The DVX System does not support the use of esxcli commands to mount or unmount Datrium datastores.

Mounting/Unmounting a DVX Datastore

1. To mount or unmount a DVX datastore, navigate to the vSphere Web Client host page.
2. In the Datrium DVX tab, click on the DVX menu icon (gear icon) and select Mount/unmount datastores7.

   The dialog displays the list of datastores. Slide the GUI control right or left to mount or unmount a host, then click on OK.

   In this example, datastores 2 through 10 are mounted. The GUI control for Datastore2 shows a lighter color, indicating that it is disabled because the host has virtual machines on it.

You cannot unmount a datastore on a host that has virtual machines on it. Move them to a different host or unregister them from the vCenter inventory.

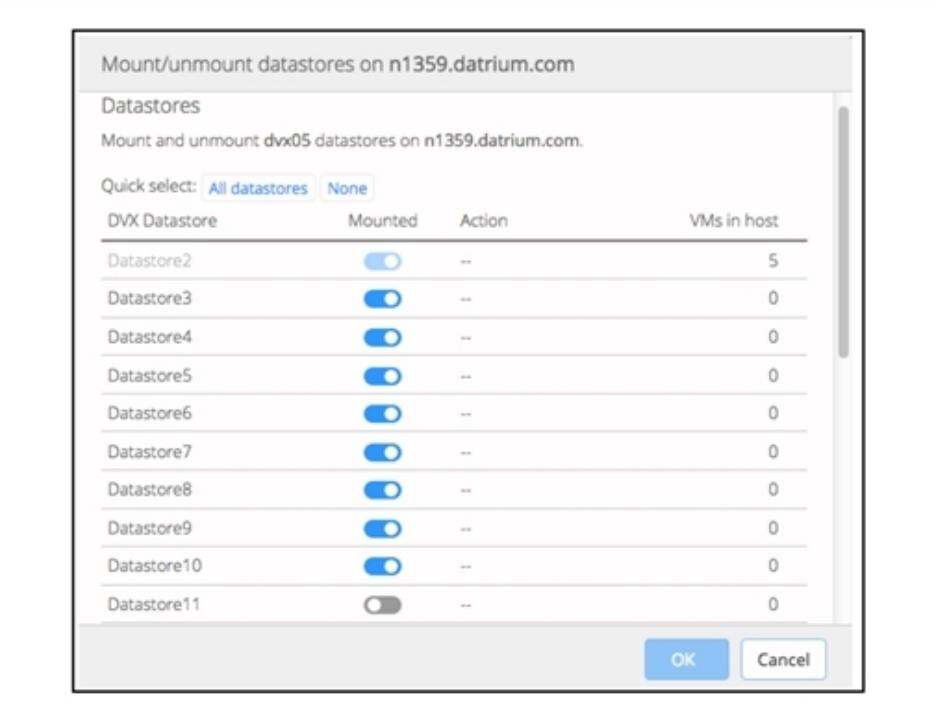

After you unmount the Datrium datastore, it might take up to 2 minutes for the host entry to be removed from the DVX GUI host page or from the output of the “hosts show” CLI command.

Creating a Datastore

You can use either the DVX GUI or DVX CLI to create a new datastore.

Creating a Datastore – DVX GUI

1. The Datastores view contains two Create datastore buttons – one at the upper right, under the gear configuration button, and a second instance underneath the datastores list.
2. Click the Create datastore button to start the process.
3. Enter a name and description for the datastore, then click Create.

4. The dialog also contains a toggle button for the export control list. You can limit access to the datastore by adding hosts to the export control list for the datastore. See Managing Access to a Datastore (Export Control Lists).

Creating a Datastore – DVX CLI

Use the DVX CLI `datastores create` command to create a new datastore.

Deleting a Datastore

You can use either the DVX GUI or the DVX CLI to delete a datastore.

> **Important:** Use care when you delete a datastore. This command removes all files from a datastore.

- You cannot delete a datastore that is mounted on any hosts.
- You cannot delete a datastore whose contents are protected by an active protection group.
- You cannot recover the deleted datastore.
- You cannot delete the last datastore. The DVX System requires a minimum of one datastore in Data Pool.

Deleting a Datastore – GUI

1. In the Datastores view, select the datastore to be deleted. The Delete and Edit buttons are enabled. Click the “Delete” button.
2. You must type the string "DELETE DATASTORE” to confirm the operation. Then click the Delete button.
3. The DVX GUI displays an informational dialog indicating the datastore has been deleted.

Deleting a Datastore – CLI

Use the DVX CLI `datastores delete` command to delete an existing datastore. The command prompts for confirmation and for the admin password.

Managing Access to a Datastore (Export Control Lists)

To limit access to a DVX datastore, use the export control list associated with the datastore. An export control list contains host Data IP addresses or IP subnet specifications.

When you create a datastore, export control is disabled for that datastore.

Export Control List – GUI

You can create and modify datastore export control lists in the “Create datastore” dialog or by editing a datastore. The following sequence is based on editing an existing datastore.

1. In the Datastores view, select a datastore and click the Edit button.
2. In the Edit datastore dialog, click on Use host export control list. The dialog expands to show a frame for the export control list.
3. Enter host data IP addresses or subnets separated by spaces.
4. A host IP address must be the address used for data traffic between the host and the DVX Data Pool. If you specify a subnet (for example, 192.168.100.0/24), any host with a data interface IP address within the subnet will be allowed access to the datastore.
5. If you have enabled the host export control list, an empty export control list prevents access to any hosts in the DVX System.
6. Click on OK to save the export control list

Export Control List – CLI

The DVX CLI provides the following export-control commands.

`datastores`	`show datastore`	Display datastore information, including the export control list.

`datastores export-control`	`add`	Add a host to the datastore export control list. Hosts in the export control list are allowed access to the datastore.
	`remove`	Remove a host from the datastore export control list. If you have enabled export control, an empty export control list prevents access to any hosts in the DVX System.
	`enable`	Enable export control for a datastore.
	`disable`	Disable export control for a datastore.

The following example shows CLI commands to add a host to a datastore export control list, then display the export control list.

```
datastores export-control add datastore --entries hostSpec
[hostSpec ...]
datastores show datastore
```

A host specification for an export control list entry is a host IP address or a subnet. The host IP address must be the address used for data traffic between the host and the DVX Data Pool. If you specify a subnet, any host with a data interface IP address within the subnet will be allowed access to the datastore.

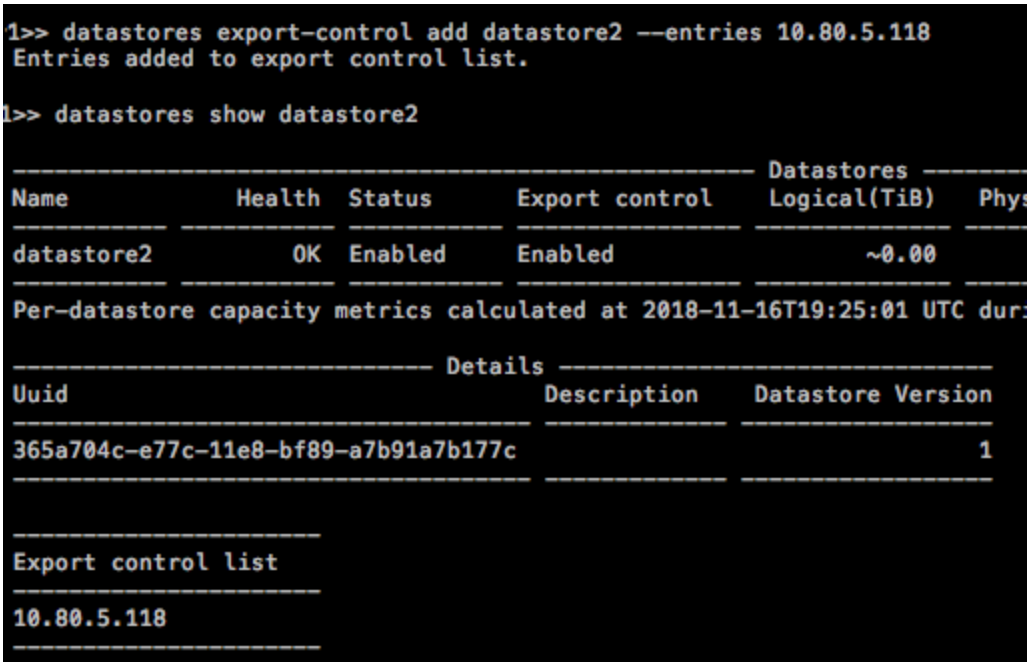

Monitoring the Datastore

You can use either the DVX GUI or CLI to monitor the datastore.

DVX GUI – Datastore

On a vSphere Web Client datastore page, select the Monitor-> DVX tab. The tab display shows the Data Node storage capacity, data reduction, and historical performance.

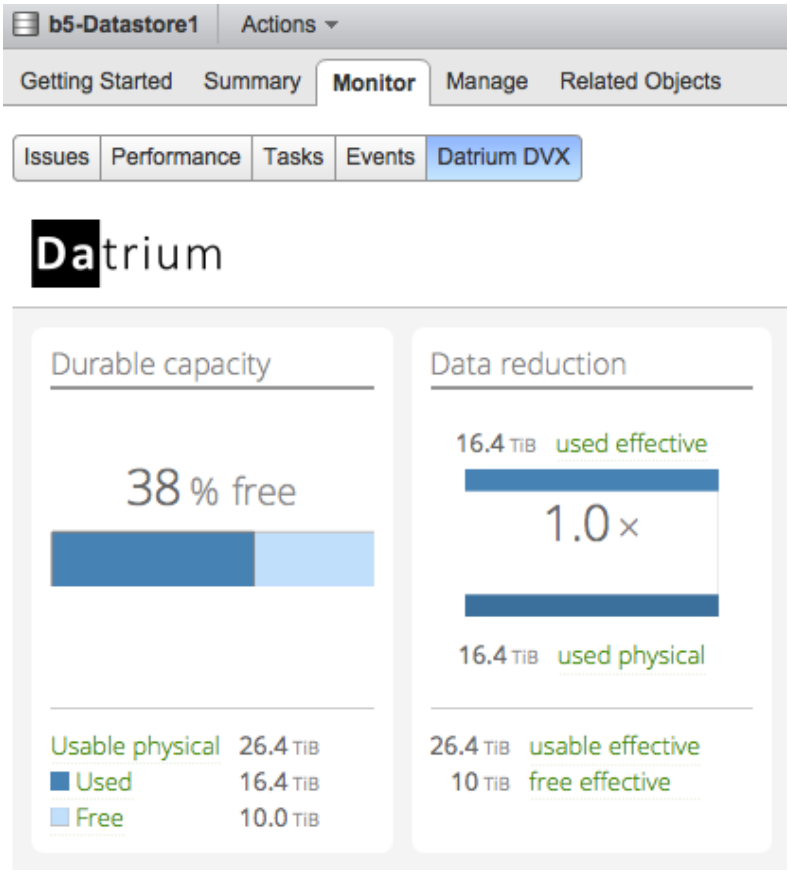

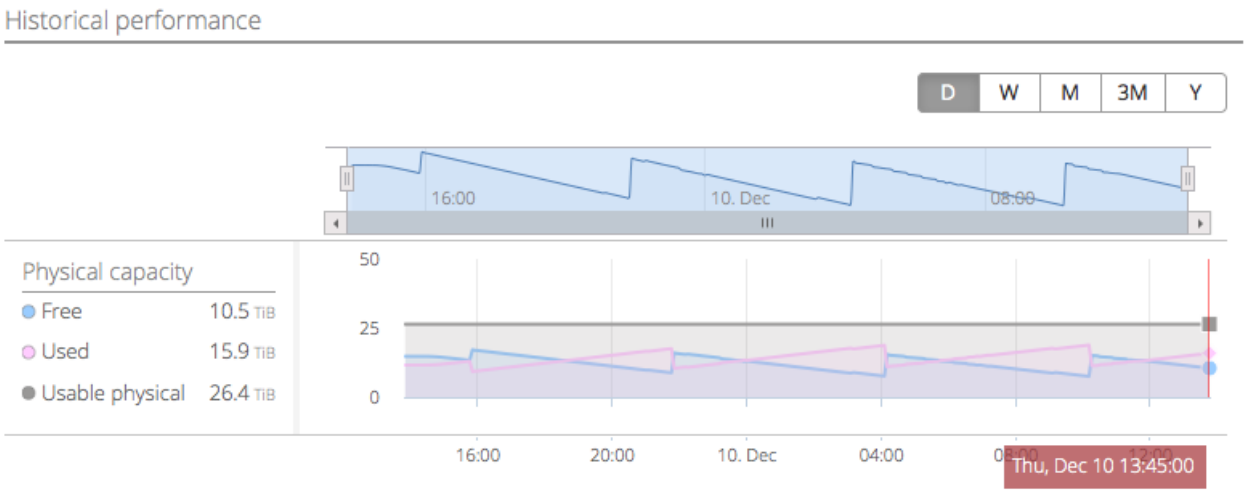

DVX CLI – Datastore

The CLI provides the following datastores commands:

`datastores` **command**	**Description**
`show`	Displays datastore status and capacity information.
`create`	Creates a datastore.
`delete`	Deletes a datastore.
`enable`	Enables the datastore. The datastore is enabled by default.
`disable`	Disables access to the datastore. Before you execute datastore disable, put all hosts that access the datastore into maintenance mode.

`datastores` **command**		**Description**
`encryption` (see Blanket Encryption)	`enable`	Enables encryption on the DVX System.
	`disable`	Disables encryption on the DVX System.
	`password set`	Changes the encryption password. For security, we recommend that you create a password that contains at least 10 characters, including a combination of uppercase and lowercase letters, numbers, and special characters.
	`rotate-key`	Changes the key that encrypts DVX data.
	`set`	Sets the encryption startup mode.
	`show`	Displays data encryption settings.
	`unlock`	Provides access to a system in locked startup mode.
`export-control` (see Export Control List – CLI)	`add`	Add a host to the datastore export control list. Hosts in the export control list are allowed access to the datastore.
	`remove`	Remove a host from the datastore export control list. If you have enabled export control, an empty export control list prevents access to any hosts in the DVX System.
	`enable`	Enable export control for a datastore.
	`disable`	Disable export control for a datastore
`sr` (see Space Reclamation (SR))	`show`	Displays DVX SR (Space Reclamation) status and results.
	`enable`	Enables SR.
	`disable`	Disables the SR feature. When SR is

`datastores` **command**		**Description**
		disabled, automatic SR cannot start and you cannot start manual SR.
	`start`	Starts a space reclamation task. If SR is currently running, the operation that you have invoked will begin after the current operation finishes. If you start an operation and it finishes after the currently scheduled time, the system-scheduled operation will begin at the next scheduled time.

DVX vCenter Server Registration

When you register the vCenter Server, the operation registers the DVX GUI plug-in on the vCenter Server. The vCenter Server will automatically retrieve the plug-in from the Data Node and install it in the Web Client.

You can have a maximum of 8 vCenter Servers registered with a single DVX System.

vCenter registration includes the following actions:

- DVX registration – Provides access to the Data Pool for DVX System setup on hosts.
- DVX GUI plug-in registration – Supports retrieval of the plug-in from the Data Node for installation on the vSphere Web Client platform.

After you register the DVX System on the VCenter Server, the Server uses the plug-in registration to obtain the DVX GUI plugin from the Data Node for installation on the Web Client.

You register a vCenter Server when you perform initial configuration of the DVX System. For a particular DVX System, you can register up to four vCenter Servers in total. To register additional vCenter Servers, use the DVX CLI command `config vcenter register`. You must provide the IP address or DNS name for the vCenter Server and the login credentials for your vCenter account. The account must have administrator privileges

> **Important:** If you register a vCenter Server that is part of a linked mode group, you must register all vCenter Servers in the group.

To see the vCenter Server registrations, use the `config vcenter show` command.

The DVX System identifies the vCenter Server association in the following contexts:

Environment	Context	vCenter link
Data Node GUI	Host and virtual machine pages	
Embedded GUI	DVX hosts page	
CLI	`config vcenter show` `hosts show`	

Powering Down the DVX System

When you power down the DVX System, the System performs an orchestrated shutdown of all Data Nodes in the Data Pool.

Before you shut down the DVX System, unmount the Datrium datastore on the DVX hosts to prevent APD events. After you restore power to the DVX System, it will reboot. When the DVX System is operational again, use the DVX GUI to remount the Datrium datastore on the hosts.

To shutdown the DVX System:

1. Log in to the management interface; use the management interface floating IP address.

   ```
   ssh admin@mgmt-floating-IP
   admin@mgmt-floating-IP's password:
   ```

2. Use the `nodes poweroff` command to shut down the Data Node. This command performs an orchestrated shutdown of all of the Data Nodes in the Data Pool.

   ```
   nodes poweroff
   ```

   After the shut down completes, Data Nodes in the Pool will show the following LED display.

   PCM LEDs – On both PCMs, the upper right LED flashes green. The lower right LED shows continuous green. The lower left LED shows amber.

   

   D12x4 Controller Notification LEDs (both controllers):

   – LED#1: off
   – Power status LED: flicker
   – Fault LED: flicker

   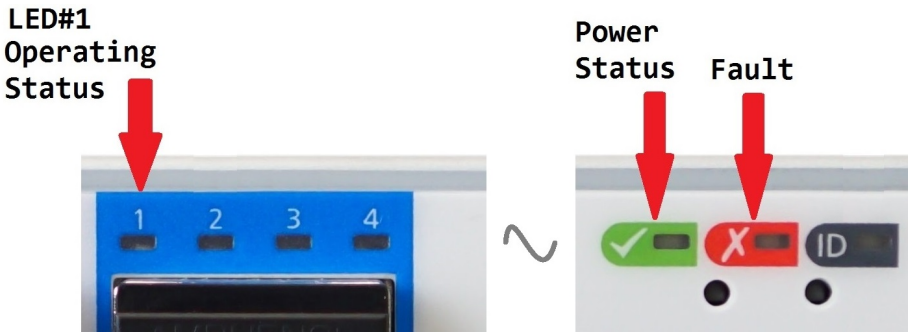

   D12x4B/D12x4C, F24x2B/F24x2D Controller Notification LEDs (both controllers):

   – Power status LED: continuous green (may occasionally flicker)
   – Fault LED: DVX4.1 and later – off
   pre-DVX4.1 – on (continuous)
   – Bootstrap LED #7: flicker
   – Bootstrap LED #0: continuous green

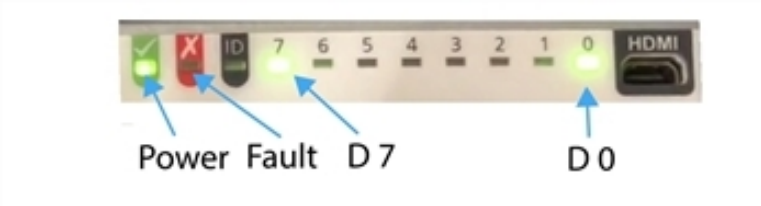

On each Data Node, turn off the PCM power switches and unplug both PCM modules. Before you turn off the power switches, make sure to wait until the controller LEDs indicate that shutdown has finished (step 2). After turning off the PCM switches, the PCM LEDs will show this continuous display for up to 1 minute.

DVX GUI Data Nodes Page

The Data Nodes Page in the DVX GUI presents a graphic representation of the Data Node components. The icons for the different components are color coded, using green, yellow, and red to indicate normal (green), caution (yellow), and error (red) conditions.

The picture below shows the DVX GUI representation of a Data Node that is configured to use both management and data networks. The configured ports are represented by green icons. The management interfaces, eth1 and eth2, are configured as a bonded port pair to support network failover. The data interfaces, eth3 and eth4, are configured to use adaptive path networking for aggregate bandwidth. Controller 1 is the active controller.

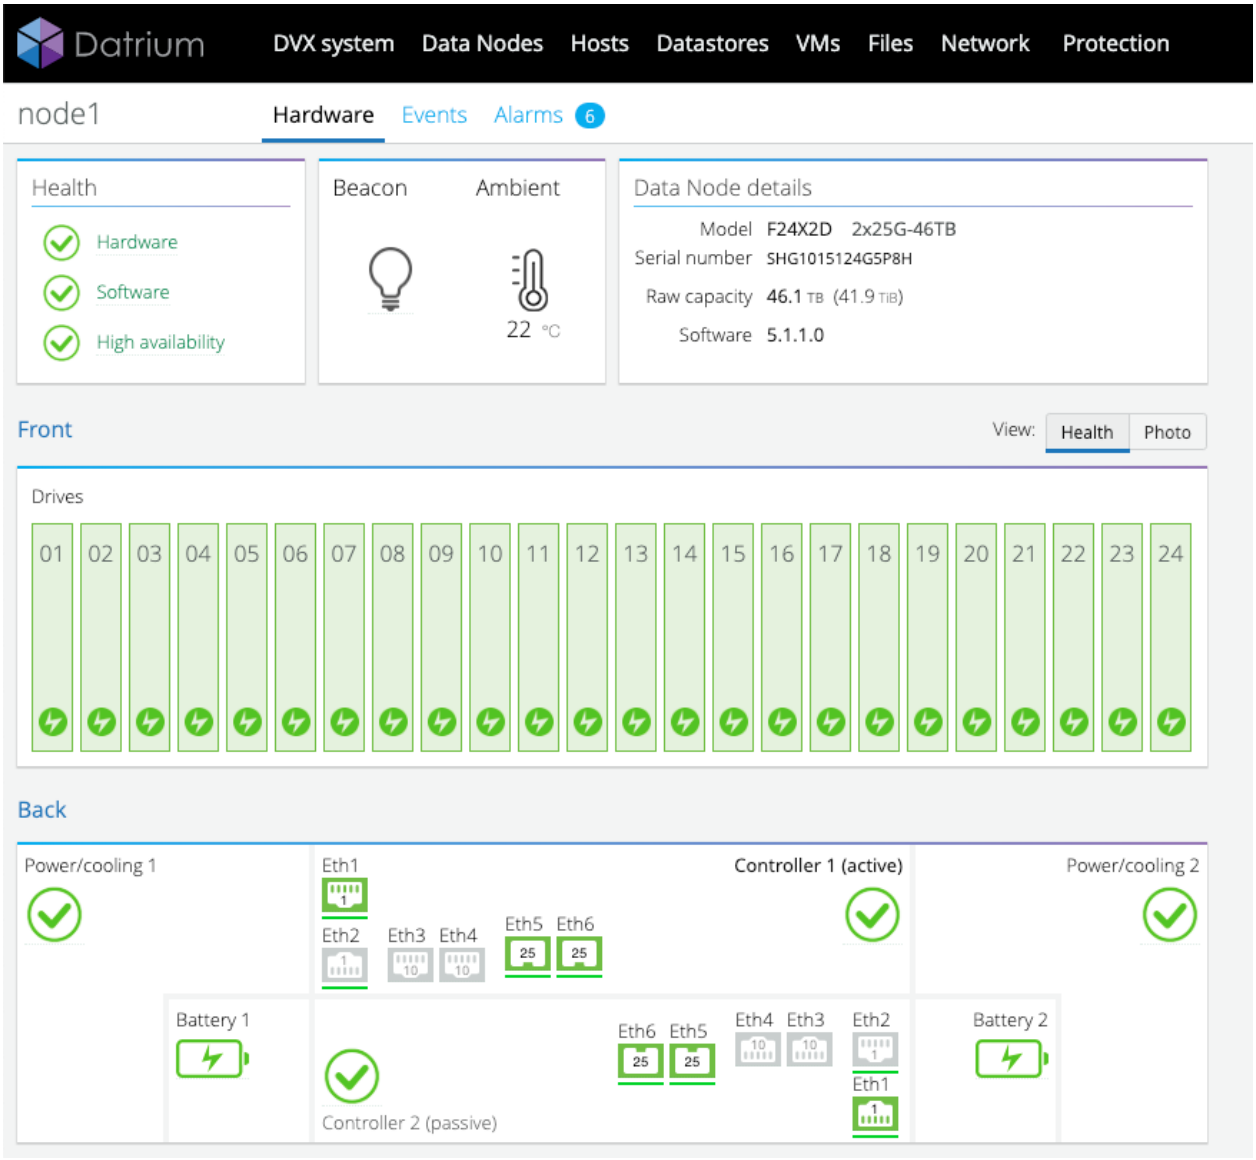

The picture below shows the front of the D12x4 Data Node.

The next picture shows a picture of the back of a Data Node. Cables are connected to the eth1, eth3, and eth4 ports on both controllers to support HA controller failover. The DVX

System uses the controller notification LEDs to indicate the controller power status and the active controller. For more information, see DVX GUI Data Nodes Page.

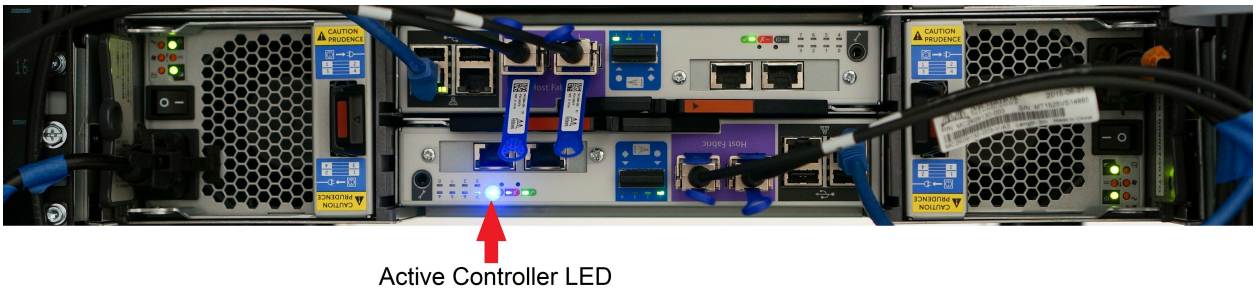

Hardware Health

The Data Node Page uses green-colored icons to indicate the health of the different Data Node components. The Hardware pane indicates general status:

- Overall hardware health. A green icon indicates that all components are healthy. If the status of any component changes, the hardware health icon will change color.
- Drive pool health. (status and health)
- High availability status. (status and health)

Data Node Beacon

A Data Node uses LEDs for beacon identification. When you turn on the beacon for a Data Node, LEDs on both front and back will flash.

- On the front, the following LEDs flash in synchrony:
  - Drive replacement LED flashes amber.
  - Node ID flashes.
- On the back, the #3 and #4 notification LEDs flash on both controller modules.

You can use either the DVX GUI or the DVX CLI to control the Data Node beacon capability.

Using the DVX GUI to Control the Data Node Beacon

The beacon icon controls the Data Node beacon capability. Place the cursor over the light bulb graphic. The tooltip window contains the beacon control button. When you turn on the beacon, the icon in the Data Node hardware GUI display flashes yellow until the Data Node LEDs start flashing. After the LEDs start flashing, the GUI icon shows a continuous yellow.

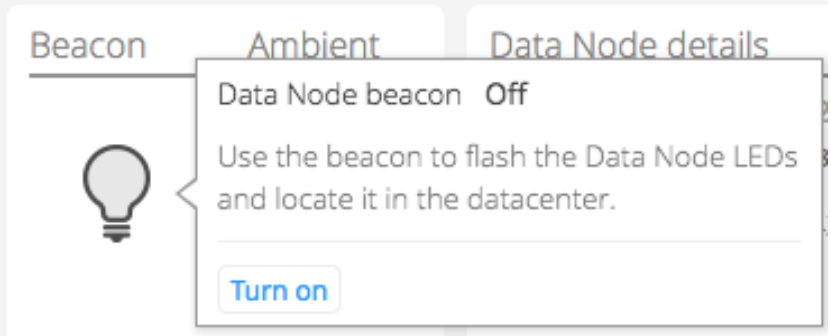

Using the DVX CLI to Control the Data Node Beacon

The CLI provides the following commands to control the Data Node beacon capability.

`nodes beacon set`	Turns on the Data Node beacon capability.
`nodes beacon unset`	Turns off the Data Node beacon capability.
`nodes show`	Displays information about the Data Node, including the beacon status.

Data Node Ambient Temperature

The Data Nodes Page in the DVX GUI shows the ambient temperature for the Data Node enclosure.

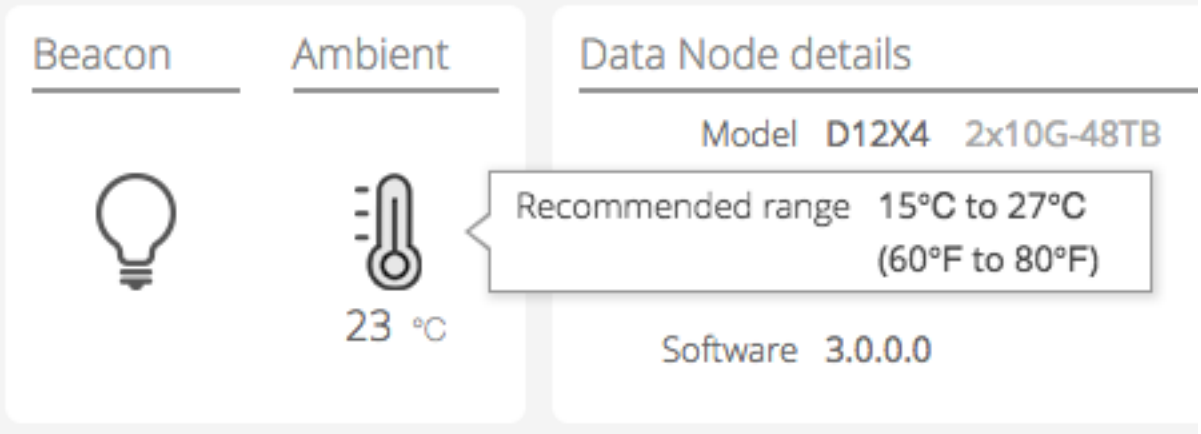

The GUI displays a thermometer icon and the ambient temperature value. The ambient temperature is measured at the front of the Data Node. It is the temperature of the exterior air that is drawn in by the Data Node fans. The acceptable operating temperature range is between range of 60°F to 80°F (15°C to 27°C). You can operate a Data Node for a short period of time when the ambient temperature is outside of the acceptable range, but you should contact Datrium Support to determine the course of action.

Data Node Fan Activity

During the first minutes of boot up the fans might change speed multiple times, including periods when they spin very fast and loud. Also when the DVX System detects certain hardware conditions the fans will spin very fast and loud. This mode of operation is precautionary and should be considered normal behavior unless it continues for longer than 7 minutes.

Data Node Response to Power Events

When there is a loss of power on both AC voltage inputs, the Data Node will switch to its own internal battery power and perform a graceful shutdown to ensure the highest level of persistent storage data integrity. The Data Node will initiate the shutdown only if both AC inputs experience a simultaneous power dropout or brownout.

Data Node Drive Replacement

When there is a drive failure, the DVX system recovers data from the drive, redistributes it to other drives, and then indicates that the drive is ready for replacement. Contact Datrium Support to arrange for a replacement drive. Do not remove the drive until you have received the replacement and can replace the failed drive immediately after you remove the failed drive. The Data Node is designed to operate with all components installed to provide adequate air flow.

> **Important:** After you remove the failed drive, you must insert the replacement drive within two minutes.

The DVX System uses the following LEDs to indicate drive replacement:

- The Node LED for Drive Replacement shows continuous amber.
- The fault LED on the failed drive module shows continuous amber.

You can use the CLI command `pool show` to display the state of individual drives. If you are using the GUI, the Hardware tab on the Data Node page shows individual drive icons. Place the mouse cursor over a drive icon to display the tooltip containing the drive status. Both the CLI and the GUI use the drive state “FAILED” to indicate that a drive is ready for replacement.

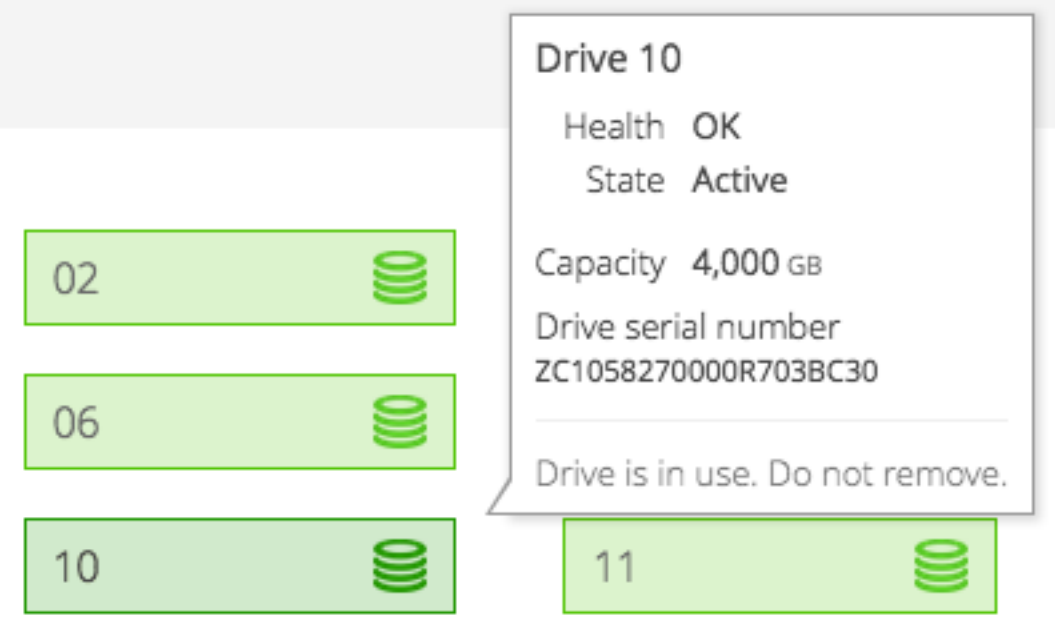

Data Node Drive Numbers

When replacing a drive, it is important to know the drive number that you are replacing. The drive numbers for each Data Node model type are as follows:

D12 Series

F24 Series

Power and Cooling Module Status

The PCM LED matrix diagrams listed below represent common situations. If you see a persistent LED display that is different than the ones shown, call Datrium Support. The diagrams are shown in the orientation of the PCM module on the left side of the back of the Data Node. The PCM module on the right side is turned up-side down.

Note: During transient states such as startup, reset, or upgrade operations, some of the PCM LEDs will turn on and off. After the operations are complete, the PCM LEDs will show a normal display.

PCM LED Patterns

LED Patterns	Notes
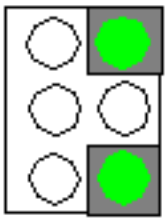	Normal operation – The upper right and lower right LEDs show continuous green.
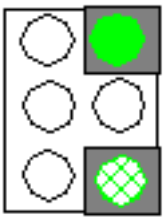	Battery Backup Unit is charging or in test mode. (Lower right LED flashes green.) • Battery charging can take up to one hour. For example, the Data Node might perform battery charging after a power loss. • The Data Node will test each of it's batteries several times a year, to ensure each battery is independently capable of supporting a power fail event for the entire Data Node. It will perform a

LED Patterns	Notes
	battery test only if there is sufficient battery capacity to support normal backup operation. It tests only one battery at a time so there is no loss of protection during the test. A test runs for approximately 11 hours per battery. • If the Data Node is exposed to more than a few power outages per month, it will run the battery test more often.
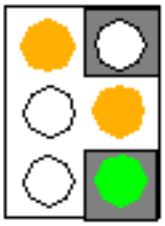	AC power loss to this PCM. • Check both ends of power cord. • Verify that the power switch is on. • Verify that the power source is operating correctly. If this does not resolve the power issue, check the DVX GUI or CLI for information. If you cannot resolve the situation, call Datrium Support.
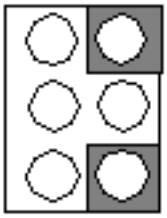	AC power loss to both PCM units. (All LEDs off.)

Network Port Status

Each network port has two associated LEDs. The link LED indicates whether there is a live connection on the port. The activity LED indicates link activity over the connection.

D12x4 Data Node Network Port LEDs

This image identifies the link and activity LEDs for a port of each type (1Gb Ethernet, 10Gb SFP+, 10Gb Ethernet).

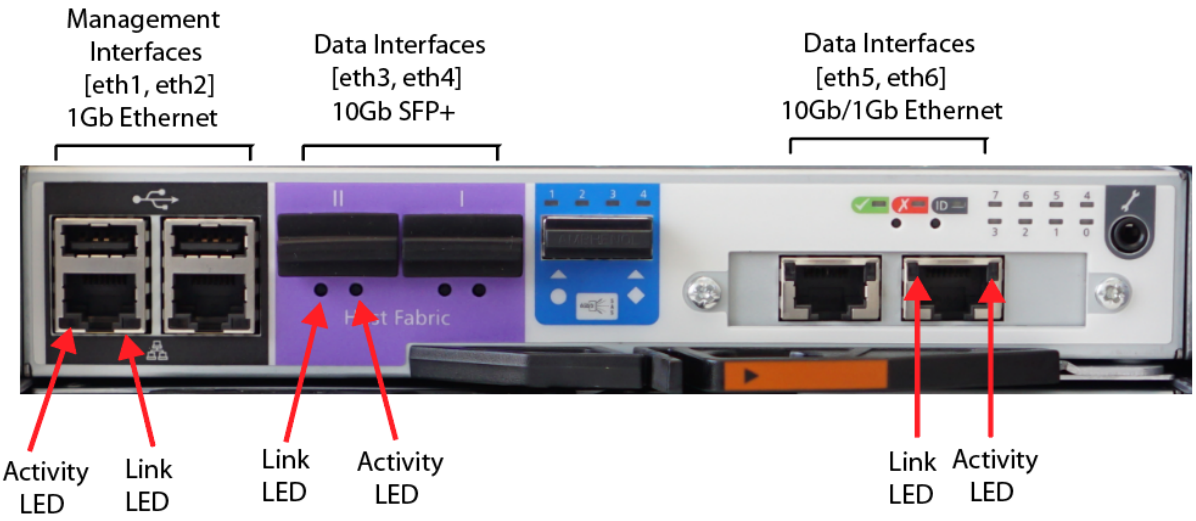

Data ports (eth3-6)

The Data Node has two SFP+ 10Gb ports and two RJ45 10G/1G Base-T ports.

- SFP+ ports (10Gb; Direct Attach copper or optical) – The link LED shows continuous green to indicate a live connection. The activity LED flashes amber to indicate link activity.
- 10G/1G Base-T (RJ45) – The link LED indicates both connectivity and speed. The link LED shows continuous green for a 10Gb connection, continuous amber for a 1Gb connection. Datrium recommends that you use a 10Gb connection for these ports. The activity LED (green) flickers off to indicate link activity.

Management port (eth1, eth2)

The management ports provide access for Datrium Support traffic, DVX GUI and CLI traffic.

- Management ports (eth1, eth2 – 1G/100Mb Base-T RJ45/USB) – The link LED indicates both connectivity and speed. The link LED shows continuous green to indicate a 1Gb

connection, continuous amber to indicate a 100Mb connection. Datrium recommends that you use a 1Gb connection for the management port. The activity LED (amber) flickers off to indicate link activity.

Port Status – DVX GUI

In the DVX GUI, the Hardware tab on the Data Node page shows a representation of the Data Node network ports. To display the port status, place the mouse cursor over a port icon to display the tool tip.

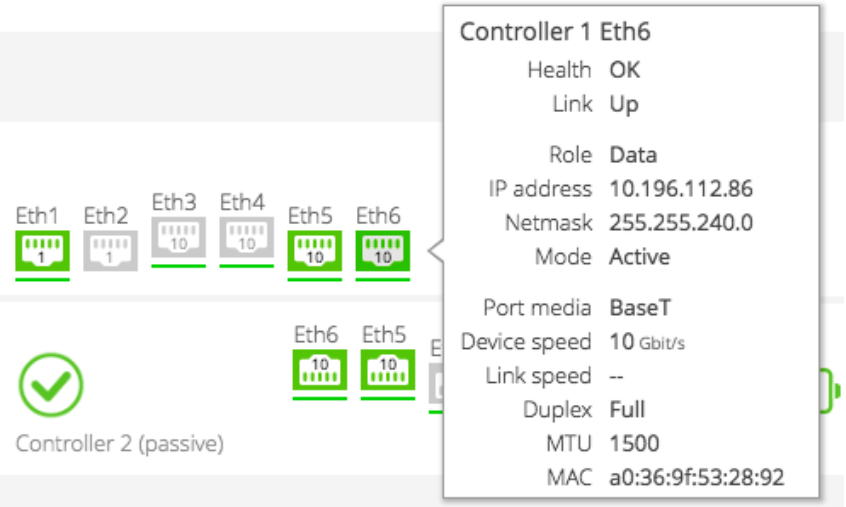

The Network tab on the Data Node page shows the network configuration, indicating the general network parameters and data and management configuration. This tab also shows the Data Node controllers and ports. The order of the port icons does not represent the actual ordering. The port icons are active – place the mouse over the icon to display the status tool tip.

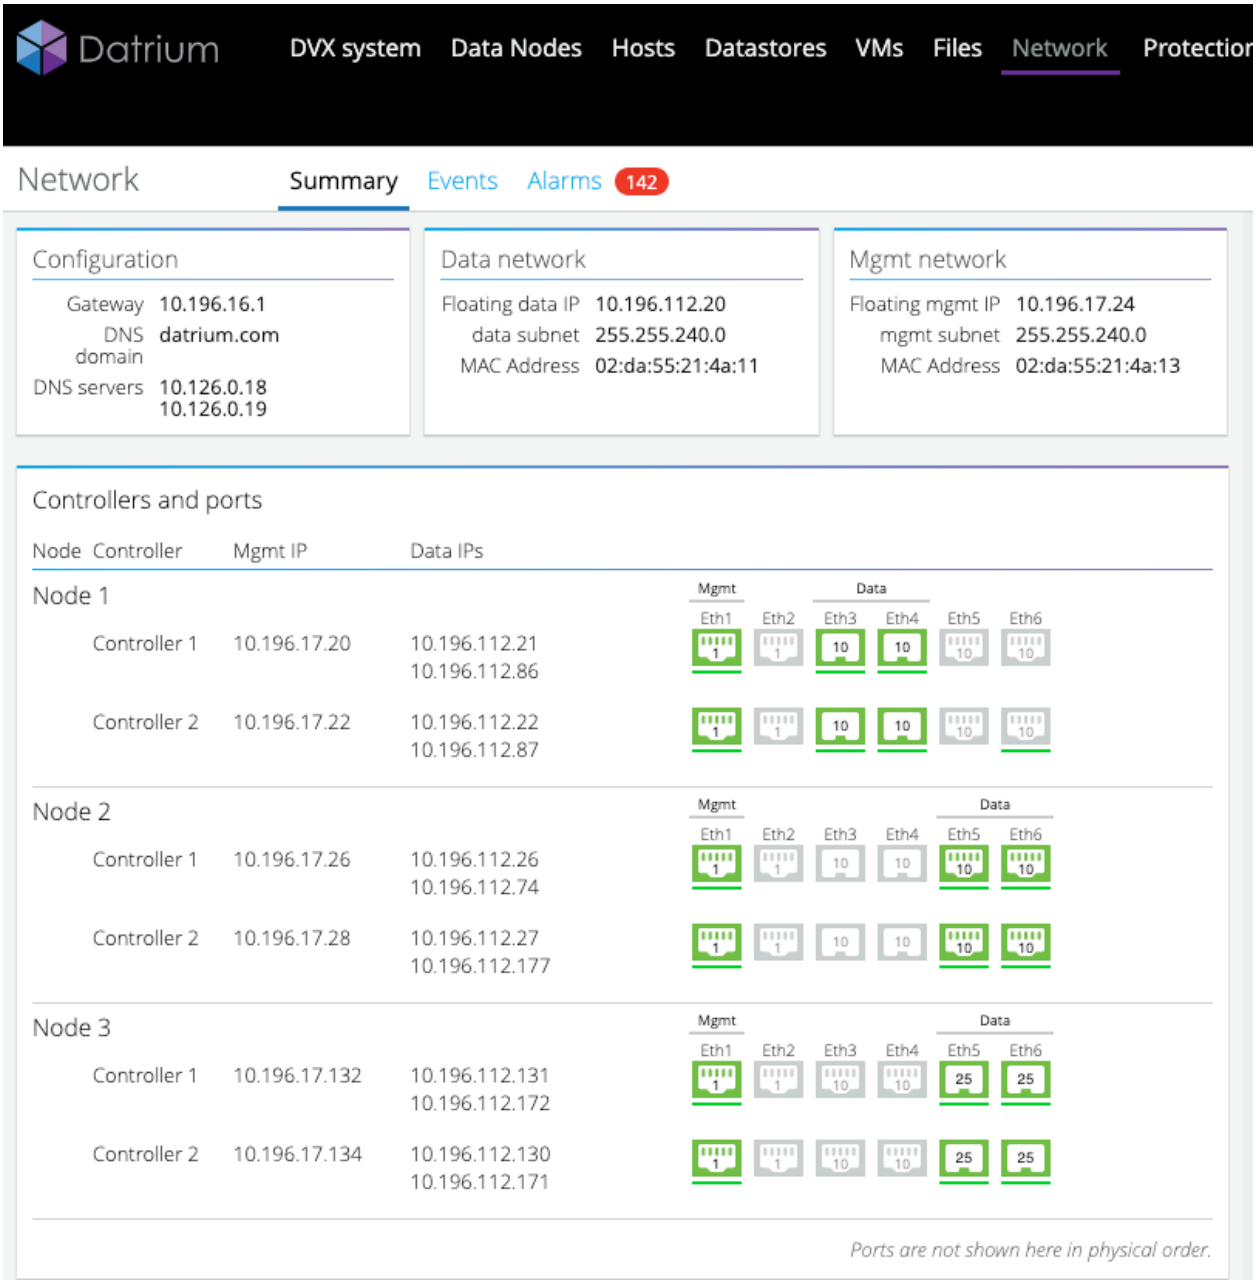

Port Status – DVX CLI

Use the `network show` command to display the network configuration and the port status.

```
>> network show

-------------------- Network ---------------------
Gateway      Domain name   DNS servers
------------ ------------- ----------------------
10.196.16.1  example.com   10.126.0.18 10.126.0.19
------------ ------------- ----------------------

----------- Data -----------
Floating IP    Data netmask
-------------- -------------
10.196.112.20  255.255.240.0
-------------- -------------

---------- Mgmt -----------
Floating IP   Mgmt netmask
------------- -------------
10.196.17.24  255.255.240.0
------------- -------------


------------------------------------------------- Ports -------------------------------------------------
Port                    Role   Mode    Link   Media      Speed(Gbps)   MTU    MAC address        IP address
----------------------- ------ ------- ------ ---------- ------------- ------ ------------------ -------------
floating.data           Data   --      --     --         --            1500   02:da:39:5f:79:2f  10.196.112.20
floating.mgmt           Mgmt   --      --     --         --            1500   02:da:39:5f:79:31  10.196.17.24
node1.controller1.eth1  Mgmt   --      True   1G  BaseT  1.0           1500   00:50:cc:7a:84:b7  10.196.17.20
node1.controller1.eth2  --     --      False  1G  BaseT  --            1500   00:50:cc:7a:84:b8  --
node1.controller1.eth3  --     --      True   10G SFP+   10.0          1500   00:50:cc:7a:84:bc  --
node1.controller1.eth4  --     --      True   10G SFP+   10.0          1500   00:50:cc:7a:84:bb  --
node1.controller1.eth5  Data   Active  True   10G BaseT  10.0          1500   a0:36:9f:53:28:90  10.196.112.21
node1.controller1.eth6  Data   Active  True   10G BaseT  10.0          1500   a0:36:9f:53:28:92  10.196.112.86
node1.controller2.eth1  Mgmt   --      True   1G  BaseT  1.0           1500   00:50:cc:7a:7a:01  10.196.17.22
node1.controller2.eth2  --     --      False  1G  BaseT  --            1500   00:50:cc:7a:7a:02  --
node1.controller2.eth3  --     --      True   10G SFP+   10.0          1500   00:50:cc:7a:7a:06  --
node1.controller2.eth4  --     --      True   10G SFP+   10.0          1500   00:50:cc:7a:7a:05  --
node1.controller2.eth5  Data   Active  True   10G BaseT  10.0          1500   a0:36:9f:52:1a:b4  10.196.112.87
node1.controller2.eth6  Data   Active  True   10G BaseT  10.0          1500   a0:36:9f:52:1a:b6  10.196.112.22
----------------------- ------ ------- ------ ---------- ------------- ------ ------------------ -------------
```

Data Node LEDs

The DVX Data Node uses a set of LED lights on the front and back panels to indicate the Data Node operating status. The LED information complements the information that you can obtain from the DVX GUI or from the DVX CLI.

- In the DVX GUI, use the events and alarms tabs on the Datrium home page. The GUI also displays a list of recent alarms in the status pane on the right side of the GUI display.
- Use the DVX CLI commands `alarms show`, `events show`, and `pool show` to display status.

The following sections describe the LEDs on the Data Node front and back panels.

Data Node Front Panel

The following image shows the location of LEDs on the front of the Data Node.

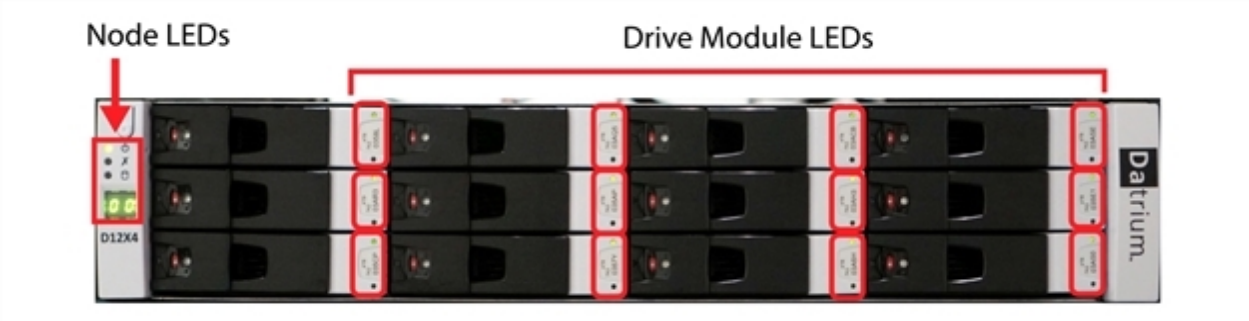

The front panel has two types of LED groups.

- Node LEDs – The LEDs on the left side of the front panel indicate the general status of the Data Node. There is also a numeric display for Data Node identification. This is the node number displayed in the DVX UI.
- Drive Module LEDs – Each drive module has two LEDs. The upper LED indicates power status and activity. The lower LED is a fault indicator.

Node Status

The left-hand panel on the front of the Data Node has three LEDs for status and a numeric display that the DVX uses as a locator beacon. The following describes the node LEDs. The picture of the panel shows the node LEDs during normal operation.

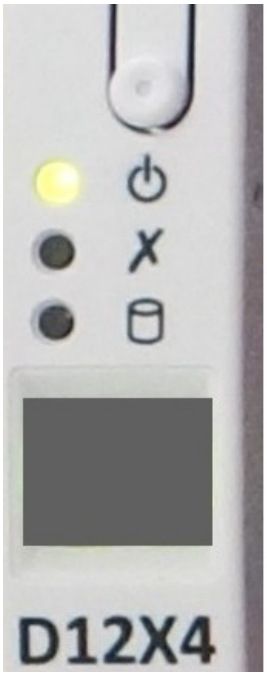

Node LEDs

LED / Color(s)	Description
Power (Power state icon) Green/Amber	Green indicates that the Data Node is powered-up. Amber indicates that the Data Node is in the process of powering up or has completed a graceful shutdown. If the LED is off, there is no available power.
Amber LED (“X” icon) Amber	Amber indicates a hardware fault. There may also be an associated fault LED display on a PCM or controller to identify the faulty module.
Drive Replacement (Disk icon) Amber	Indicates that one or more drive modules are ready to be replaced. In addition, the faulty drive will show a lighted

LED / Color(s)	Description
	amber fault LED.
Node ID Green	Displays “00” when you power on the Data Node. At the end of software startup, the Data Node displays the node ID. During locator beacon operation, the node ID flashes.

Drive Module Status

Each Drive Module has two LEDs that indicate power, activity, and fault conditions. The following image describes the Drive Module LEDs.

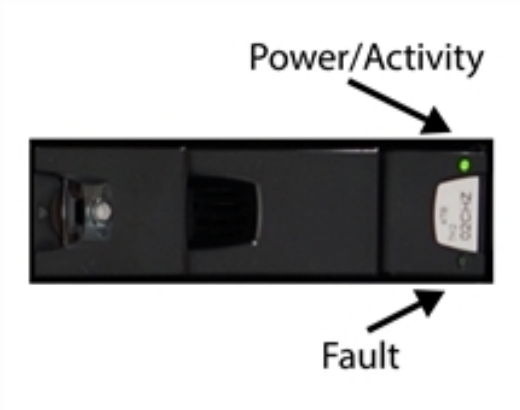

Drive Module LEDs

LED	Description
Power/Activity	Continuous green indicates that the drive module is powered-up. The LED flickers off to indicate activity. In some fault conditions, the Data Node turns the LED off.
Fault	Continuous amber indicates that you should

LED	Description
	replace the drive module. When a drive fails, the Data Node also shows continuous amber for the Node drive replacement LED.

Data Node Back Panels

The Data Node uses LEDs on the back panels to provide an indication of Data Node operation and status. The back panel contains the following LED groups:

- PCM (Power Cooling Module) – Indicates power, fan, and battery module status.
- Notification – Indicates Data Node operational status.
- Boot activity – Used for startup operations.
- Network port status – Indicates port connectivity and link activity.

The following image shows the location of LED groups on the back the Data Node.

D12x4 Data Node (Back) LED Groups

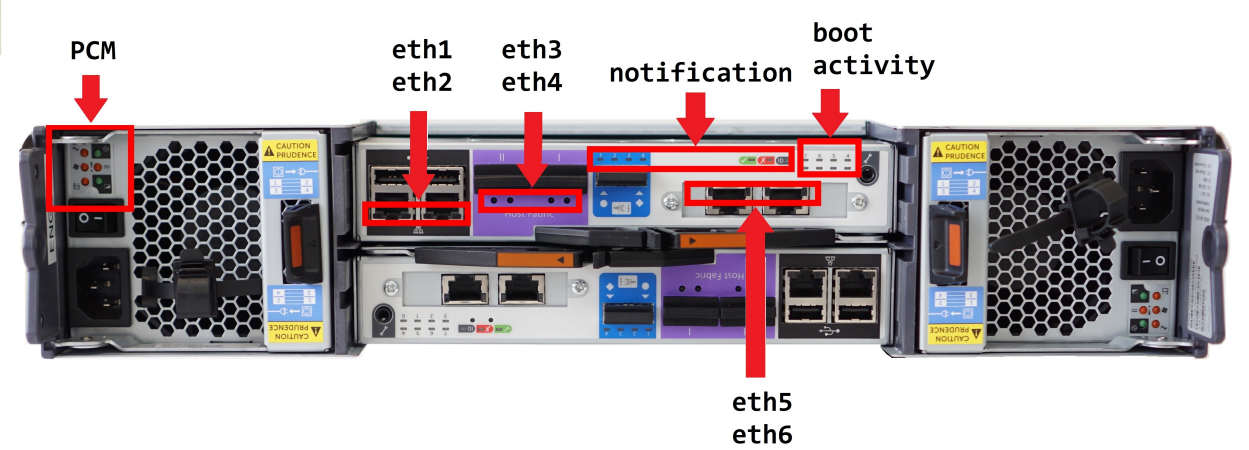

Controller Module Status

The Data Node uses the Notification LED group to indicate Controller Module status and to indicate whether the Controller is operating as the active Controller for the Data Node.

Controller Module Status LEDS – D12x4B,D12x4C/F12x2B,F12x2D

The following image shows the LED group in the orientation of the upper Controller Module. The lower Controller Module is turned up-side down.

LED	Color	Description
1	Green	Operating Status – The Data Node software has begun its startup process.
2	Amber	Not used.
3	Green	LED flashes during locator beacon operation.
4	Amber	LED flashes during locator beacon operation.
Power	Green	Power state icon. Indicates that the controller is fully powered. The LED flashes during normal power-up and graceful shutdown. In some fault conditions, the Data Node turns the LED off.
Fault	Amber	“X” icon. Indicates a fault or attention condition. The LED flickers during normal power-up or graceful shutdown.
Active	Blue	“ID” label. Indicates that the controller is currently functioning as part of the DVX system. **Note**: You should never remove a controller if the Active Controller LED is blue.

Controller Module Status LEDs – D12x4

The following image shows the LED group in the orientation of the upper Controller Module. The lower Controller Module is turned up-side down.

LED	Color	Description
1	Green	Operating Status – The Data Node software has begun its startup process.
2	Green	Not used.
3,4	Green	LEDs flash during locator beacon operation.
Power	Green	Power state icon. Indicates that the controller is fully powered. The LED flashes during normal power-up and graceful shutdown. In some fault conditions, the Data Node turns the LED off.
Fault	Amber	“X” icon. Indicates a fault or attention condition. The LED flickers during normal power-up or graceful shutdown.
Active	Blue	“ID” label. Indicates that the controller is currently functioning as part of the DVX system. **Note**: You should never remove a controller if its Active Controller LED is blue.

Transient States

During transient states such as startup, restart, reset, or upgrade operations, the device performs a sequence of operations before starting normal operations. The startup procedure can take up to eight minutes. There might be extreme cases that require a few minutes longer. The actual time taken depends on the following factors:

- Cold boot or warm restart – During a warm restart the Data Node performs only a part of the cold boot sequence. A warm restart generally requires less time.
- Data Node state after the previous shut down – If the Data Node did not shut down gracefully, it will perform additional tasks to achieve a consistent state that can support normal operation

Controller Boot Activity LEDs

During the hardware boot procedure, the Data Node LEDs reflect the boot activity. The boot activity LEDs on the back of the Data Node will flash in various patterns as it performs hardware tests and initialization. This might take up to six minutes. While the boot activity LEDs are on or flashing, you can ignore all other LED activity. When hardware initialization is complete, the Data Node turns off the boot activity LEDs and starts the Data Node software.

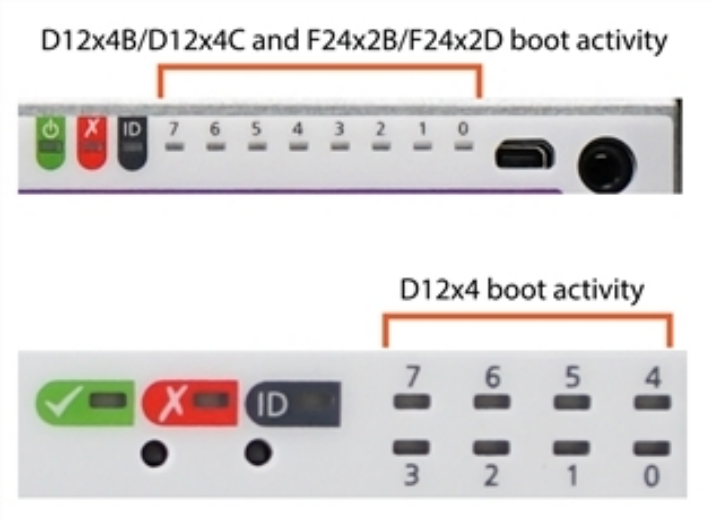

Relocating a Data Node

A Data Node weighs approximately 70lb (32kg). Do not try to move it by yourself. Make sure that the target location is ready for installation of the Data Node (rack brackets, available and sufficient power, available Ethernet connections).

Before you shut down the Data Node, unmount the Datrium datastores on the DVX hosts to prevent APD events. After you relocate the Data Node and then restore power to it, the Data Node will reboot. When the Data Node is operational again, use the DVX GUI to remount the Datrium datastore on the hosts.

To relocate a Data Node.

1. Power off the Data Node.
2. Log in to the Data Node. If you have a separate management network, use management interface floating IP address. Otherwise, use the floating IP address for the data interface.

   ```
   ssh admin@floating-ip
   password: datrium#1
   ```

3. Power down the Data Node.

   ```
   nodes powerof
   ```

4. After the shut down completes, the Data Node will show the following LED display.

   PCM LEDs – On both PCMs, the upper right LED flashes green. The lower right LED shows continuous green. The lower left LED shows amber.

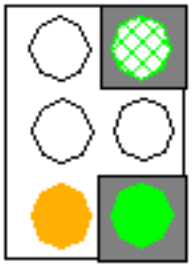

Controller Notification LEDs
(both controllers):
– LED#1 is turned off.
– Power status LED flashes.
– Fault LED flickers.

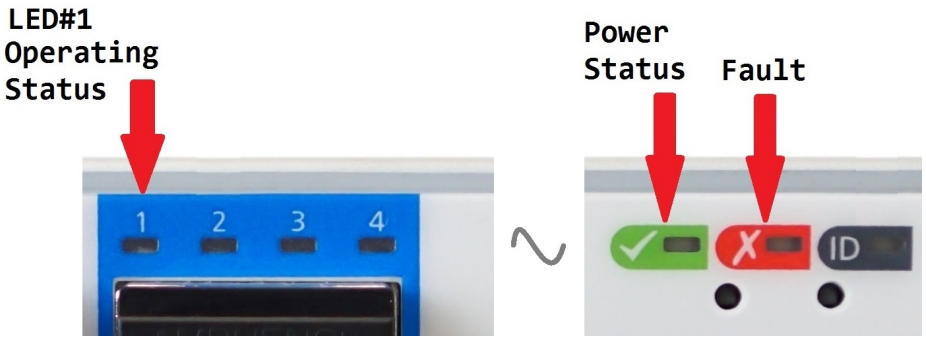

5. Turn off the PCM power switches and unplug both PCM modules. Make sure to wait until the LEDs indicate that shutdown has finished.
6. To make it easier to carry, remove the PCM and controller modules.
7. Install the Data Node chassis in the new location: Insert the PCM and controller modules; connect power to both PCM modules; connect the controllers to the network.
8. Extract the chassis.
9. Turn on the PCM power switches. When power is applied to the Data Node, it will boot automatically.
10. During the first minutes of boot up the fans might change speed multiple times, including periods when they spin very fast and loud. Also when the DVX System detects certain hardware conditions the fans will spin very fast and loud. This mode of operation is precautionary and should be considered normal behavior unless it continues for longer than 7 minutes.

DVX System Audit Trail

The DVX System maintains an audit trail of user operations generated by the DVX UI elements – DVX GUI, DVX Web Client GUI, DVX CLI, and DVX API. The System captures events that correspond to the beginning of operations. In the case of login operations, the DVX System also captures login failures.

To display an audit trail in the DVX GUI – on the DVX GUI Dashboard or Real Time view, click on the Events tab and then click on Audit.

DVX GUI – Audit Trail

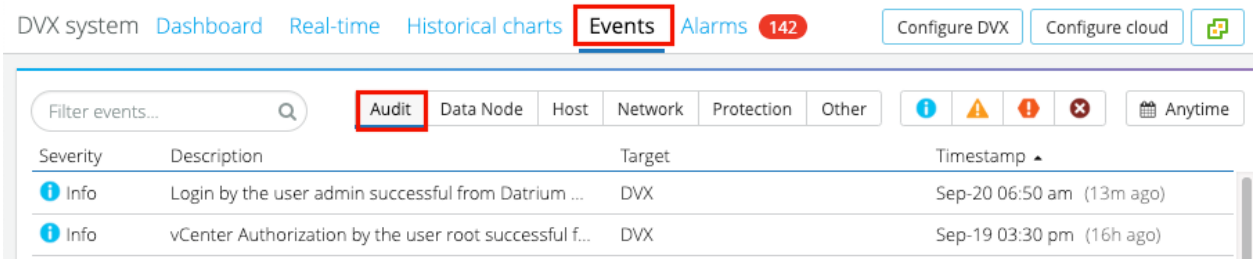

To display an audit trail in the DVX CLI – specify the event type “AuditEvent” with the events show command.

DVX CLI – Audit Trail

```
dvx05.node1.controller2>> events show --event-type AuditEvent
------------------------------------------------------------------------------
            Time stamp  Details
------------------------ -----------------------------------------------------
2018-02-22T23:05:19 UTC  'vms snapshots clone' by user 'admin' from CLI running at 10.
2018-02-22T23:05:09 UTC  'vms snapshots clone' by user 'admin' from CLI running at 10.
2018-02-22T23:05:02 UTC  'vms take_snapshot' by user 'admin' from CLI running at 10.2.
2018-02-22T16:10:04 UTC  'alarms clear' by user 'root' from vSphere Datrium Plugin run
```

DVX System Network Management

The DVX Data Node is a high-bandwidth network storage appliance that provides adaptive pathing networking for aggregate data bandwidth, management port redundancy, and controller failover capabilities. The Data Node supports the use of separate data and management subnets.

This section contains the following topics:

DVX Network Features

Adaptive pathing networking	The DVX System provides adaptive pathing networking for increased data traffic capacity over teamed 10G data interfaces on the Data Node. The DVX System can determine the viability of network paths and it will spread traffic across interfaces.
Port redundancy	The Data Node provides port redundancy for both data and management network interfaces. Data port redundancy – The DVX System uses the aggregate bandwidth of two or more data ports. If there is a network failure on one of the ports, the DVX System will continue to use the bandwidth that is available on the remaining active interface(s). Management port redundancy – The DVX System supports the use of the two management ports in an active/passive bonded port configuration.
Controller failover	The Data Node supports controller failover in response to controller failure or loss of Data Node or Compute Node connectivity.

	Controller failure – The Data Node has a pair of controllers that operate in active/passive configuration for high availability. If the active controller fails, the DVX System promotes the passive controller to active status and uses the newly active controller to continue to serve data requests. Data Node connectivity – With a Data Pool of more than one Data Node, a DVX System supports dynamic controller failover based on connectivity between the Data Nodes. The DVX System will fail over controllers on Data Nodes to maintain connectivity between all Nodes in the Data Pool. Compute Node connectivity – The DVX System supports dynamic controller failover based on the following criteria: • The passive controller must have connectivity to more Compute Nodes than the active controller. • The passive controller set of Compute Nodes must include the same set of Compute Nodes that are connected to the active controller. • Connectivity between Data Nodes in the Data Pool is not compromised. If all three of these conditions are satisfied, the DVX System will failover to the passive controller.
Data Node Failover	Data Node Fault Tolerance provides failover support in the event of a Data Node failure in a DVX Data Pool consisting of 3 or more Data Nodes. When Data Node Fault Tolerance is activated, the DVX System will continue serving and storing data normally in the following network failure situations: • Network interface on one controller fails, where connectivity is lost between one Data node controller and the rest of the network. • Network interface on a single Data Node fails, and node is isolated on the network. • Network path between two Data Nodes fails. • Network path between hosts and a Data Node controller fails.

	• Switch fails where nodes are separated to multiple groups, or all Data Node became isolated.
Data and management subnets	The Data Node provides 10G data ports and 1G management ports. The Data Node is designed to support separate subnets for data and management.

Data Node Network Status and Configuration Display

The Hardware tab in the Data Nodes view shows a graphical representation of the back of the Data Node.

The GUI uses icons to represent the ports and color to indicate port status. Green indicates that there is link activity and Compute Node connectivity.

To display port information, place the mouse cursor over a port icon to display a tool tip. The port information includes the following information:

Category	Possible Values	Category	Possible Values
Health	OK, Warning, Critical, Unknown	Media type	BaseT, SFP+
Link	Up, Down	Device speed	25Gbits, 10Gbits, 1Gbit
Role	Data, Mgmt	Link speed	25Gbps, 10Gbps, 1Gbps
IP address	address	MTU	1500, 9000
netmask	mask	MAC	address
Mode	Active, Passive (mgmt)		

DVX GUI – Data Node Network Interfaces

The following image shows a tool tip for the eth4 data interface.

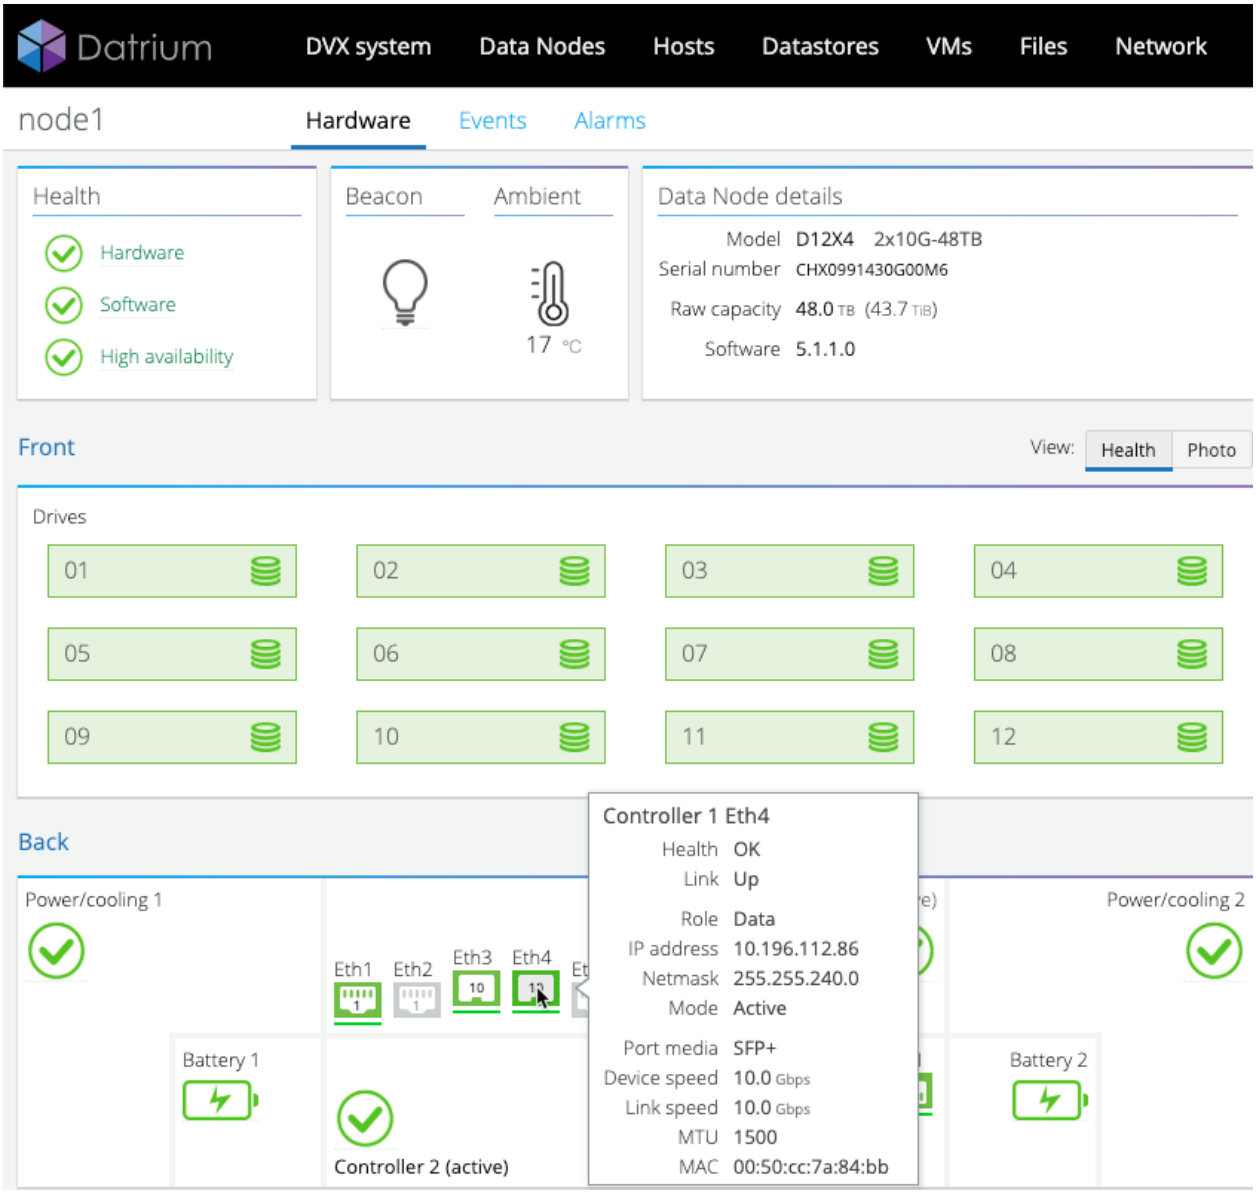

To see the network configuration in the DVX GUI, click on the Network tab, which shows the following information:

Category	Information
Configuration	Gateway Domain DNS servers
Data network	Floating IP address

Category	Information
	Netmask Interfaces (eth3+eth4, eth5+eth6, or eth5+eth6+eth7+eth8) Teaming (Active-active) MAC address
Management network	Floating IP address Netmask Interfaces (eth1, eth2, or eth1+eth2) Teaming (None) MAC address
Controllers and ports	Node information for each Node in the Data Pool: IP addresses for data ports IP address(es) for management port(s) Port information (mouse over for tooltip)

DVX GUI – Network Configuration

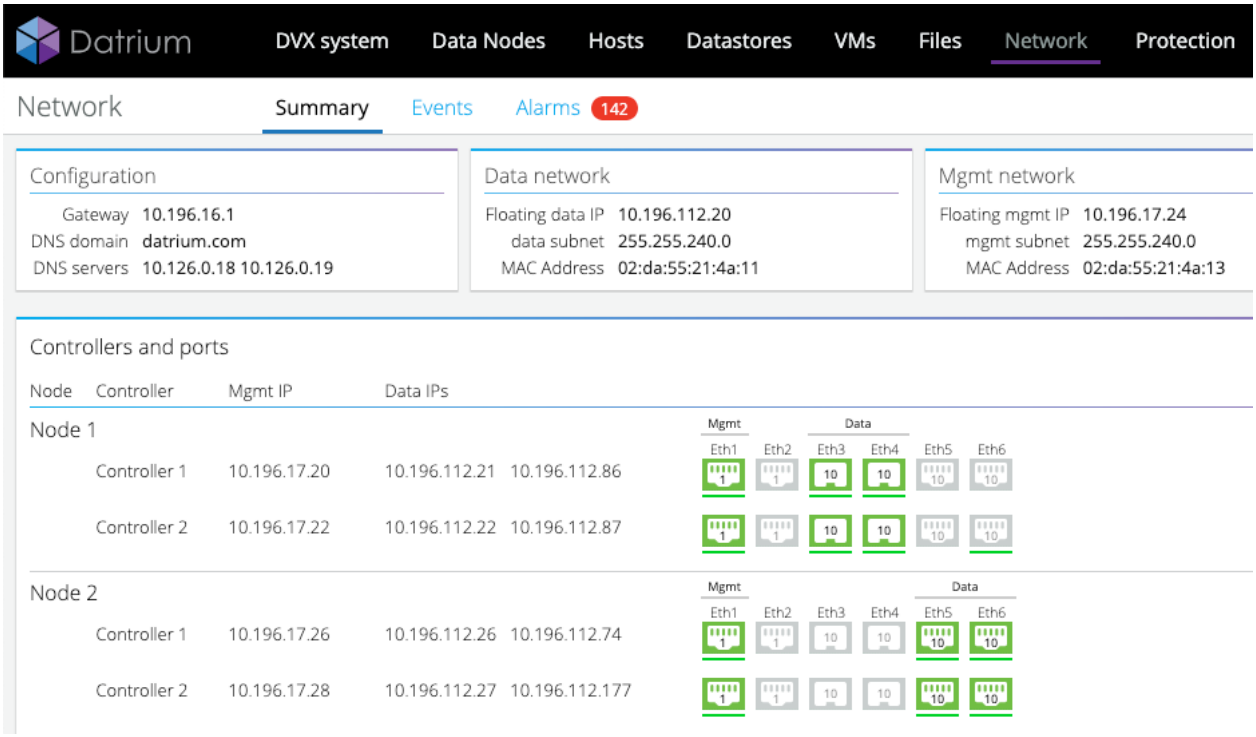

The DVX CLI `network show` command displays the following information:

Category	Information
Network	Gateway Domain name DNS servers
Data	Floating IP Data netmask Interfaces (eth3+eth4 or eth5+eth6) Teaming (Active-active)
Management (Mgmt)	Floating IP Mgmt netmask Interface(s) (eth1, eth2, eth1+eth2)

Category	Information	
	Teaming (None)	
Ports	Role (Data, Mgmt) Mode (Active, Passive) Link (True, False) Media (25G SFP+, 10G BASET, 10G SFP+, 1G BASET)	Speed, Gbps (25.0, 10.0, 1.0) MTU(1500, 9000) MAC address IP address

Configuring the Network

The following information applies to an existing Data Node network configuration. You can use the DVX GUI or the DVX CLI to configure the network. For information about installing the DVX System and creating the initial network configuration, see the *Datrium DVX Software Configuration* guide.

- DVX GUI – Network Configuration
- DVX CLI Network Commands
- Changing DVX System IP Addresses
- Separating Data and Management Traffic
- Mounting a Different Datrium Datastore
- Effects of Data Node Network Configuration

DVX GUI – Network Configuration

To modify the network settings, click on the “Configure network” entry on the system menu.

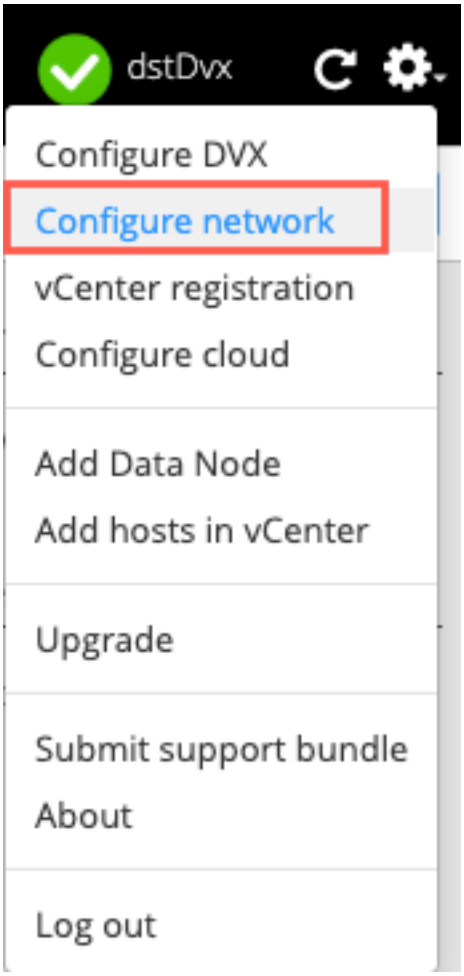

The GUI displays a dialog where you can modify the following settings:

- Global network settings – gateway, Data Node domain, and DNS servers.

  The DVX System uses the gateway for access to the Datrium Support server at autosupport.datrium.com. The gateway must be on the same subnet as the Data Node interface that you are using for management traffic.

- Data and management network settings – floating IP addresses and subnet masks

  See also Changing the Data Floating IP Address.

- Data Node network interfaces and corresponding port IP addresses

  See also Separating Data and Management Traffic.

DVX CLI Network Commands

The following table summarizes the DVX CLI network commands.

network	commit-setup	Commit a new network configuration.
	show	Displays the Data Node network configuration.
	setup	Configure Data Node IP addresses and the DVX network environment (domain name, gateway, DNS servers, and jumbo frames for data traffic).
	setup prepare-cluster	If you have moved your Data Nodes to another network, you may need to establish controller communication between Data nodes. You can run the network setup command with this option to establish Data Node communication before resetting the network configuration. You only need to run this command if you run `network setup`, and the command times out.
	tools	ping, traceroute, and nslookup commands.
	test	Run a set of network access tests.

Changing DVX System IP Addresses

It may be necessary to change one of the DVX System IP Addresses for any number of reasons – for example, a network topology change or moving a DVX System to a different site. The following sections describe how to change a host IP address or one of the Data floating IP addresses.

It may also be necessary to update IP addresses for the gateway or DNS Server(s) in your DVX network configuration. You can use either the DVX CLI `network setup` command or the DVX GUI.

Changing the ESXi Host IP Address

To change the host IP address, use the following procedure so that the DVX System recognizes the change.

1. Use the vSphere Web Client to unmount the Datrium datastore. (Mounting/Unmounting a DVX Datastore.)
2. Change the host IP address. In the vSphere Web Client:
3. Click on Host → Manage.
4. Click on the “Networking” tab.
5. Click on VMkernel adapters.
6. Select the appropriate VMkernel adapter for Data Node data traffic.
7. Click on the pencil icon to edit the adapter settings.
8. Click on IPv4 settings and update the network IP information for the new data network.
9. Use the DVX GUI to remount the Datrium datastore.

Changing the Data Floating IP Address

When you change the data floating IP address in your Data Pool, after the change *all controllers* on all Data Nodes must be failed over before you can resume normal operations with the DVX system. You can use the DVX GUI or CLI for this task.

Use the following procedure to change the data floating IP address for either the data or management interfaces:

1. Use the DVX GUI to unmount the Datrium datastore. (Mounting/Unmounting a DVX Datastore.)
2. Log in to each Compute Node and stop the Datrium service.

   ```
   localcli datrium stop
   ```

3. Change the data floating IP address on the Data Node. You can use either the DVX CLI `network setup` command or the DVX GUI. When you change the floating IP address for the Data Node interface that you use for management traffic, the DVX

System will automatically re-register the vCenter Server.

4. After you change the floating IP address, you need to run the DVX CLI `network setup` command or the DVX GUI. When you change the floating IP address for the Data Node interface that you use for management traffic, the DVX System will automatically re-register the vCenter Server.
5. Log out from your vCenter Server and then log in again.
6. Next, you need to failover both controllers on every Data Node in your Data Pool. You can perform this task from either the DVX GUI or CLI.
7. From the DVX GUI:
    a. Ensure that the Health of all Data Nodes shows as green in the Data Nodes summary page.
    b. For each data node in the Data Nodes summary page:
    c. Click on the data node to view details.
    d. On any controllers that are marked "active" or "pool active", point your cursor over the green check mark and click on the "Failover controller" link to trigger failover on the controller.
    e. Wait for 2 minutes.
    f. Repeat on ALL Data Nodes.
    g. After these steps are done on all Data Nodes, return to the Data Nodes summary page in the DVX GUI and wait until the Health status of all data nodes becomes green.
8. OR, from the DVX CLI:

    a. Ensure that Health of all nodes is OK and HA state is Redundant by using the `nodes show` command and viewing the output.
    b. For each Data Node in the Data Pool, run the `nodes failover` command nodeX.controller1. E.g.,

    ```
    nodes failover nodeX.controller1
    ```

    c. Wait about 2 minutes.

    d. Run the nodes failover on the next controller on the next Data Node. E.g.,

    ```
    nodes failover nodeX.controller2
    ```

e. Wait about two minutes.

f. Repeat for all controllers on all Data Nodes.

g. After these steps are carried out on all Data Nodes, run the `nodes show` command and wait until the Health of all Data Nodes is OK and the HA state is Redundant.

9. Last, use the DVX GUI to remount the Datrium datastore.

Separating Data and Management Traffic

As of version 3.0, the DVX System requires that you configure both data and management interfaces. If you have been using combined data and management traffic over the data interface with previous versions of the System, use the following procedure to change to using separate data and management interfaces.

If you intend to continue to use the same IP address for the Data Node data interface, you only need to configure the management interface. Use the network setup command to invoke the setup wizard.

```
network setup
```

If you intend to change the IP addresses for the Data Node data interface when you separate the data and management traffic, use the following procedure:

1. Configure the management interface. (See above.)
2. Use the procedure Changing the Data Floating IP Address.

Mounting a Different Datrium Datastore

To change the Data Node that an ESXi host uses for virtual machine storage, use the following procedure:

1. If the target Data Node is not registered with the vCenter associated with the ESXi host, log in to the Data Node and register it with vCenter:

   ```
   config vcenter register ipAddr| dnsName -user userName
   [ -password password]
   ```

The vCenter Server might not load the plug-in. You must restart the vSphere Web Client service to force the Web Client to load the plug-in. Use the following command to restart the Web Client:

```
service vsphere-client restart
```

2. Use the vSphere Web Client to unmount the datastore.
3. Log in to the host and stop the Datrium service.

   ```
   localcli datrium stop
   ```

4. Make sure the host and Data Node are both recognized by the vCenter Server.
5. Remount the datastore. To remount the datastore, use the DVX GUI.

Network Traffic Payload Size (Jumbo Frames)

The DVX System supports the use of jumbo frames for network traffic over the Data Node data interfaces. If you use jumbo frames, every device on the network path between the Data Node and the host must be configured for jumbo frame traffic. This includes:

- Data Node data interface.
- Switches between the Data Node and the hosts.
- VMkernel NICs on ESXi hosts.

To enable or disable jumbo frame payloads, use the CLI command “network setup”. This command runs a wizard that includes network configuration.

- Enable jumbo frames to use a Maximum Transmission Unit (MTU) of 9000 bytes on the Data Node data interface.
- Disable jumbo frames to set the MTU to 1500 bytes.

Effects of Data Node Network Configuration

Data Node network configuration causes a brief interruption in data service. If the resulting Data Node port or IP address is different than the previous configuration, the Data Node restarts all data services to provide access through the newly configured interface. If you have an existing CLI session, you must start a new CLI session.

A change in the Data Node network interface may result in a change in the datastore on the ESXi host(s) that use the Datrium datastore. The Datrium mount point definition includes an IP address or DNS name that identifies the Data Node. This IP address is the floating IP address for the data interface.

- If your mount point definition uses a DNS name, then the DNS server will handle any change to the Data Node network interface. You can change the Data Node interface and no further actions are required.

- If you are using a DNS name in the Datrium mount point definition, the stability of the Datrium mount point depends on the consistency of the DNS server operation. If the DNS server does not provide consistent results, the DVX System may be unable to resolve the DNS name.
- If your mount point definition uses an IP address, a change to the Data Node network interface will create a new datastore definition. This affects any virtual machines on the datastore.

Use the following procedure to ensure that the virtual machines will use the correct datastore.

1. Power off the virtual machines and remove them from the vSphere inventory.
2. Before you change the Data Node floating IP address, unmount the Datrium datastore and unregister the vCenter Server (config vcenter unregister).
3. Stop the Datrium service.

   ```
   localcli datrium stop
   ```

4. Use the network setup wizard to change the Data Node network interface. See the description of network setup in the *Datrium DVX Software Configuration* manual.
5. Register the vCenter Server (config vcenter register).
6. Update the DVX system mount point definition on the ESXi host(s) so that it uses the correct IP address. (add host procedure)
7. Add the virtual machines to the vSphere inventory.

Access to the Internet Gateway

The DVX System uses the management interface for access to the gateway. Use the following CLI command to verify access to the gateway and the Datrium Support server.

```
support test
```

If there is a problem, you might see an HTTP error:

- If you are using a firewall, make sure that you have configured access appropriately.
- If you are using an HTTP proxy in your environment, use the `config web proxy set` command to identify the proxy server and port for HTTP access.

DVX Network Concepts

The Data Node uses teamed 10G or 10G/25G ports for data traffic, depending on controller type, and a single 1G port or a bonded pair of 1G ports for management traffic.

Data traffic includes all communication between the Data Node and the DVX Hyperdriver on a host.

Management traffic includes Datrium Support traffic, and DVX GUI and CLI traffic. The management interface must be on the same subnet at your internet gateway. Datrium Support traffic requires a gateway for access to the Datrium Support portal (https://support.datrium.com).

The following sections provide information about DVX networks.

Data Node Network Ports

The following sections describe the network ports for the different controller types.

For information about monitoring link activity, see Network Port Status.

D12x4B / D12x4C/ D12x10D / F24x2B / F24x2D 2x25G Data Node Network Ports

The D12x4 2x25G model Data Nodes (D12x4B and D12x4C), D12x10D model Data Node, and the F24x2 2x25G model Data Nodes (F24x2B and F24x2D) support 10G or 25G data traffic.

- To configure the data interface, you select a teamed pair of interfaces of the same media type – 10BASE-T (eth3+eth4) or SFP+/SFP28 (eth5+eth6; these interfaces support 10G or 25G traffic).
- To configure the management interface, you select a single interface (eth1 or eth2) or the bonded pair of interfaces (eth1+eth2).

The following table shows the port names and types. Ports are referenced by interface name – eth1..eth6. The port names are shown for controller 1. Controller 2 is turned upside down in the chassis.

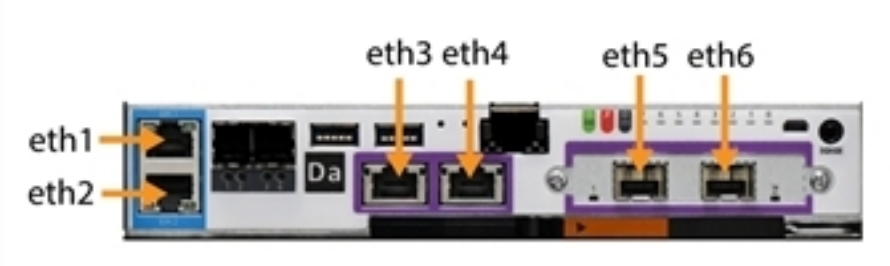

CLI Port Name	GUI Port Name	Port Type
node1.controller1.eth1 node1.controller1.eth2	Controller 1 Eth 1 Controller 1 Eth 2	1G RJ45 Ethernet. Used as management ports for the DVX GUI and CLI, and for Datrium Support remote access.
node1.controller1.eth3 node1.controller1.eth4	Controller 1 Eth 3 Controller 1 Eth 4	10G BASE-T RJ45 Ethernet ports used as data ports.
node1.controller1.eth5 node1.controller1.eth6	Controller 1 Eth 5 Controller 1 Eth 6	10G SFP+/25G SFP28 ports used as data ports. These ports accept direct attach twin axial copper connections or optical

CLI Port Name	GUI Port Name	Port Type
		connections. The SFP interfaces support a data rate according to the capabilities of the switch and the cable that connects the Data Node to it. To support speeds up to 25G with twin-axial copper connections, both the cable and switch must be rated for 25G traffic. If either the cable or switch operates at 10G, the Data Node will operate at 10G. Datrium offers 10G/25G dual speed optical transceivers as an option for these 25G-capable Data Nodes.

D12x4B 4x10G Data Node Network Ports

The D12x4B 4x10G Data Nodes support 10G data traffic.

- To configure the data interface, you select a team of interfaces of the same media type – 10BASET (eth3+eth4) or SFP+ (eth5+eth6, eth7+eth8, or eth5+eth6+eth7+eth8).
- To configure the management interface, you select a single interface (eth1 or eth2) or the bonded pair of interfaces (eth1+eth2).

The following table shows the port names and types. Ports are referenced by interface name – eth1..eth8. The port names are shown for controller 1. Controller 2 is turned upside down in the chassis.

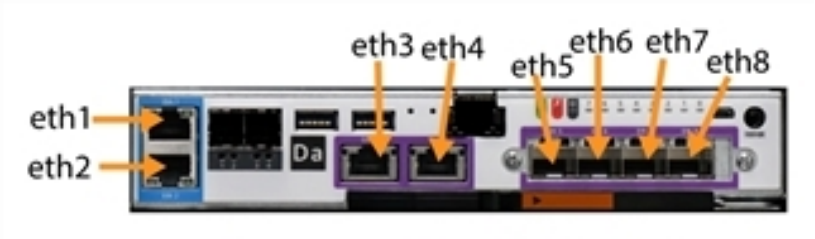

CLI Port Name	GUI Port Name	Port Type
node1.controller1.eth1 node1.controller1.eth2	Controller 1 Eth 1 Controller 1 Eth 2	1G RJ45 Ethernet. Used as management ports for the DVX GUI and CLI, and for Datrium Support remote access.
node1.controller1.eth3 node1.controller1.eth4	Controller 1 Eth 3 Controller 1 Eth 4	10G BASE-T RJ45 Ethernet ports used as data ports.
node1.controller1.eth5 node1.controller1.eth6 node1.controller1.eth7 node1.controller1.eth8	Controller 1 Eth 5 Controller 1 Eth 6 Controller 1 Eth 7 Controller 1 Eth 8	10G SFP+ ports used as data ports. These ports accept direct attach twin axial copper connections or optical connections. When you use the SFP+ ports, you can use one of the following combinations of these ports: eth5+eth6 eth7+eth8 eth5+eth6+eth7+eth8 Datrium offers 10G optical transceivers as an option for these 10G Data Nodes.
CLI Port Name	GUI Port Name	Port Type
node1.controller1.eth1 node1.controller1.eth2	Controller 1 Eth 1 Controller 1 Eth 2	1G RJ45 Ethernet. Used as management ports for the DVX GUI and CLI, and for Datrium Support remote access.

CLI Port Name	GUI Port Name	Port Type
node1.controller1.eth3 node1.controller1.eth4	Controller 1 Eth 3 Controller 1 Eth 4	10G SFP+ ports used as data ports. These ports accept direct attach twin axial copper connections or optical connections. The ports require Datrium-supplied adapters that accept the copper or optical modules. (Melanox Copper/Direct-Attach module or Melanox SFP+ 10GBASE-SR optical module.)
node1.controller1.eth5 node1.controller1.eth6	Controller 1 Eth 5 Controller 1 Eth 6	10G BASET RJ45 Ethernet ports used as data ports.

Data Node High Availability

The DVX Data Node provides high availability through controller redundancy, network port redundancy, and Data Node redundancy with Resiliency :

- Controller Redundancy uses a pair of controllers in an active/standby configuration to provide continuous access to data.
- Network Port Redundancy uses active-active data port pairs to provide aggregate bandwidth. If there is a network path failure, the DVX System will use the remaining active port. The Data Node also supports an active/standby bonded port-pair configuration for the management interface.

Data Node Resiliency provides high availability of the DVX Data Pool in the event of a data node failure in a scale-out DVX of 3 or more Data Nodes. For more information, see Data Pool Resilience for (HA).

Controller Redundancy

The DVX System provides redundant Data Node controller components to support continuous access to data. The Data Node provides an active/standby HA (High Availability) capability.

The Data Node supports controller failover in response to controller failure. It also provides a dynamic controller failover capability that is based on Data Node connectivity and Compute Node connectivity.

- Controller failure – The Data Node has a pair of controllers that operate in active/passive configuration for high availability. If the active controller fails, the DVX System promotes the passive controller to active status and uses the newly active controller to continue to serve data requests.
- Data Node connectivity – With a Data Pool of more than one Data Node, a DVX System supports dynamic controller failover based on connectivity between the Data Nodes. The DVX System will fail over controllers on Data Nodes to maintain connectivity between all Nodes in the Data Pool.
- Compute Node connectivity – If the following conditions are satisfied, the DVX System will failover to the standby controller.
    - The standby controller must have connectivity to more Compute Nodes than the active controller.
    - The standby controller set of Compute Nodes must include the same set of Compute Nodes that are connected to the active controller.
    - Connectivity between Data Nodes in the Data Pool is not compromised.

To support controller redundancy, you must connect cables to the ports that you will use on both controllers. When you configure data interfaces, you identify two or more data interfaces. When you configure the management interface, you identify one or both management interfaces. A single DVX network interface represents the same port on both controllers.

The picture below shows the change in data port usage if there is a controller failover. The teamed port pair eth5+eth6 is configured on the Data Node. Controller 1 is the active controller.

If controller 1 fails or if controller 2 has network connectivity to more hosts (including all the hosts with connectivity to controller 1), controller 2 becomes the active controller and the data interface on controller 2 assumes the active role. Data Node communication with ESXi hosts can continue without reconfiguration.

Floating IP Address

The Data Pool uses floating IP addresses to support the High Availability capability in the event of controller failover. The Data Pool supports one floating IP address for the data ports and a second floating IP address for the management ports.

- The Hyperdriver on a host uses the data floating IP address to communicate with the Data Pool.
- The DVX GUI uses the management floating IP address to communicate with the Data Pool.
- Use the management floating IP address to create DVX CLI and DVX GUI connections to the Data Pool.

Regardless of any controller failover that might occur in a Data Pool, the DVX System maintains the floating IP addresses for continued access.

Network Port Redundancy

The Data Node provides adaptive pathing for data ports and bonded pair redundancy for management ports.

- Adaptive pathing uses two or more data interfaces of the same connector type for aggregate bandwidth over the interfaces. If the network path for one of the interfaces fails, the DVX System will continue to use the bandwidth on the remaining active interface(s).
- Bonded pair management ports provide an active/passive redundant interface. If the active port fails, the passive interface becomes active and the system fails over to the newly active interface.

The table below shows the interface team possibilities for DVX traffic. You can configure only one data interface team on a Data Node. You cannot mix ports of different speeds or

connector types. For example, you cannot configure both eth3 and eth6 ports at the same time.

Data Node Model	Traffic type	Interface Teams	Speed and Connector Type
D12x4B 2x25G D12x4C 2x25G D12x10D 2x25G F24x2B 2x25G	Data	eth3+eth4	2x10G BASE-T RJ45
		eth5+eth6	2x10G SFP+/25G SFP28 (optical or direct attach copper)
F24x2D 2x25G	Management	eth1+eth2	1G RJ45
D12x4B 4x10G	Data	eth3+eth4	10G BASE-T RJ45
		eth5+eth6	2x10G BASE-T RJ45
		eth7+eth8	2x10G BASE-T RJ45
		eth5+eth6+eth7+eth8	4x10G BASE-T RJ45
	Management	eth1+eth2	1G RJ45
D12x4	Data	eth5+eth6	2x10G BASE-T RJ45
		eth3+eth4	2x10G SFP+ (optical or direct attach copper)
	Management	eth1+eth2	1G RJ45

If there is a controller failover, the management interface on controller 2 supports network failover as well.

Network Topology

The DVX System provides aggregate data bandwidth and continued operation in the event of a network failure.

Active Data Paths

The following figure shows the data paths between hosts and the active controller on the Data Node. This topology uses the following components:

- A Data Node with teamed 10G data ports.
- Two switches that use inter-switch linking.
- Two or more hosts that use NIC teaming for the ports that will send traffic to the Data Node.

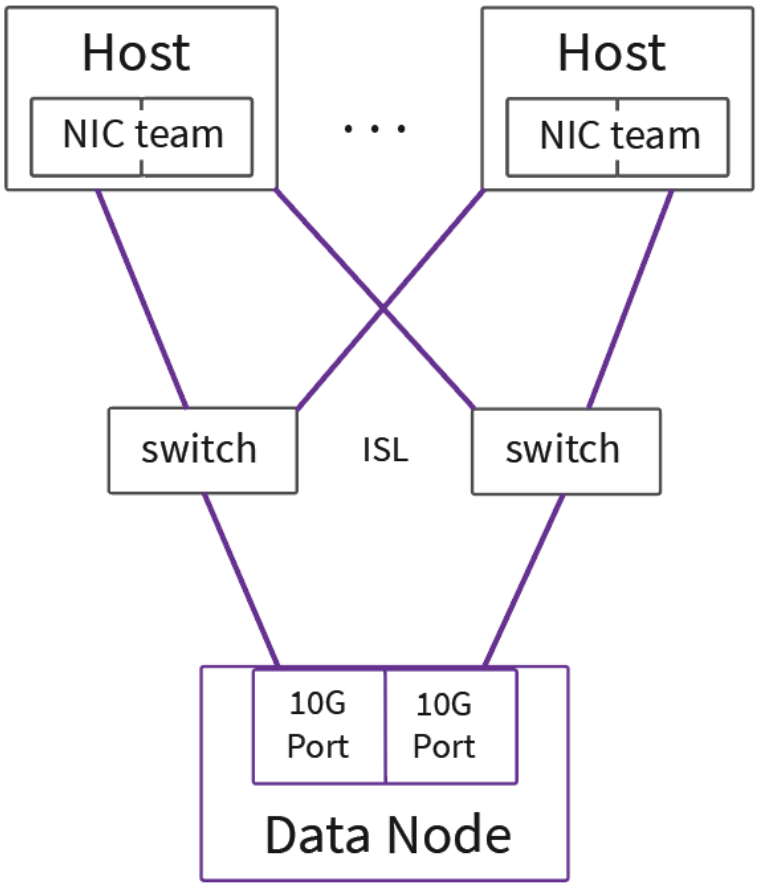

The DVX System distributes traffic across Data Node data interfaces

- The DVX Hyperdriver on a host uses the active NIC port to send data traffic to both ports on the Data Node.

- The use of inter-switch linking supports the distribution of traffic.
- The DVX ports are teamed to support aggregate pathing. The DVX System determines active paths and uses all of the available network interfaces on the active controller.

The DVX System supports uninterrupted data traffic in the event of a network failure along a data path. With redundant switches, the host can use the standby NIC port to send traffic to the reachable data port.

Requirements to support network failover for data traffic:

- Connect the ports to multiple switches that use inter-switch linking.
- The Data Node data interfaces must be on the same subnet as the hosts.
- Use NIC teaming on the host. Use NICs that will support the desired data speed.
- Enable beacon probing on ESXi hosts.
- On Linux hosts, use the network interface active-backup bonding option.

Redundant Topology

The following picture shows the network topology for an environment that has separate subnets for data and management, with support for redundant data and management network interfaces. The gateway is on the management subnet.

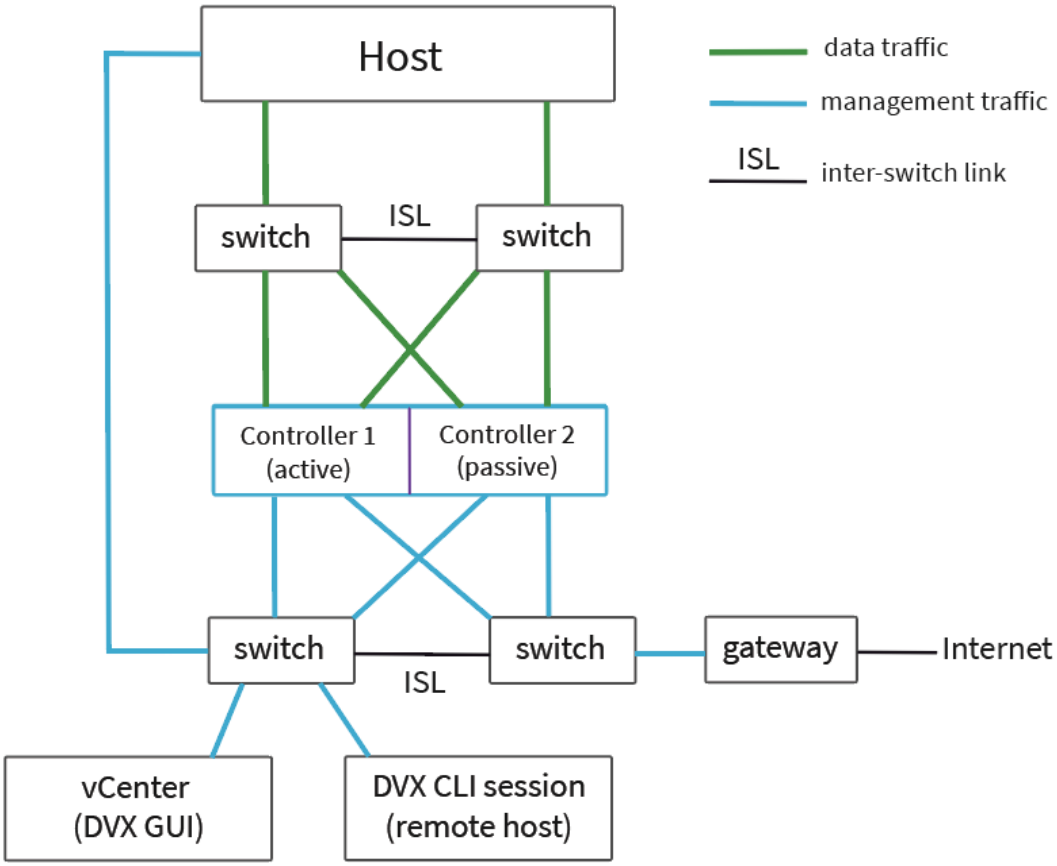

To support controller failover, there are both data and management connections to both controllers on the Data Node.

- The host uses the data subnet for communication with the Data Node.
- The host uses the management subnet for communication with the vCenter Server.
- The vCenter Server uses the management subnet for communication with the Data Node.

To support network failover:

- The Data Node uses teamed data ports and bonded management ports on both controllers.
- Switches within a subnet use inter-switch linking.
- The host uses NIC teaming for the data ports.

The following picture shows the switch connections for teamed data ports.

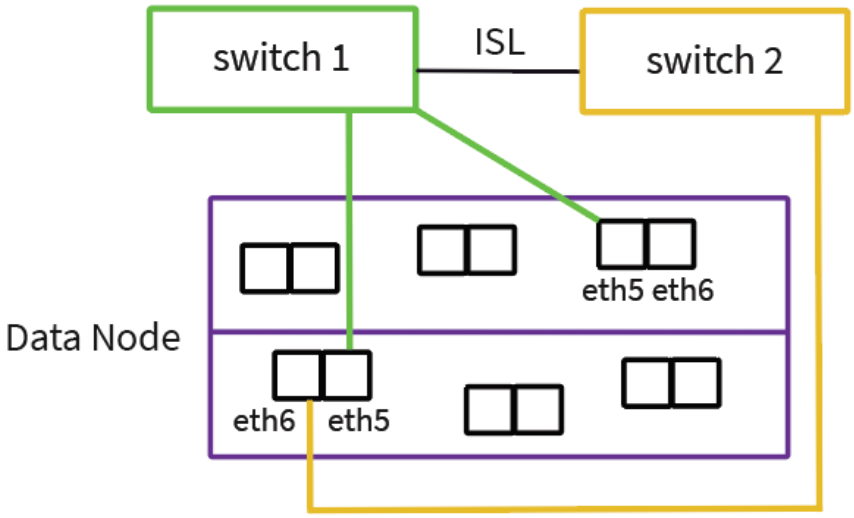

Teamed port connections:

- Connect all ports for a single interface to the same switch.
- Use inter-switch linking between switches that are connected to the Data Node.

In the figure above, both eth5 ports are connected to switch 1 and both eth6 ports are connected to switch 2. The DVX System will use both ports for data traffic. If there is a network failure on one of the ports, the other port in the team carries the load. If there is a controller failover (controller 1 fails over to controller 2), the Data Node will use the teamed ports on the second controller. If there is a network failure in one of the ports in that team, the DVX System will use the active port.

Management Interface – Bonded Port Pair:

- When you configure a management bonded port pair, the first interface in the pair is considered the primary interface. When a network link is available on the primary interface, that interface is the active interface. If there is a network failure on the eth1 interface, the eth1 interface fails over to eth2. When the eth1 link is restored, the eth2 interface fails over to the primary eth1 interface.

The following image shows the change in network port usage if there is a management network failure.

- Controller 1 is the active controller.
- The eth1 and eth2 management interfaces are configured on the Data Node.
- The management ports on both controllers are bonded. Each bonded pair has a single IP address.

In this situation, the Data Node is using the eth1 port for management traffic.

If there is a network failure, the eth2 port becomes the active management port and DVX GUI/CLI communication with the Data Node and Datrium Support operations can continue without any additional network configuration. If there is a network failover during a CLI session, you will have to log in again.

If there is a controller failover, the management interface on controller 2 supports network failover as well.

Network Support for Replication

The following picture shows the data paths used in DVX snapshot replication.

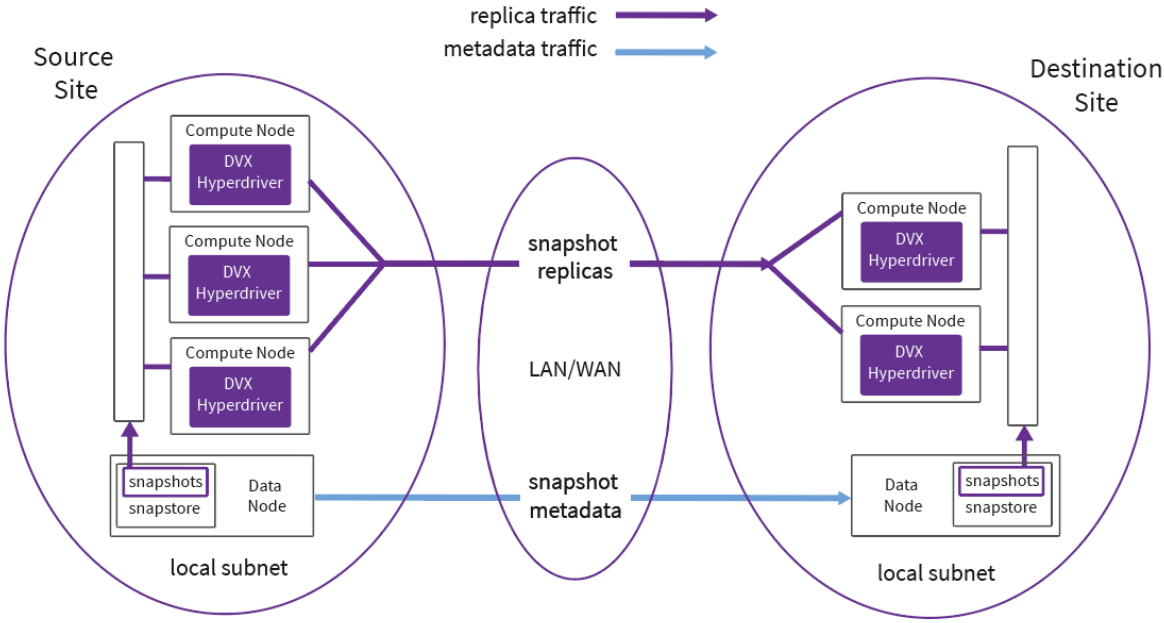

The DVX Systems use Compute Node resources to transfer snapshot replicas.

- Each source site host must have network access to at least one IP address at each destination site host.
- Any firewalls must allow access through port 1525 for host-to-host communication.
- You can designate host ports for replication traffic.
- The DVX System uses the Data-Node-to-Data-Node connection to transfer snapshot metadata.
- The source site Data Node must have network access to the replica site Data Node.
- Any firewalls must allow access through port 4105 for Data-Node-to-Data-Node communication.

Specify the management floating IP address when you define the replica site. (See the *DVX System Management* manual.) The DVX System will use the management interfaces for Data Node communication between the source and destination sites. The management floating IP address must be routable from all destination hosts.

In an individual DVX System site, source or destination, the Data Node and hosts are on the same subnet. Datrium recommends that data traffic between the host(s) and Data Node in a single DVX System should be on the same subnet.

DVX System Performance

The DVX Hyperdriver uses host flash and host compute resources to perform storage operations. If there are insufficient host resources, performance management is achieved through relocation of virtual machines to hosts with more resource headroom. You can use the DVX User Interface (GUI or CLI) to monitor storage activity and determine how to optimize efficient use of resources.

The following sections provide information about monitoring and managing the DVX System.

- Monitoring Performance (DVX GUI)
- Monitoring Performance (DVX CLI)
- VM Performance
- Host Flash Monitoring and Management
- Latency

The following elements in your environment have an effect on DVX System performance:

Host flash cache	The DVX System uses host flash to accelerate read and write performance. The host flash architecture also improves performance during boot storms that involve similar operating system images on the same host. The first cold boot request transfers boot blocks from the Data Node to host flash; subsequent boot requests will be resolved in the deduplicated host flash cache. The amount of host flash affects system performance. If you use enough host flash to hold all active whole vDisks (thin, after compression and deduplication across all vDisks on that host), you can eliminate host flash cache misses.
Host compute resources	The DVX System uses host compute resources to process I/O requests. The DVX performs compression, deduplication, and serialization operations on the host.
Network bandwidth	The DVX System is intended to operate over a high-speed

	Ethernet network. The Data Node provides teamed data interfaces that support up to 20 gigabit aggregate bandwidth. To create a network environment that will support the best storage performance, the host(s) and Data Node should be on the same subnet. The storage subnet should be isolated from general network traffic to avoid saturation of the subnet.
Data Node capacity and aggregate bandwidth	The Data Node can store up to 29TB of deduplicated, compressed data, supporting up to 1500 powered-on virtual machines distributed across 32 hosts.
Individual virtual storage activity	Some virtual machines may use more resources than other virtual machines. To balance the load, you can use vSphere vMotion to move resource-intensive virtual machines to hosts that have spare capacity.

Monitoring Performance (DVX GUI)

The DVX GUI provides the following pages and tabs for performance monitoring:

Performance Data	DVX or vSphere Page (Tab)
Aggregate data (DVX System) – Performance graphs (real-time, 24 hr., host, virtual machine) for IOPS, throughput, latency, and cache hit rate.	DVX Real Time (Summary)
Historic charts (Throughput, IOPS, latency, hit rate, flash cache, physical capacity)	DVX Dashboard (Historic charts) Host page (Monitor->Datrium DVX->Historic charts) VM page (Monitor->Datrium DVX->Historic charts)
Top host per metric (aggregate)	DVX Real Time (Summary)
Top host per metric	DVX Real Time (Historic charts)
Top virtual machine per metric	Host page (Monitor->Datrium DVX->Summary)

Performance Data	DVX or vSphere Page (Tab)
Virtual machine performance (aggregate – hit rate, IOPS, throughput, latency, queue length)	Host page (Monitor->Datrium DVX->Summary)
Virtual machine (hit rate, IOPS, throughput, latency, queue length)	VM page (Monitor->Datrium DVX->Summary)
IOPS (Cache read, Data Node read)	Host page (Monitor->Datrium DVX->Summary) Host page (Monitor->Datrium DVX->Historic charts) VM page (Monitor->Datrium DVX->Summary)
Flash map	DVX Real Time (Summary) Host page (Monitor->Datrium DVX->Summary) VM page (Monitor->Datrium DVX->Summary)
Flash data reduction	DVX Real Time (Summary) Host page (Monitor->Datrium DVX->Summary) VM page (Monitor->Datrium DVX->Summary)
Summary statistics (host)	DVX Hosts page – list of hosts
Summary statistics (virtual machine)	VMs list page (Monitor->Datrium DVX->VMs)

Monitoring Performance (DVX CLI)

You can use the following CLI commands to monitor the DVX System:

`dvx perf`	Displays storage pool and host statistics.
	Displays real-time statistics about host I/O performance.

`hosts show`	Displays summary statistics for each host.
`pool perf`	Displays real-time Storage Pool statistics.

VM Performance

The following sections provide information about tuning virtual machine performance

- Load Balancing
- DVX Performance Modes

Load Balancing

In the DVX System, virtual machines use storage resources on the local host. The DVX Hyperdriver uses host resources to perform storage operations – storage performance is local and isolated from other hosts in the environment.

To adjust virtual machine storage speed, use VMotion to migrate the virtual machine to a host with better compute or flash headroom. For example, you can quickly identify hosts that have a cache shortage by looking at the host list on the “Hosts in DVX” page.

The virtual machine list on the host page (Datrium DVX -> VMs tab) shows cache usage for individual machines. To determine a candidate for migration, look for virtual machines with active data that fits within the headroom of a host that is over-provisioned. Ideally this will be a virtual machine with a low cache hit rate, indicating a higher level of network access.

To obtain the best performance, make sure that the local host cache hit rate is 100%. In particular, if your hit rate is less than 98% make changes to your environment to improve the hit rate. There is a significant improvement as the hit rate increases to 99%, and an even more significant improvement as the hit rate achieves 100%. Consider the following:

- Use VMotion to migrate virtual machines to hosts with available cache capacity.
- Add SSD(s) to increase cache capacity on individual hosts.
- Add host(s) to increase the total cache capacity for the DVX System.

DVX Performance Modes

The DVX System defines two performance modes. The performance mode determines the amount of host CPU resources that the Hyperdriver will use:

- **Fast**. The default performance mode. The Hyperdriver uses a minimum of 3 cores on a host. If the host has more than 16 cores, the Hyperdriver uses up to 20% of the available host cores, to a maximum of 16 cores.
- **Insane**. Maximized performance. You can use the DVX GUI to invoke Insane Mode. In Insane Mode, the Hyperdriver uses up to 40% of the cores on the host, to a maximum of 16 cores.

These figures reflect running ESXi hosts with hyperthreading enabled. Datrium recommends that you run ESXi hosts with hyperthreading enabled. For information about how to use the vSphere Web Client to enable hyperthreading, see the VMware documentation for vSphere Resource Management.

Setting the DVX Performance Mode

The performance mode buttons (Fast, Insane) are in the upper right corner of Datrium DVX tab on the host pages. Use the DVX GUI plugin to the vSphere Web Client to set the mode.

Performance Modes

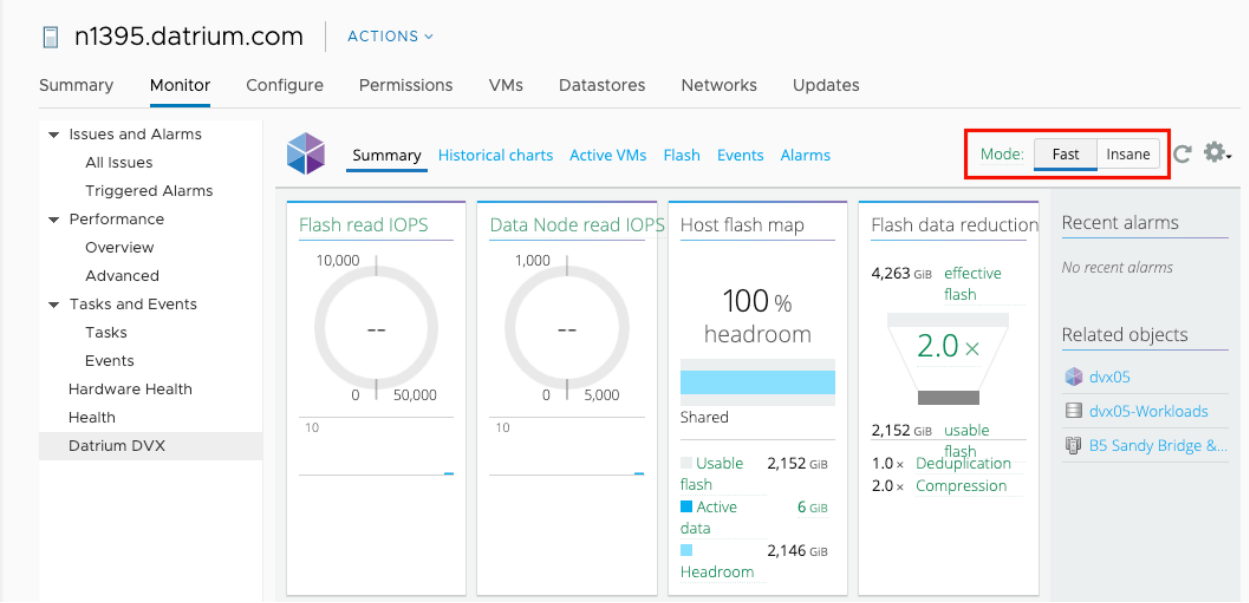

Effective use of the DVX performance modes depends on the state of the DVX System.

- Is the virtual machine load distributed across the hosts in the system? All hosts should be running with 100% cache hit rate.
- If the total number of cores on the host is close to the minimum required for DVX Hyperdriver installation (8 cores), the Insane Mode performance gain will be minimal.
- If the DVX minimum core usage (up to 20% in Fast Mode) is close to the DVX maximum of 16 cores (hyperthreading enabled), Insane Mode performance gain will be minimal or none. For example, if the total number of cores on the host is 80, Fast Mode and Insane Mode are equivalent.

Using Insane Mode

The Hyperdriver operates in response to virtual machine demand:

- The Hyperdriver uses all of the CPU resources that it has reserved only at peak IOPS. If the virtual machines are not running at peak IOPS, the Hyperdriver consumes only a corresponding fraction of the reserved CPU resource.
- The Hyperdriver will not consume more CPU resource than it has reserved, regardless of I/O load.

To determine when to use Insane Mode, you can monitor host IOPS, host queue length, and average CPU utilization. The “Aggregate VM performance” table on the Datrium tab on a host page shows both host IOPS and host queue length. The esxtop utility shows the average CPU utilization.

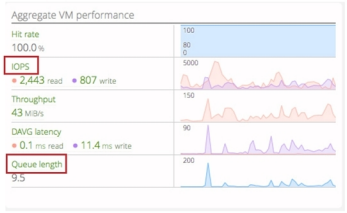

Host IOPS

Insane Mode is useful during periods when host I/O peaks. To monitor host I/O, use the “Aggregate VM Performance” statistics for virtual machine IOPS on the Datrium tab on a host page. Also, use the historic IOPS chart (Monitor->DVX->Historic Charts).

Host Queue Length

When the queue length indicates an I/O bottleneck, you can use Insane Mode to increase the resources used for I/O processing. The queue length threshold is dependent upon your particular workload. To monitor the host queue length, use the “Aggregate VM Performance” statistics for “Queue Length” on the Datrium tab on a host page.

Average CPU Utilization

You can use Insane Mode when the total average CPU utilization is at 80% or below. Use the vSphere esxtop tool to monitor CPU utilization. The PCPU UTIL measurements in the esxtop display include the average CPU usage of the host.

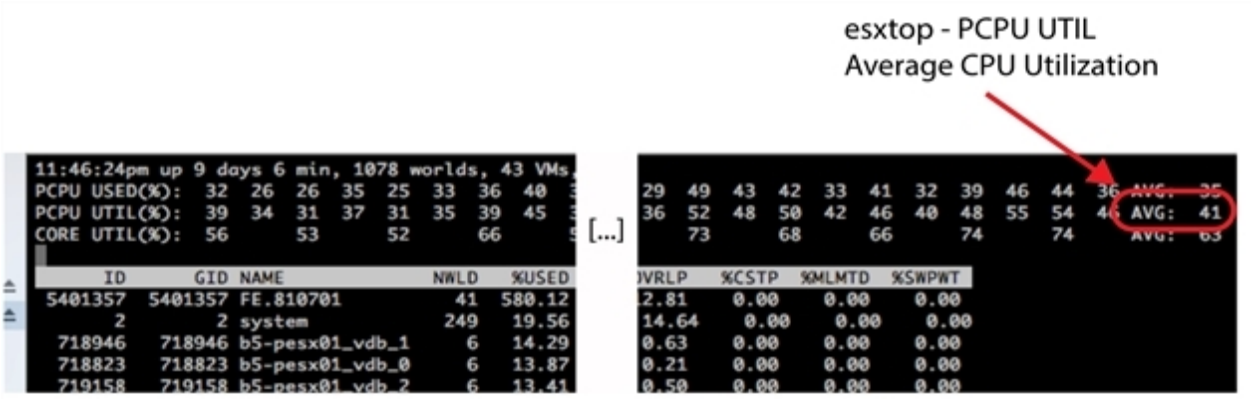

It is ideal for the average CPU utilization to be at or below 80%. Above 80%, performance will start to degrade. CPU utilization above 85% will cause significant performance problems with latency and throughput.

When there is a demand for the DVX storage resources, the Hyperdriver uses CPU resources up to the limit associated with the current performance mode.

- Fast Mode uses up to 20% of the CPU resources. Fast Mode is the default. Before Hyperdriver installation, the average CPU utilization should be at or below 60%.

- Insane Mode uses up to 40% of the CPU resources. To turn on Insane Mode, the host average CPU utilization (virtual machines plus Hyperdriver) should be at or below 60%.
- If your virtual machines will use 60% or more of the CPU resources on a regular basis, using Insane Mode might have a negative effect on I/O performance.

CPU Usage

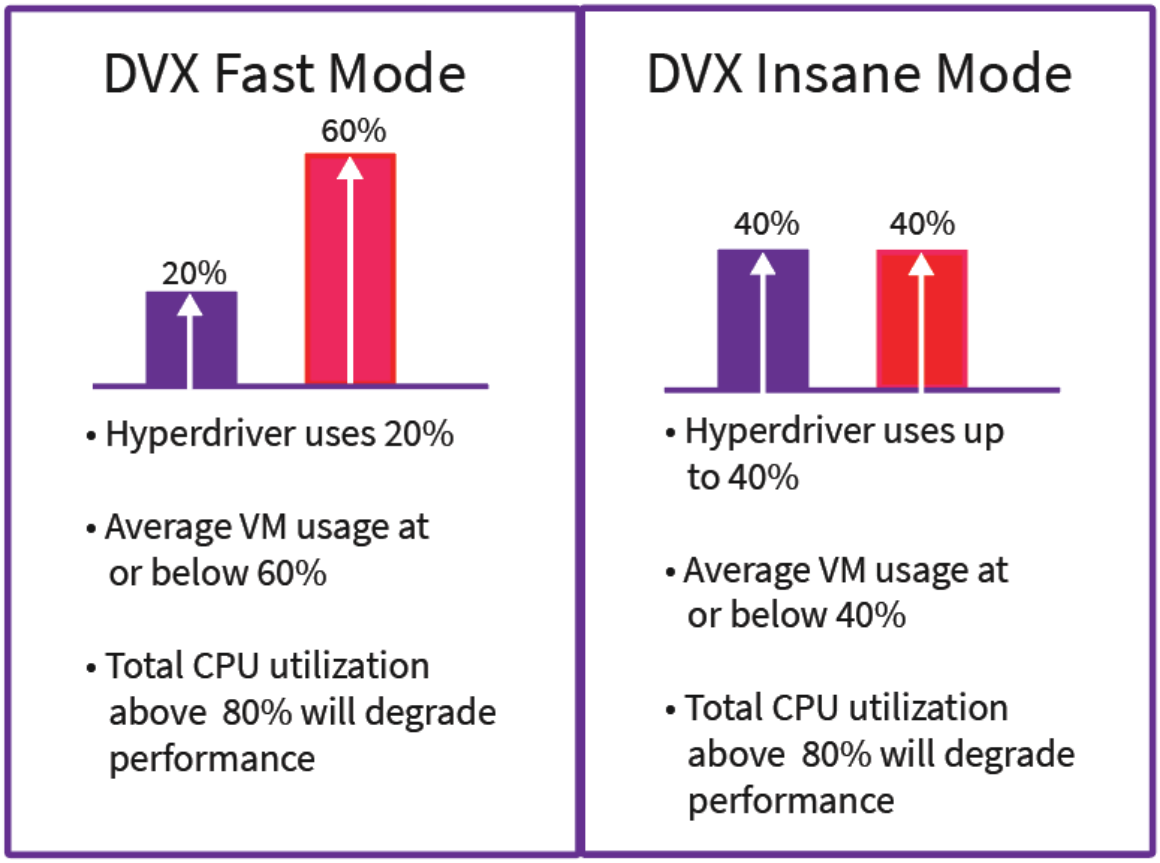

Host Flash Monitoring and Management

The DVX Hyperdriver uses host flash as a local cache for virtual machine storage. The following image represents Hyperdriver operations on virtual machine data and its use of host flash for data storage.

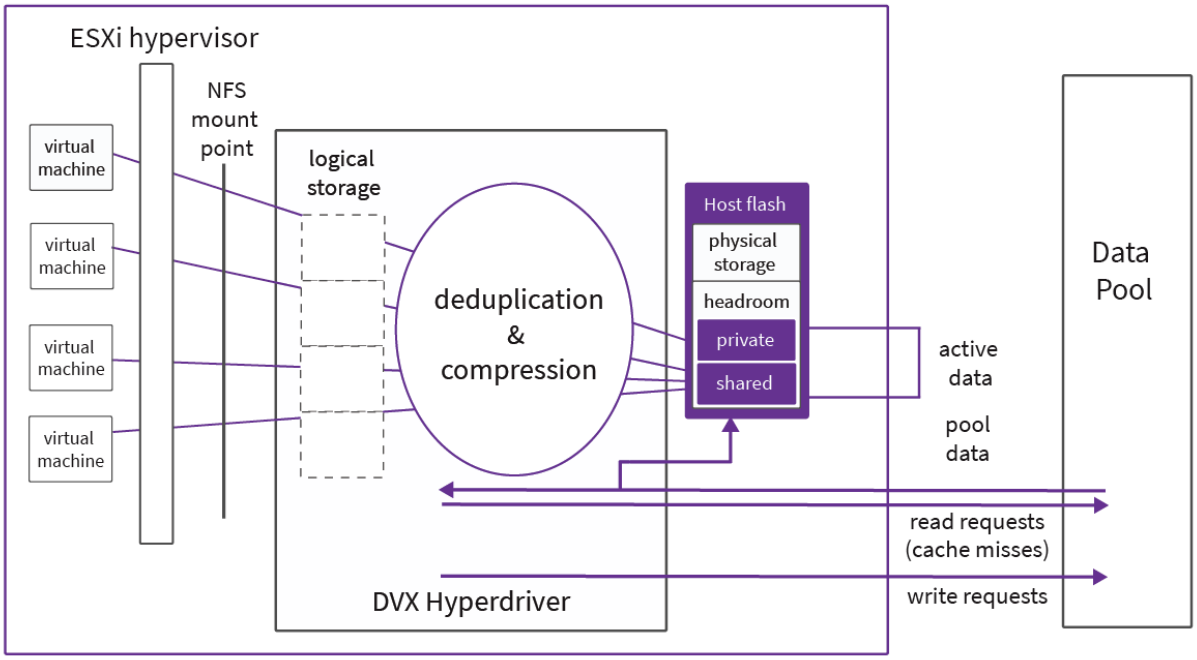

Virtual machines store and access data through the ESXi interface to the DVX NFS mount point. Storage requests are handled by the DVX Hyperdriver. The Hyperdriver sends write requests directly to the Data Pool, then to local flash. Read requests that are not resolved in local flash are sent to the Data Pool.

The Hyperdriver performs inline deduplication and compression on virtual machine data that it stores in the local cache. The DVX System provides statistical data based on effective and physical storage measurements.

- Effective storage is virtual machine data that has not been deduplicated or compressed.
- Physical storage is host flash storage that contains active data. Active data is recently accessed data. The active data is divided into private data and shared data. Private data is accessed by a single virtual machine. Shared data is deduplicated data that is accessed by more than one virtual machine.

Active Data – Host Flash Map

The DVX System provides data about the dynamic state of host flash storage. On a DVX GUI host page, the host flash map shows the following measurements:

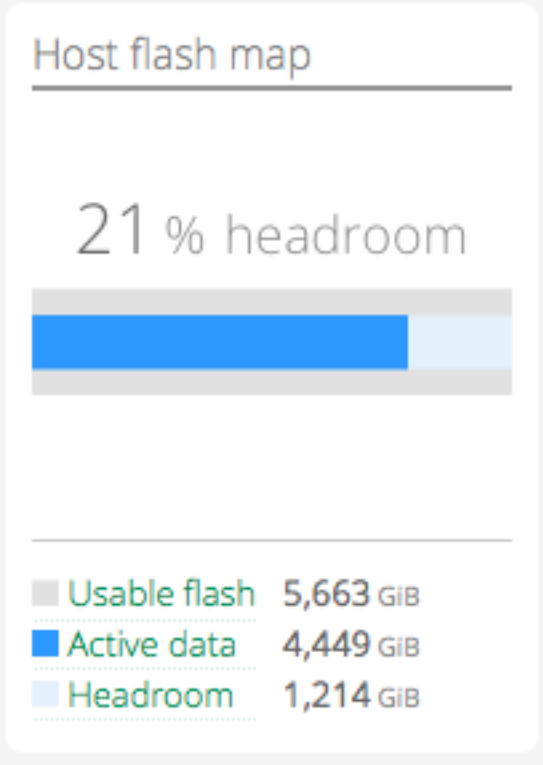

- Usable flash – the total physical capacity of the DVX-dedicated flash on the host.
- Active data – the physical size of recently accessed data stored on flash.
- Headroom – the percentage of unused host flash.
- Shortage – the percentage of additional flash needed to satisfy read requests locally.

1. If the host flash map indicates that there is a shortage, you can use the following solutions:
    - Add more flash to the host. The best performance is achieved if 100% of read transactions are resolved in the local cache.
    - The DVX System can use a maximum of 10 SSDs on a single host.
    - The DVX System can use up to a maximum of 16TB on a single SSD.
    - The minimum size of an individual SSD for DVX use is 400GB raw capacity.
    - Each host can support up to 32TB of total flash (across all SSDs). If the cumulative flash size on a host is over 32TB, the DVX system will use only 32TB.
2. Use VMotion to move one or more virtual machines to one or more hosts that have sufficient headroom.

When you move a virtual machine to a different host that has sufficient capacity, it will resolve the capacity shortage for that particular virtual machine. It might or might not resolve the flash capacity shortage on the source host. If the moved virtual machine was using data that was shared with other virtual machines, the other virtual machines on the source host still access that data, keeping it in the flash. You might have to move more than one virtual machine to free up flash capacity.

To select a virtual machine for migration, look for virtual machines that do not use shared data.

The virtual machine private data and capacity measurements are available on the virtual machine page in the VM flash map and Flash data reduction charts.

- Private data is in the Flash data reduction tooltip.
- To display the tooltip, place the mouse cursor over the data reduction factor.

The target host must have sufficient headroom for the virtual machine data.

- If the virtual machine is running with sufficient headroom (in the flash map, the active data is less than or equal to the local data capacity), compare the assigned flash measurement with the target host headroom.
- If the virtual machine is running with flash shortage (in the flash map, the active data is greater than the local data capacity), compare the VM flash map active data measurement with the target host headroom.

Historical Charts (Virtual Machine)

To get a more accurate picture of usage, navigate to a virtual machine page and select “Historic charts”. The DVX GUI displays various historical statistics. You can change the time window to see the cache hits for that virtual machine during a particular period.

If you have a virtual machine that has a consistently large workload, it might be appropriate to dedicate a host to that virtual machine and use enough flash capacity to support the activity.

Latency

In an efficiently configured system with adequate resources, the majority of read requests should be resolved in the host SSD cache.

- If cache hits are reduced, you will see more network traffic and increased latency. The DVX System reports latency as the device average response time or DAVG directly observed by virtual machines. In the vSphere context, DAVG is the latency for an I/O request – DAVG latency includes both network latency and storage device latency. The DVX System also reports host DAVG matching the native ESXi metric reported by esxtop.
- You will normally see increased network traffic during cold startup, when the DVX Hyperdriver must read from the Data Pool to warm up the cache.

DVX Device Average Response Time (DAVG)

DVX DAVG statistics can represent two types of measurements:

Local access – Latency associated with storage requests that are resolved locally in the host flash. If the host has sufficient flash for the virtual machine activity, the DVX Hyperdriver can use the active data in local flash to satisfy read requests. DAVG measurements reflect access to local flash as seen by the virtual machines.

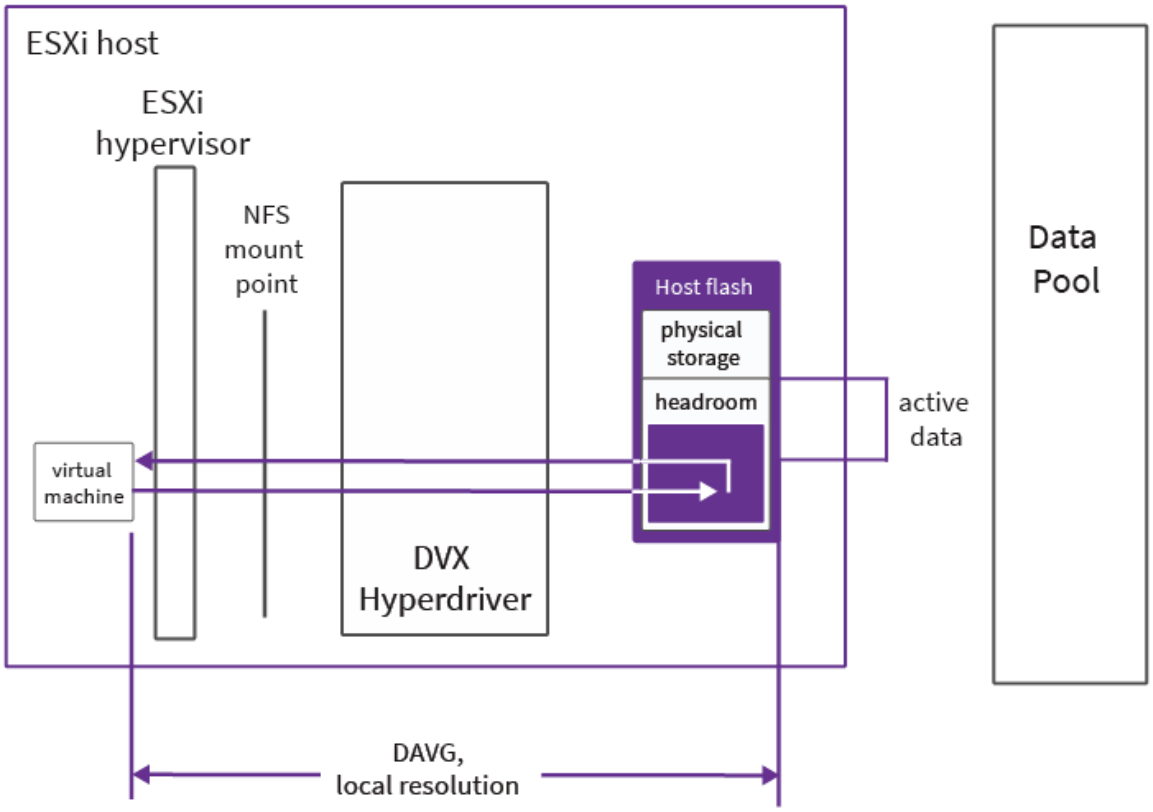

Network access – Latency associated with storage requests that require storage pool access over the network. The DVX Hyperdriver sends write data directly to the Data Node. The transaction is not acknowledged until the data is written to NVRAM on a Data Node in the Data Pool. If the Hyperdriver cannot resolve a read request in the local flash active data, it must retrieve the data from the Data Pool. In both cases, the DAVG measurements reflect the network access to the Data Pool as seen by the virtual machine.

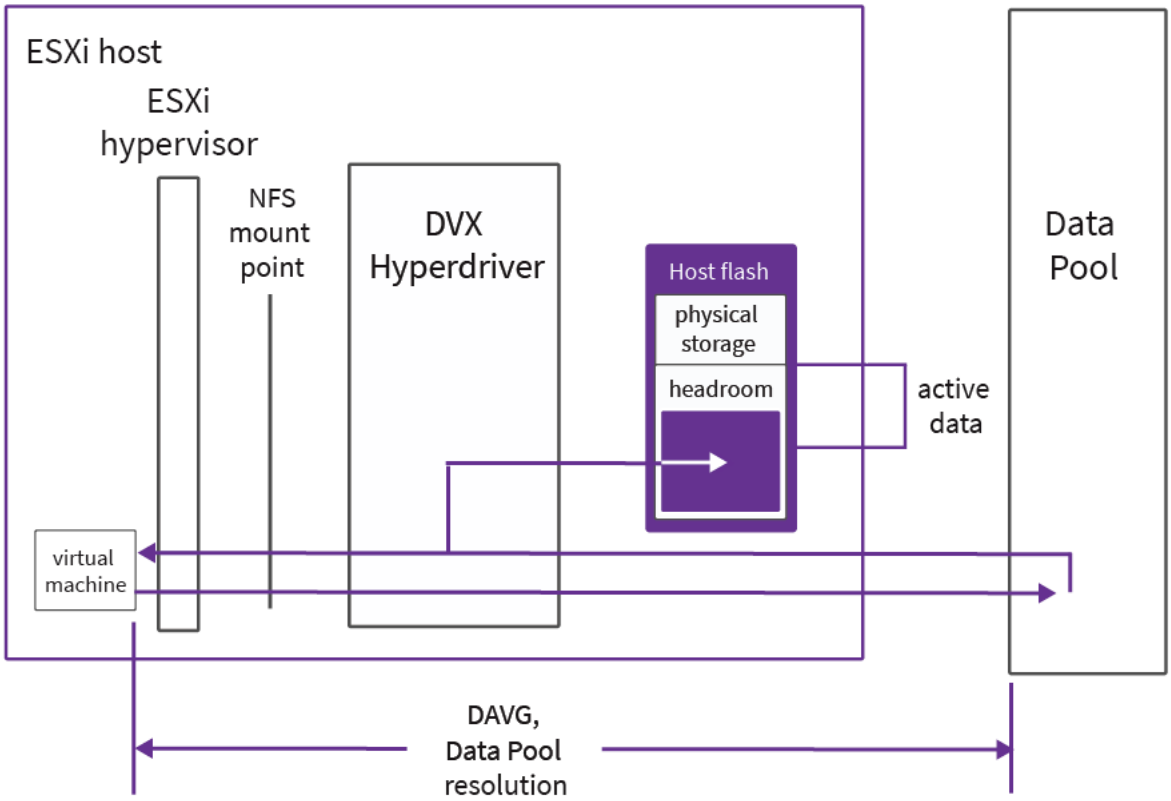

ESXi Performance

The following information relates to virtual machine performance in the ESXi environment.

In situations where your system is running under load – for example a bootstorm, or virtual machines with applications that use multiple virtual disks and generate lots of I/O – you might see NFS out-of-memory warnings in the vmkernel.log file:

```
NFS: 5440: Failed to convert sgArr to NFSIoInfo: Out
of memory
```

This indicates that too many vDisks are open and there are too many outstanding I/O requests, causing guest I/O commands to fail.

- Because some guest I/O requests may fail, some guest operating system or file system may not be able to handle such I/O request failures, for example, during boot, or the virtual disk may become inaccessible from inside the virtual machine. If that happens,

you might have to power off and on the virtual machine again.

- Virtual machine performance might be affected, including the inability to power on virtual machines.

If this occurs during a bootstorm, this situation can resolve itself over time. If you see these warnings during normal operations, when you have virtual machines running under load, you should use vMotion to rebalance the load across your hosts.

DVX Host Management

The following sections provide information about DVX host management.

Adding a Host to a DVX System

When you add a host to the DVX System, you perform the following tasks:

1. Install the Hyperdriver VIB on the host and use the pre-day-zero host check to verify host suitability for DVX use. Install the VAAI VIB if desired.
2. Configure SSDs for DVX use.
3. Mount the Datrium datastore.

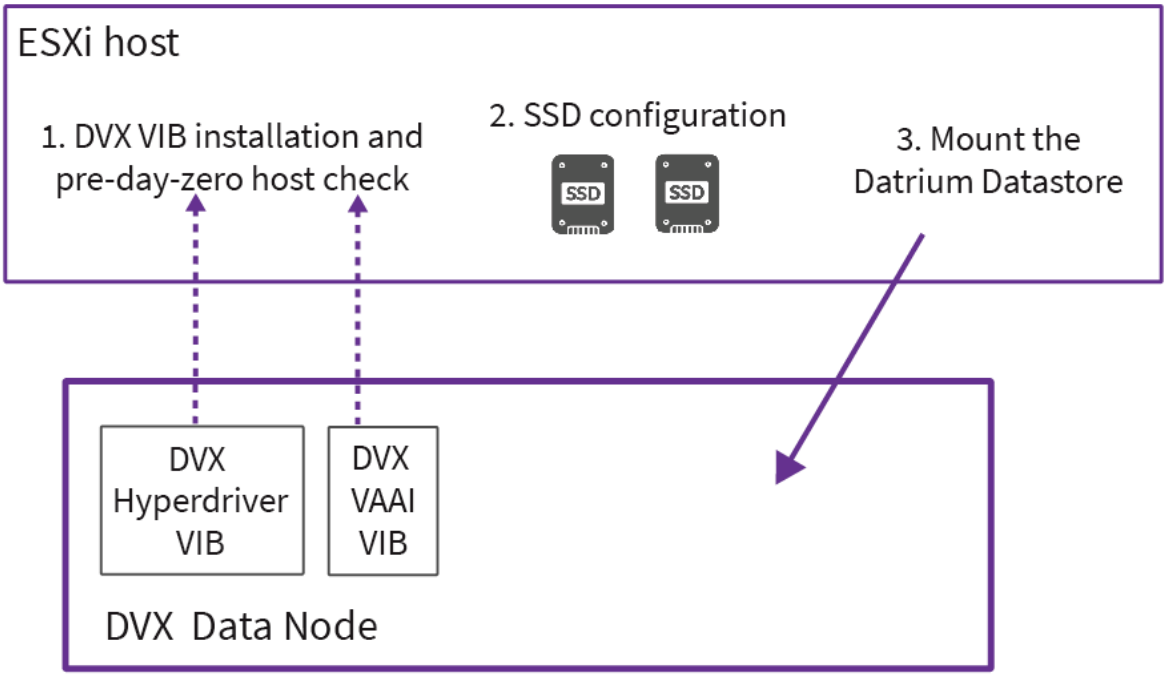

You use the DVX GUI plug-in to perform most of the host setup. The VAAI VIB is not required for DVX operation. Install it to use the DVX support for cloning. You must install the VAAI VIB manually.

DVX Use of Host Resources

The DVX Hyperdriver uses host CPU, memory, and SSD resources:

- The Hyperdriver requires a minimum of 8 cores on a host. The DVX System will not install the Hyperdriver on the host if it does not have at least 8 cores.
- The Hyperdriver uses a minimum of 14.9 GiB of host memory. It can use up to a maximum of 85.9 GiB of host memory, depending on the amount of host flash that is dedicated for DVX use.
- The Hyperdriver requires at least one host SSD. Datrium recommends a minimum of two SSDs on each host. See SSD Requirements.

Host CPU Usage

The DVX System defines two performance modes:

- **Fast**. The default performance mode. The Hyperdriver reserves a minimum of 3 cores on a host. The Hyperdriver uses up to 20% of the available host cores, to the maximum of 16 cores reserved for DVX use.
- **Insane**. Maximized performance. You can use the DVX GUI to invoke insane mode. (See DVX Performance Modes.) In Insane Mode, the Hyperdriver uses up to 40% of the cores on the host, to a maximum of 16 cores reserved for DVX use. If the total number of cores on a host is close to the minimum required for DVX use, the gain in performance will be minimal.

The DVX System reserves CPU resources on the host according to the performance mode and the number of cores on the host. If there are remaining CPU resources, a virtual machine can use them according to its own resource parameters.

> **Note:** Users can reserve resources using resource pools in vCenter, which will map to resource pools on hosts. That might cause some hosts to run out of available CPU resources for reservation. If you try to add such a host to a Data Pool, the DVX System cannot reserve enough CPU on the host and it will fail to mount the host.

Host Memory Usage

The DVX System uses a minimum of 14.9 GiB of memory on each host. The amount of host memory used increases with each additional terabyte of raw flash memory that is dedicated for DVX use, up to a maximum of 85.9 GiB.

The following image hows some examples of the amount of host memory that the Hyperdriver uses based on the amount of raw flash that is dedicated for DVX use. It also includes a chart representing Hyperdriver memory usage.

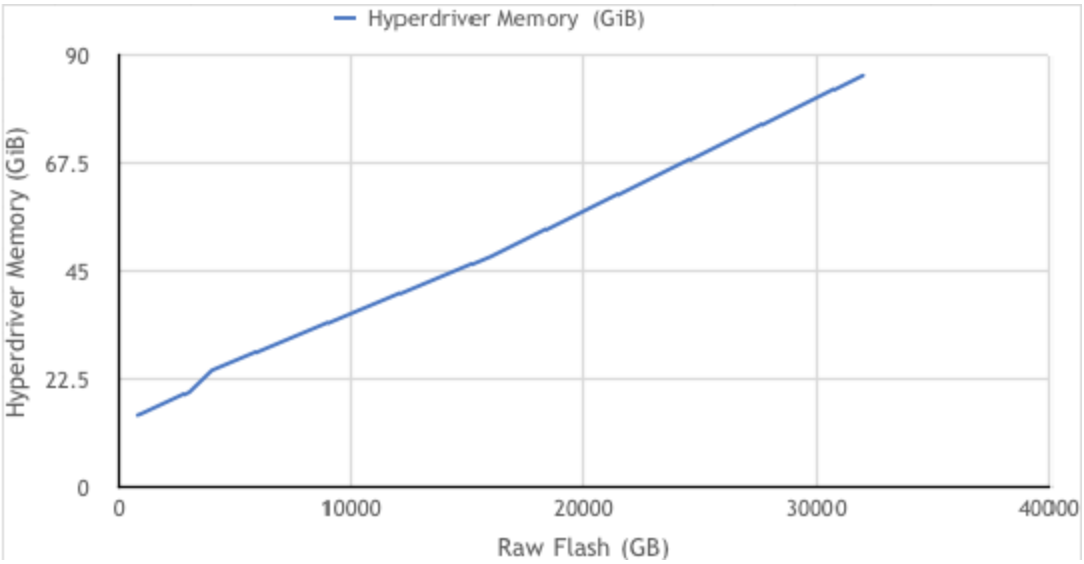

DVX Hyperdriver Host Memory Use Based on Raw Flash Size	
Raw Flash (GB)	**Hyperdriver Memory (GiB)**
800	14.9
1000	15.3
2000	17.5
3000	19.7
4000	24.3
16000	48.1
32000	85.9

SSD Requirements

- The DVX system requires a minimum of 1 SSD on a single host.
- The DVX system can use a maximum of 10 SSDs on a single host.
- The minimum size of an individual SSD for DVX use is 400GB raw capacity.

- The DVX System can use up to a maximum of 16TB raw capacity on a single SSD.
- Each host requires a minimum of 800GB raw capacity for the total amount of flash across all SSDs.
- Each host can support up to 32TB raw capacity for the total amount of flash across all SSDs. If the cumulative host flash size is over 32TB, the DVX system will use only 32TB.

The DVX System will use only internal host SSDs. For example, you cannot use a USB SSD for DVX storage.

Pre-Day-Zero Host Check

Before you add the host to the DVX System, you can use Datrium software to check host suitability for DVX use. Log in to the host and use the following command to install the datrium-hyperdriver-esx VIB on the host.

The following commands install the Hyperdriver VIB from the Data Node, depending on which version of vSphere you are using.

Install the Hyperdriver VIB on the host for vSphere 6:

```
esxcli software vib install -d http://mgmt-floating-
IP-address/static/ESXVIBHyperDriver6
```

Install the Hyperdriver VIB on the host for vSphere 7:

```
esxcli software vib install -d http://mgmt-floating-
IP-address/static/ESXVIBHyperDriver7
```

The datrium-hyperdriver-esx VIB provides a Datrium extension to `esxcli`and to `localcli`.

After installing the Hyperdriver VIB, you must reboot the host to use the Datrium extension to `esxcli`. You can use `localcli`before rebooting.

To perform a pre-day-zero host check, log in to the host and run one of the following two commands:

```
esxcli datrium check
```

or

```
localcli datrium check
```

> **Note:** After installing the hyperdriver VIB on the ESXi host, sometimes the Datrium namespace is not added to the esxcli command. In this case, you will see the following error:
>
> ```
> esxcli datrium
> Error: Unknown command or namespace datrium
> ```
>
> This can be fixed by restarting hostd on the ESXi host. To restart hostd, SSH into the host as root, and then run this command:
>
> ```
> /etc/init.d/hostd restart
> ```
>
> Alternatively, you can run the `localcli` command to access the datrium namespace without the need to restart the hostd. For example:
>
> ```
> localcli datrium
> ```

```
>> esxcli datrium check
Test                              Value                                                Result
--------------------------------  ---------------------------------------------------  ------
CPU                               24 cores                                             OK
Memory                            255.87 GiB                                           OK
NIC speed                         10 Gbps or faster                                    OK
ESXi version                      6.0.0 (3620759)                                      OK
ESXi boot type                    visor-thin                                           OK
Lockdown mode                     Not enabled                                          OK
RAM disk                          Can create 1000 MiB RAM disk                         OK
Root partition                    32512 KiB free                                       OK
Scratch location                  /vmfs/volumes/573f4787-5ca51e2b-1478-5cb9019ad89c    OK
Persistent logging                Yes                                                  OK
vSphere Flash Read Cache (vFRC)   Not enabled                                          OK
Content-Based Read Cache (CBRC)   Not enabled                                          OK
vSAN                              Not enabled                                          OK
Incompatible VIBs                 Not found                                            OK
RAID controller CLI               HP hpssacli installed                                OK
Incompatible storage drivers      Not found                                            OK
```

The host check command indicates the results of various tests. For optimal DVX System operation, all results should be `OK`. You should resolve any situations indicated by any other result. The following table hows the minimum requirements for DVX use.

ESXi Host Criteria for DVX Use (Minimum requirements)	
Host Criteria	**Minimum Requirement**
CPU	8 CPU cores in total
Memory	10 GiB RAM
NIC speed	one vSwitch with at least one 10+ Gbps NIC
ESXi version	ESXi 6.5 or later
ESXi boot type	installable or embedded
Lockdown mode	Not enabled
RAM disk	Host has enough memory to for 1600 MiB RAM disk space
Root partition	10 MiB free space
Scratch location	Scratch location is configured and writable
Persistent logging	Host has configured persistent vmkernel logging
Encryption AES-NI	CPU supports AES-NI (Advanced Encryption Standard New Instructions); AES-NI should be enabled
vSphere Flash Read Cache (vFRC)	Not in use
Content-Based Read Cache (CBRC)	Not enabled
vSAN	Not enabled
Incompatible VIBs	No incompatible third-party VIBs such as the one from PernixData
RAID Controller CLI	For hosts with MegaRaid, Dell or HP RAID controllers, the vendor-specific CLI utility should be installed: • MegaRaid: storcli

ESXi Host Criteria for DVX Use (Minimum requirements)	
Host Criteria	**Minimum Requirement**
	• Dell: perccli • HP: hpssacli
Incompatible storage drivers	No incompatible storage drivers running, e.g. the lsi_mr3 driver on ESXi 5.5

Host Selection and Configuration for DVX System Use

To configure host access to DVX System storage on a Data Node, use the DVX GUI to perform the following tasks:

1. Host Selection – Choose the host(s) that will use DVX storage.
2. Host Configuration (SSD Selection):

    - Identify the DVX System to be used.
    - Select the host flash that the DVX Hyperdriver will use. You can select flash drives for DVX use and you can deselect flash drives to remove them from DVX use.

After you provide this information, the DVX System will install the Hyperdriver software on the host(s), partition the flash drive(s) for Hyperdriver use, and mount the Datrium datastore that will provide access to the Data Node.

If you are using the DVX VAAI VIB, you must install that manually. See DVX VAAI VIB Installation.

Host Selection

You can configure a single host or you can configure multiple hosts at the same time.

Single Host Selection

To configure a single host:

1. Navigate to the vSphere Web Client host page for the host that you want to configure.
2. Select the host Monitor tab, then select the Datrium DVX tab.
3. If the host is not configured for DVX use, the GUI displays an installation dialog. To continue, click on the “Install” button. The GUI installs the Hyperdriver on the host and then displays the DVX host configuration dialog (see Host Configuration (SSD Selection)).

Multiple Host Selection

You can configure multiple hosts at the same time:

1. Select from all hosts managed by the vCenter Server.
2. On the DVX dashboard, click on the gear icon in the upper right corner to display the drop down menu. The “Add hosts” menu entry displays the DVX host selection dialog.
3. Select from the set of hosts in a cluster.
4. Navigate to the vSphere Web Client cluster page.
5. Select the cluster Monitor tab, then select the Datrium DVX tab.
6. Click on the gear icon in the upper right corner to display the DVX system menu. The “Configure hosts” menu entry displays the DVX host selection dialog.

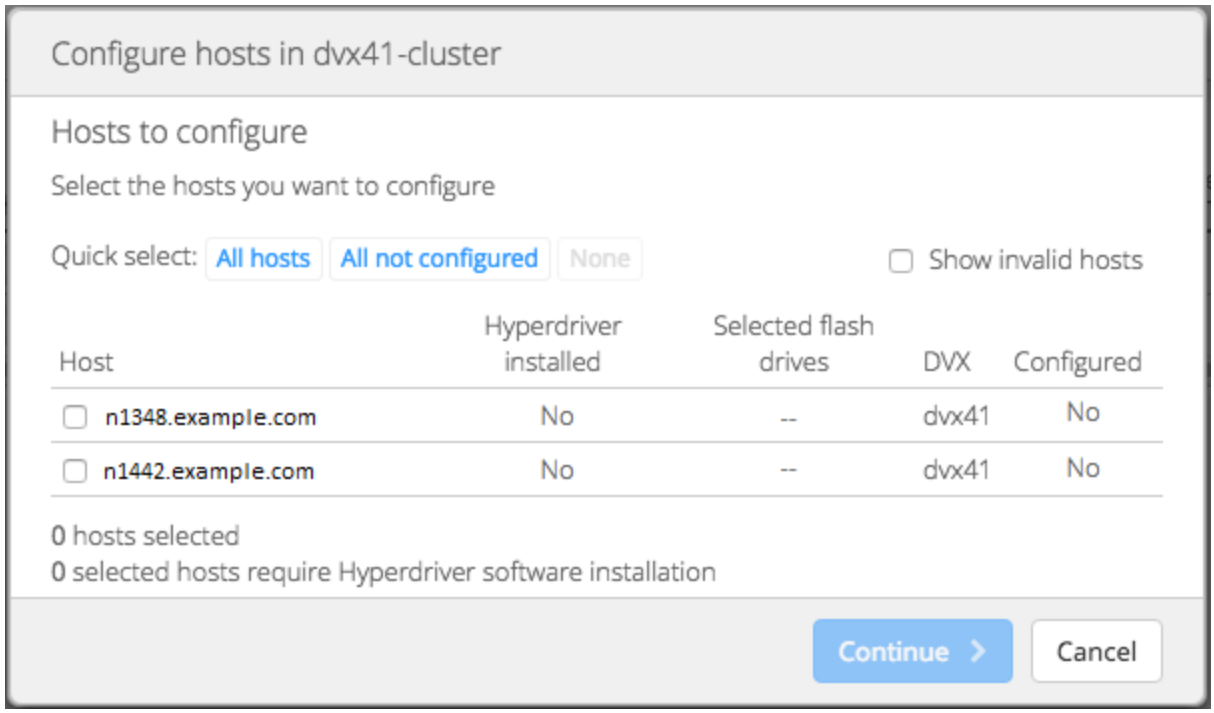

7. Select the hosts that you want to configure, then press the “Install Hyperdriver ... and continue” button. The DVX GUI displays a dialog that indicates Hyperdriver installation is in progress.
8. After the installation is complete, the DVX System creates partitions on the host boot drive and displays the host configuration dialog for SSD selection.

Host Configuration (SSD Selection)

The Host Configuration dialog shows the available datastore(s) and host SSDs for selection.

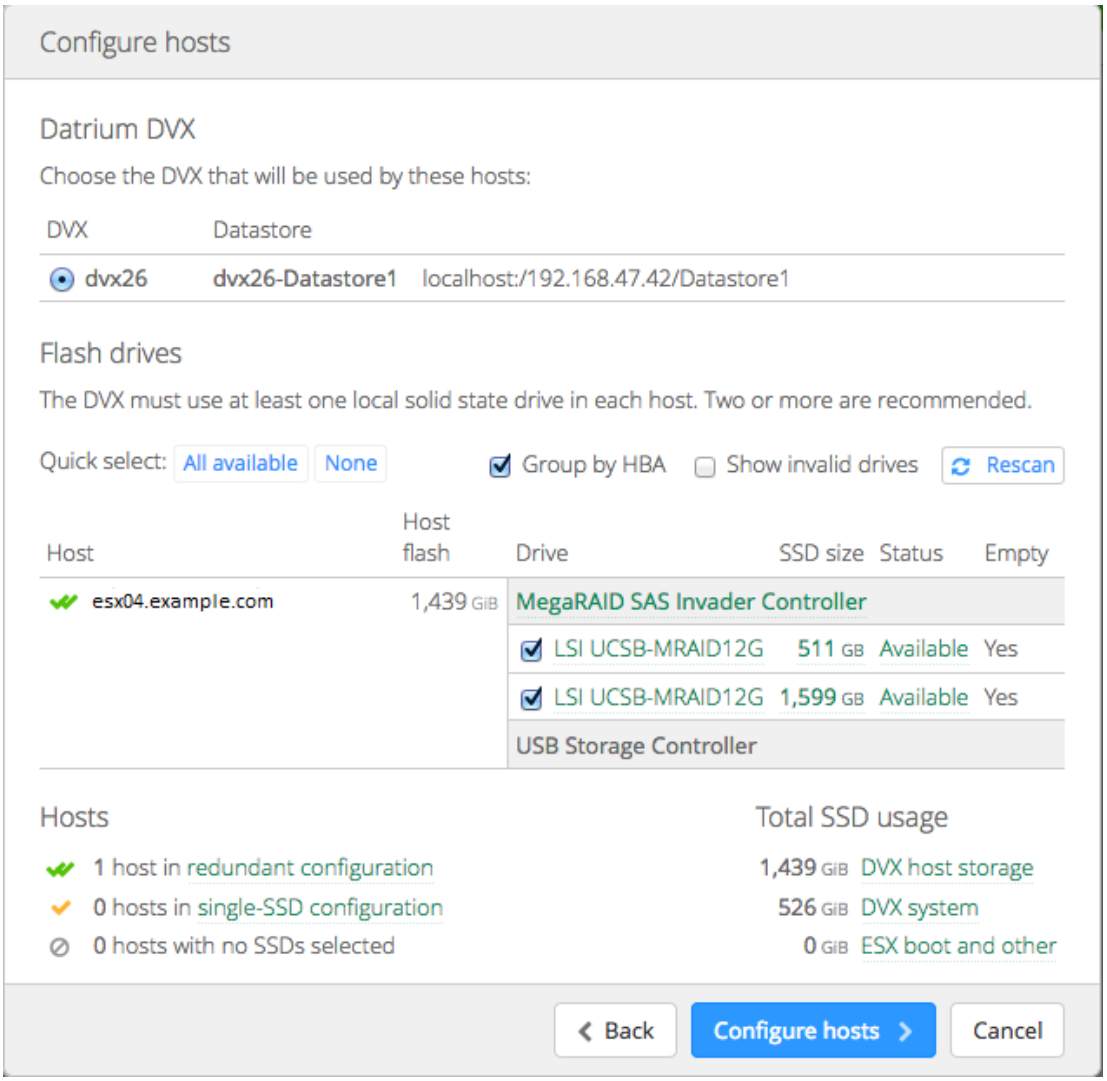

Details of the dialog

DVX datastore selection – If there is only one, it is preselected.

- Flash drives selection
- Quick select buttons – Select all available drives or deselect all drives.
- Grouping by HBA – Displays controller drive identification and groups SSDs accordingly.
- Show invalid drives – Displays drives that are not available for DVX use.
- Rescan – Use the Rescan button to update the SSD display after adding or removing a drive, or to synchronize the DVX display with the vSphere environment.
- When you select the SSDs, the GUI displays the “Host flash” and “Total SSD usage” amounts. The figure shows the SSD selections and the corresponding flash calculations.
- Drive description – The host table indicates the drives available on each host. Any SSDs that are currently used by the DVX System will be selected. The following table hows the host table columns.

Column	Description
Host	Host identification.
Host flash	Total amount of host cache dedicated for DVX use, in gibibytes.
Drive	Drive name.
SSD size	Size of the drive in gigabytes.
Status	Drive status (DVX use, Available, Too small, Unavailable, or Failed).
Empty	Partition status (“--” – configured for DVX use, “No” – SSD contains a VMFS partition, “Yes” - SSD contains a raw partition)

- Hosts:
  - Redundant configuration – A host is in redundant configuration if it has two or more SSDs with DVX cache partitions. A host is in single-SSD configuration if it has only one SSD with a DVX cache partition. Datrium recommends that you use at least two SSDs for the Datrium cache.

- Total SSD usage – Indicates the amount of storage space available for virtual machine storage and the amount of DVX System space, in gibibytes.

Host Configuration Procedure:

1. Select a DVX System. If there is only one, it will be pre-selected.
2. Select the drive(s) that you want to use. On a single host, the DVX System can use up to a maximum of 10 SSDs and a maximum total host SSD capacity of 32TB.
3. The DVX GUI provides two approaches to flash drive selection:

    - "All available drives" – This is an active policy that determines how the Hyperdriver will use flash now and in the future. If you choose this policy, the Hyperdriver will use all drives that are currently installed, and any SSDs that you add to the host in the future, up to the maximum number of drives.
    - "None" – For each host, you must select one or more drives from the list of currently installed drives.

    > **Note:** If a host does not have any SSDs installed, the SSD selection dialog indicates this. In this case, install SSDs and then press the Rescan button. If the new SSDs are not recognized, reboot the host, and then refresh the GUI browser window to continue with SSD selection

4. Press "Configure hosts". The DVX System will install Hyperdriver software on the host (s), partition the selected flash drive(s) for Hyperdriver use, and mount the Datrium datastore.
5. If any of the selected drives contain VMFS partitions, the GUI will prompt for confirmation. If you proceed, the DVX System deletes the VMFS partitions and uses the space for DVX partitions. After the DVX has initialized the SSD(s), it will not use any other disks on the system.

> **Important:**

- Do not perform any operations on these SSDs such as using them to create a datastore.
- Do not use the Datrium partition on the SSDs for virtual machine storage.

- Do not change the names of the partitions. The DVX Hyperdriver generates names for the partitions. The Hyperdriver must be able to use the names that it generates.
- Do not use DVX partitions for the ESXi scratch location

ESXi Configuration

To use an ESXi host with the DVX System, you might have to modify the ESXi configuration.

ESXi Persistent Scratch Location

VMware recommends that you configure a persistent scratch location. Datrium recommends that you do not use Datrium reserved partitions for the scratch location. For information, see the VMware Knowledge Base article about configuring a persistent scratch location.

If the ESXi persistent scratch location is not configured and an ESXi host is using DVX storage, when the host reboots ESXi might automatically use Datrium reserved partitions for the scratch location. To prevent this, set the scratch location to a directory on persistent storage that is not on the dedicated Datrium SSD.

If ESXi uses a Datrium reserved partition for the scratch location DVX performance might be affected. In addition, the Datrium reserved partitions might persist after the following actions:

- Removing the SSD from DVX use.
- Removing the Hyperdriver from the host.

ESXi Syslog

Do not use the Datrium datastore as the location of the ESXi syslog location. Set the `Syslog.global.logDir` ESXi advanced configuration attribute to a different, persistent location that is not used by the DVX System.

Virtual Disk Provisioning

Datrium recommends that you use the default thin VMDK format. Unlike block storage with VMFS, in the DVX System virtual machine write requests do not require block zeroing. All blocks default to zero in thin provisioned VMDKs on DVX datastores.

The DVX System supports thick-provisioned VMDK format only for Oracle RAC deployment. For all other situations, use thin-provisioned virtual disks. See Oracle RAC Support.

The maximum size of a virtual disk that you can create for a virtual machine to be used in the DVX System is 32TB.

Driver Compatibility

Make sure that you have the following drivers on your ESXi host.

Driver	Description
VMware Tools PVSCSI V1.2.3.0	In rare cases, earlier versions of this driver are known to cause problems with ABORT command handling. If you see a problem, update the PVSCSI driver from the latest stable kernel from https://www.kernel.org/ (version 1.0.3 or above) or the version of the PVSCSI driver contained in the vSphere 6.0 version of VMware Tools (version 1.2.3.0 or above).
LSI MegaRAID megaraid_sas	If you have an LSI MegaRAID controller on your host, do not use the lsi_mr3 driver. Instead, you should use one of the following megaraid_sas drivers (available as VMware downloads): vSphere 5.5U2 – VMware ESXI5.5 scsi-megaraid-sas 6.608.11.00-1OEM SAS Driver for Megaraid SAS vSphere 6.0U1 – VMware ESXi 6.0 scsi-megaraid-sas 6.608.11.00-1OEM SAS Driver for Megaraid SAS

System Time Considerations

Do not adjust the system time for either the ESX host or the associated Datrium storage node controller after initial setup. You can change the system times as necessary before the first time you mount the Datrium datastore. After you mount the datastore, time modifications might cause problems.

To support time synchronization of the Data Node and your vSphere environment, you should configure NTP on your ESXi hosts. See the VMware KB article "Configuring Network Time Protocol (NTP) on ESX/ESXi hosts using the vSphere Client (2012069)". (http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2012069)

DVX VAAI VIB Installation

The VAAI VIB will require ESXi host reboot. The DVX System will not support VAAI (vSphere Storage APIs - Array Integration) storage operations until the next time you reboot the ESXi host.

Installation of the DVX VAAI VIB

The following procedure provides a brief overview of how to use vSphere Update Manager (vSphere 6) or vSphere Lifecycle Manager (LCM) (vSphere 7) to install the DVX VAAI VIB on an ESXi host. For more information about using VUM, see the appropriate VMware documentation.

1. Add a download to the VUM/vLCM repository. Use the Datrium VAAI VIB zip file:

   ```
   Datrium-vaai_2.2.1.0-38226.zip
   ```

2. Validate and apply the download, then download the VIB.
3. Create a separate baseline for the VAAI VIB.
4. Install the Datrium VAAI NAS plugin (Hosts & Clusters page, Update Manager tab).
5. Attach the baseline.
6. Remediate.

Automatic Download and Installation of the DVX VAAI VIB

To use the VUM capability for automatic update, add the DVX VAAI VIB zip file to the patch repository:

```
datrium-vaai_2.2.1.0-38226.zip
```

esxcli Installation of the DVX VAAI VIB

Use the following steps to copy the DVX VAAI VIB and install it on the ESXi host.

Install the VAAI VIB on the host for vSphere 6:

```
esxcli software vib install -d http://mgmt-floating-
IP-address/static/ESXVIBVAAI6
```

Install the VAAI VIB on the host for vSphere 7:

```
esxcli software vib install -d http://mgmt-floating-
IP-address/static/ESXVIBVAAI7
```

DVX Host Resilience

The following sections describe the DVX host environment for resilient operation.

Peer Cache Data Access After SSD Failure

You can continue to run virtual machines and applications at high performance on the host even when one, many, or all local SSDs have failed.

- When an SSD fails on a host that has multiple SSDs, the DVX System reconfigures datrium-reserved partitions on the remaining SSDs and continues serving I/O requests from local flash. See Recovery from Host Drive Failures.
- When the last host SSD fails, the DVX System uses Peer Cache Mode to serve I/O from other hosts in the DVX System and from the Data Pool as needed. Virtual machines and applications will continue to run under Peer Cache Mode until the virtual

machines vMotion to a new host and the local caches are warm, or until you replace the failed SSDs. In this case, you should replace the failed SSDs as soon as possible to restore the normal operating environment

The DVX System uses Peer Cache Mode by default.

To enable or disable flash sharing on a particular host, select “Host settings” on the DVX host page system menu.

When you enable flash sharing on hosts in a DVX System, the DVX System distributes the load between all the hosts that are available for host flash sharing. For example, on host A that has total SSD failure, one virtual machine might obtain its data from host B while another virtual machine might obtain its data from host

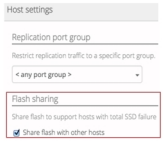

When you configure your DVX System environment for Peer Cache support, you must enable at least two hosts for flash sharing. If a flash-sharing host becomes unavailable for some reason, the DVX System can locate another flash-sharing host to continue serving data to a “flash-less” host that has lost its flash resources. If you have only one host enabled for flash sharing and you unmount the Datrium datastore on that host, any flash-less host that is using the flash-sharing host resources will encounter an APD (All Paths Down) situation.

When you must reboot a flash-sharing host for ESXi maintenance, use the following procedure:

1. Disable flash sharing on the host. The DVX System will continue to use other flash-sharing hosts to provide data to hosts that require it.
2. Put the host into maintenance mode.
3. Perform the maintenance task(s).
4. Take the host out of maintenance mode.
5. Enable flash sharing on the host. The DVX System will resume using the host for peer cache support.

Peer cache operation will affect host performance monitoring:

- Flash-less host latency calculations do not include the network latency between the flash-less and flash-sharing hosts. The latency for a flash-less host is the same as the latency calculation for the host that is sharing its flash.
- The DVX System counts virtual machines running on a flash-less host twice. The DVX UI (GUI and CLI) shows those virtual machines as active on both the flash-less host and on the shared-flash host.

DVX Partitions on Host Drives

The DVX Hyperdriver runs on each ESXi host that uses Datrium storage. When you add a host to a DVX System, you select drives for DVX use, and the DVX System creates partitions on those drives. The DVX creates two kinds of Datrium partitions on local host drive(s) – reserved partitions and cache partitions. The DVX System will not use an ESXi scratch drive.

Datrium Partition on ESXi hosts	Comment
datrium-reserved-sys datrium-reserved-core	The DVX System creates two types of reserved partitions on a host. The name of the reserved partitions start with the prefixes “datrium-reserved-sys” and “datrium-reserved-core”. The datrium-reserved-sys partition requires 16 gigabytes of storage space. The datrium-reserved-core partition requires 130 gigabytes of storage space. DVX puts the datrium-reserved partitions on the ESXi boot drive or on selected SSDs. It will not create datrium-reserved partitions on external drives, HDDs, SAN storage, or boot SD cards.
DVX Cache	The DVX cache partition is a raw partition. The DVX System creates one or more cache partitions on a host. The DVX System can put a cache partition on an SSD boot drive and it can use additional SSDs for the cache.

The location of the partitions depends on the ESXi deployment type and the boot drive type. The following table hows the support for ESXi deployment types.

ESXi Deployment Type	DVX Support
Embedded (visor-usb)	USB boot drive – DVX creates reserved and cache partitions on host SSD(s).
Installable (visor-thin)	SSD boot drive, single SSD – DVX creates datrium-reserved and cache partitions on the SSD boot drive. SSD boot drive, multiple SSDs – DVX creates the datrium-reserved partitions on the SSD boot drive if space is available. It creates a datrium-reserved-sys partition on a second SSD to act as a backup to the primary partition. It creates the cache partition(s) on all selected SSDs including the boot drive.
PXE (visor-pxe)	DVX does not support PXE deployment.
Shared SAN (boot from SAN)	DVX creates datrium-reserved partitions on one or more host SSDs.

Each ESXi host requires at least one SSD for DVX use. Datrium recommends a minimum of two SSDs on each ESXi host. The additional SSDs increase the reliability of the system.

Virtual Flash Resource Management in vCenter

The vCenter Server reports on host virtual flash resources, including host flash that is dedicated to DVX System use.

vSphere does not use DVX host flash space for any caching, although the display implies that the space is available for vSphere use. After you have selected host flash for DVX use, you can ignore any vCenter representation of DVX flash resources.

To verify DVX System use of host flash resources, use the `esxcli datrium check` command on the host. The command output indicates that the vSphere flash read cache is not enabled. For information about this command, see Pre-Day-Zero Host Check.

```
>> esxcli datrium check
Test                              Value                                                Result
--------------------------------  ---------------------------------------------------  ------
CPU                               24 cores                                             OK
Memory                            255.87 GiB                                           OK
NIC speed                         10 Gbps or faster                                    OK
ESXi version                      6.0.0 (3620759)                                      OK
ESXi boot type                    visor-thin                                           OK
Lockdown mode                     Not enabled                                          OK
RAM disk                          Can create 1000 MiB RAM disk                         OK
Root partition                    32512 KiB free                                       OK
Scratch location                  /vmfs/volumes/573f4787-5ca51e2b-1478-5cb9019ad89c     OK
Persistent logging                Yes                                                  OK
vSphere Flash Read Cache (vFRC)   Not enabled                                          OK
Content-Based Read Cache (CBRC)   Not enabled                                          OK
vSAN                              Not enabled                                          OK
Incompatible VIBs                 Not found                                            OK
RAID controller CLI               HP hpssacli installed                                OK
Incompatible storage drivers      Not found                                            OK
```

You can also use the `esxcli storage vmfs extent list` command, which will show the datrium-reserved partitions on the drive(s).

Recovery from Host Drive Failures

Recovery from a host drive failure is determined by ESXi deployment type, the number of SSDs dedicated to DVX storage, the setting for Peer Cache Data Access After SSD Failure, and vSphere.

Deployment Type	Virtual Machine Recovery
Installable ESXi	DVX host configuration creates the datrium-reserved partitions on the SSD boot drive. If the boot drive fails, vSphere HA (High Availability) will restart the virtual machines from the failed host on another host.
Embedded ESXi and Shared SAN	**Single SSD** – If the drive containing the datrium-reserved partitions fails and there are no hosts available for peer cache data access, the Datrium datastore will be inaccessible and the virtual machines on the host will become unresponsive. vSphere recovery operations: • vSphere 5.5 – Manual intervention is needed to restart the virtual machines on another ESXi host. • vSphere 6.0 and later – Make sure that the "Protect against Storage Connectivity Loss" option is enabled. This feature supports

Deployment Type	Virtual Machine Recovery
	automatic restart of virtual machines by vSphere HA. **Multiple SSD** – If the drive containing the datrium-reserved-sys and datrium-reserved-cores partitions fails, the DVX System uses the additional SSD(s). The Datrium datastore will be temporarily unavailable on that host for up to 2 minutes. Two drive configuration – The DVX System creates datrium-reserved partitions and a DVX cache partition on one drive. The second drive contains a DVX cache partition and a datrium-reserved-sys partition that is a mirror of the corresponding partition on the primary drive. If the -sys and -cores SSD fails, the DVX System builds a datrium-reserved-cores partition on the remaining SSD, but after the failure there is no redundancy. Three or more SSDs – The DVX System can survive the failure of the drive that contains the -sys and -cores partitions as long as there are two SSDs containing datrium-reserved-sys partitions.

Virtual Machine Settings

The following section describes virtual machine settings that affect DVX System operation.

vSphere HA – VM Component Protection (APD)

vSphere HA and VM Component Protection (VMCP) supports recovery from datastore All Paths Down (APD) events. If you are running vSphere 6.0 or later, you can use VMCP to determine if host access to the Datrium datastore is in an APD state. If VMCP is enabled and vSphere detects APD events on a host, vSphere HA can restart that host's virtual machines on another host that is connected to the Datrium datastore. To use VMCP to detect and respond to datastore APD events, use the VM component settings shown in the following table to configure the vSphere HA properties on the vCenter Server.

VM Component Protection

	Protection against Storage Connectivity Loss	ON
Virtual Machine Response		
	VM restart priority	High
	Response for host isolation	Power off and restart VMs
	Response for Datastore with Permanent Device Loss (PDL)	ON
	Response for Datastore with All Paths Down (APD)	Power off and restart VMs (aggressive)
	Delay for VM failover for APD	3 minutes
	Response for APD recovery after APD timeout	Disabled
	VM monitoring sensitivity	High

Datrium recommends the vSphere default of a three minute delay for VM failover for APD.

Oracle RAC Support

The DVX System supports an Oracle RAC installation. For Oracle RAC, use thick-provisioned eager zeroed VMDK format. The DVX System supports this format only for Oracle RAC. The DVX System support for Oracle RAC requires the DVX VAAI VIB version 2.1.0.0.

The DVX System supports Oracle RAC only on SSD-based Data Nodes.

To enable DVX System support for Oracle RAC operation, for each virtual machine in the Oracle RAC installation, each shared disk (VMDK) must be thick provision eager zeroed and have the multi-writer flag set. The following sequence shows how to set the multi-writer flag.

- The virtual machine must be powered off.
- In the Virtual Hardware settings, set Sharing to Multi-Writer. Use this setting only for Oracle RAC installations.
- Select Edit Settings in the VM Hardware frame of the virtual machine Summary tab.

- Select a disk to be used for the Oracle RAC installation. The following figure shows the disk type and sharing settings in the Virtual Hardware tab of the Edit Settings dialog.

Mounting/Unmounting a DVX Datastore

To mount or unmount a DVX datastore, navigate to the vSphere Web Client host page.

1. In the Datrium DVX tab, click on the DVX menu icon (gear icon) and select “Mount/unmount datastores”.
2. Select or deselect the datastore to mount or unmount it. Note that you cannot unmount a datastore on a host that has virtual machines on it.

Removing an SSD from DVX Use

To remove an SSD from DVX use and maintain DVX storage activity, the host must have sufficient remaining storage drive capacity for DVX use after the SSD has been removed. (See SSD Requirements.) If you intend to remove the SSD(s) that contain DVX partitions when there is not sufficient remaining space, call Datrium Support.

Before you remove an SSD from DVX use, you must do one of the following:

- Put the host in maintenance mode
  OR
- Unmount the Datrium datastore. (Use the vSphere Web Client to unmount the Datrium datastore.)

To remove an SSD from DVX use, display the host configuration dialog and deselect the SSD.

- If you deselect the SSD that contains the DVX Hyperdriver reserved partitions, the DVX System will move the Hyperdriver partitions to another drive. If that is not possible, you will not be able to deselect the SSD. In this case, add an SSD to your host to allow relocation of the Hyperdriver partition.

- If the SSD to be removed from DVX use is used as a scratch location, the DVX System will not remove the Hyperdriver partitions from the drive. The scratch location will be maintained.

Moving a Host to a Different DVX System

Use the following steps to move an ESXi host from one DVX System to a different DVX System:

1. Unmount the Datrium datastore on the host.
2. Uninstall the Hyperdriver from the host. For information, see the Datrium KB article "DVX Hyperdriver Software Uninstall Procedure".
3. Use the GUI to select the target Data Node and go through the GUI-based host configuration to add the host to the Data Node.

Datrium Data Protection

The DVX System provides copy data management and protection for virtual machines and files in use, at rest and in transit. To protect data, the DVX UI provides operations for individual virtual machines and files and for *protection groups*. A protection group contains one or more virtual machines and/or files from any host in your DVX System.

The DVX System supports the following data protection operations:

Snapshots	A snapshot represents one or more virtual machines and/or files at a specified point in time.
Clones	A clone is an independent copy of a virtual machine or file. You can create a clone of a live virtual machine, a file, or of snapshots of virtual machines, files, and protection groups. You can also clone a virtual disk.
App-consistent Snapshots and Clones	You can use the Datrium VSS Agent to create application-consistent DVX clones and DVX snapshots.
Elastic Replication	You can create a single replica of a snapshot or you can create a sequence of protection group replicas based on a snapshot schedule. The DVX System also supports replication to Cloud DVX in AWS.
Restore Virtual Machines, Files, and Protection Groups	You can use snapshots to restore virtual machines, files, and protection group contents to a particular point in time. If you are using replication, you can use replica site snapshots for recovery if the source system goes down. You can retrieve single files from the guest operating system environment of a virtual machine at a local site or a remote site. You can use the DVX System as an array-based solution that provides replication support for VMware SRM operations. You can use Datrium ControlShift to orchestrate DR (Disaster

	Recovery) operations for on-premises private clouds and public cloud environments.
Blanket Encryption	The DVX System provides Blanket Encryption for data in use, at rest, and in transit.

Snapshots

The figure below shows examples of the DVX snapshot types:

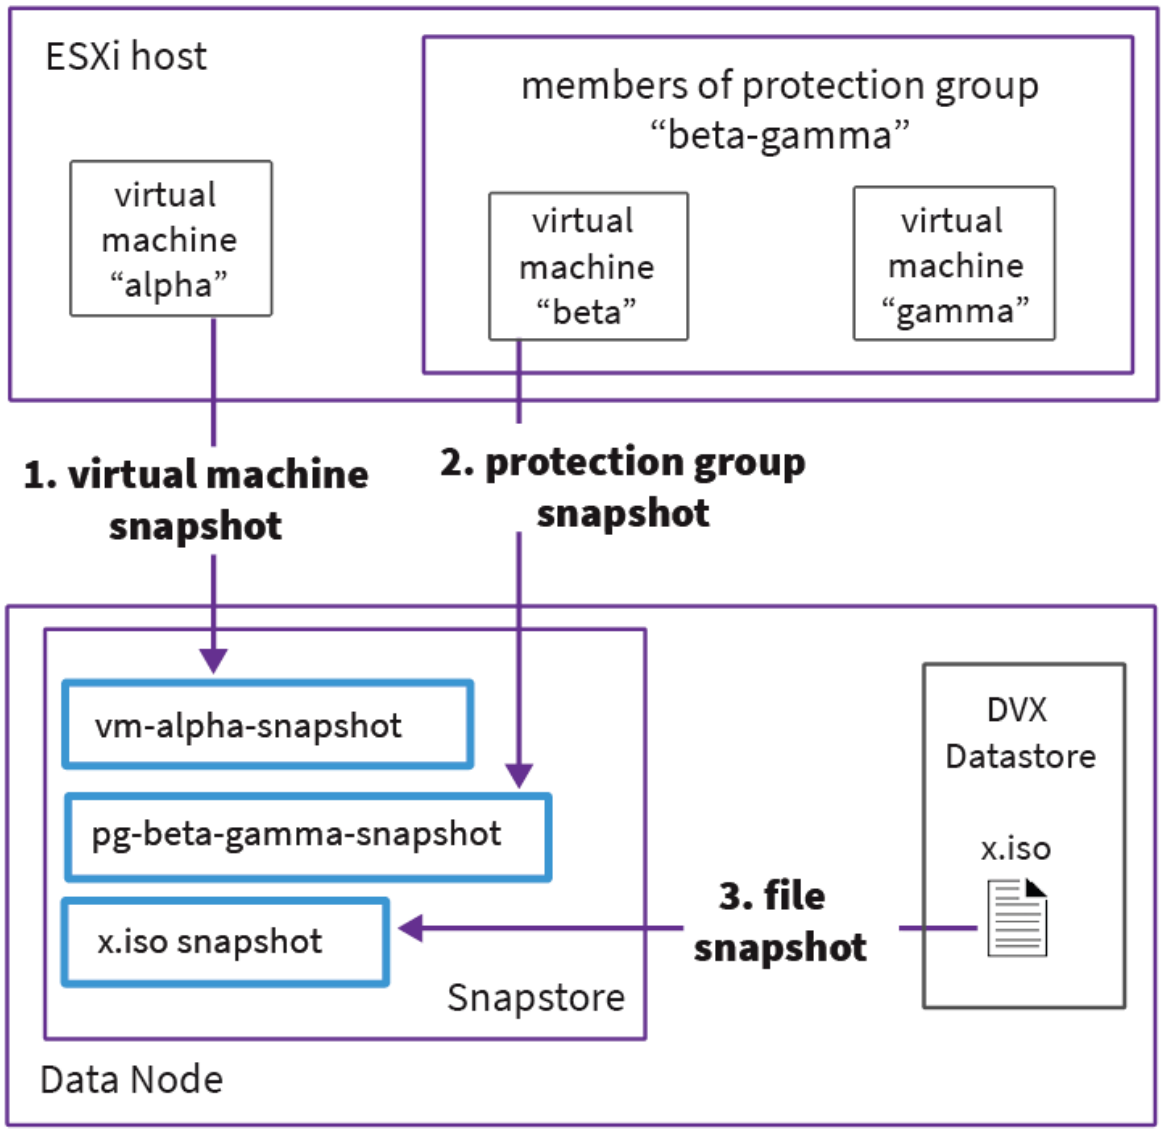

1. A virtual machine snapshot — (“vm-alpha-snapshot”).
2. A protection group snapshot — (“pg-beta-gamma-snapshot”) contains the virtual machines “beta” and “gamma”.
3. A file snapshot — (“x.iso-snapshot”) is based on the file in the DVX datastore.

When you take a snapshot, the DVX System creates the snapshot in the DVX Snapstore. DVX snapshots are not visible in the vSphere Web Client or in the DVX Datastore. To manage DVX snapshots, use the DVX UI.

The DVX System supports two types of snapshots:

- A crash-consistent view of the virtual machine data at a single point in time. When you take a protection group snapshot, the DVX System pauses all I/O to all virtual machines and files in the protection group at the same I/O instant, even if the virtual machines are on different hosts. Protection group contents are snapped simultaneously. A DVX snapshot does not include the contents of memory or transactions in progress.
- An application-consistent view of the virtual machine data at a single point in time. This is available only on Windows Server 2008, 2012, and 2016 virtual machines. To obtain app-consistent snapshots, you must install the Datrium VSS Agent on the Windows virtual machine.

  When you take an app-consistent snapshot of a virtual machine, the Datrium VSS agent uses Microsoft VSS (Volume Shadow Copy Service) to quiesce applications. When activity has ceased, the DVX System pauses I/O to the virtual machine and takes the snapshot. For more information, see App-consistent Snapshots and Clones.

A virtual machine snapshot contains all the files that are necessary to power on the virtual machine. The following files are *not* included in a DVX virtual machine snapshot.

- -digest-flat.vmdk
- -digest.vmd
- .vswp
  .log
- .hlog

- -aux.xml
- any file ending with a tilde (~)
- non-VMware files that have been stored in the same directory as the virtual machine files

If you want snapshots of these additional files, you must identify the file(s) when you take a snapshot. See File Search Pattern.

The DVX System can support a maximum of 2000 protection group snapshots per group, and a total maximum of 1.2 million virtual machine snapshots.

There are two methods of taking snapshots:

- Scheduled Snapshots
- Manual Snapshots

Scheduled Snapshots

To set up recurring snapshots, create a schedule for protection group snapshots. The following sections describe how to create a protection group and snapshot schedule.

Create a Protection Group

A protection group is a collection of one or more virtual machines and/or files. Protection group membership can be any combination of static and/or dynamic elements.

- Dynamic membership — You can specify search patterns to identify virtual machines and files for snapshots (dynamic VM pattern or dynamic file pattern). When you specify a pattern, the DVX System stores that pattern and evaluates it every time that it takes a snapshot of the protection group (by schedule or manually).
- Static membership — You can select virtual machines and files. The selected virtual machines and/or files are added to the protection group immediately. Whenever the DVX System takes a snapshot of the protection group (by schedule or manually), it will take a snapshot of the static group members.

The protection group defines snapshot schedules, snapshot retention policies, and replication and replica retention policies for the protection group contents. You can have a maximum of 50 protection groups in a DVX System.

Creating a Protection Group – GUI

To create a protection group:

1. Select the “Protection” view in the DVX GUI.
2. Click on the “Create” button in the live groups section of the Protection view. The GUI displays the “Create protection group” dialog.
3. The dialog consists of two pages. The first presents fields for the protection group name and group membership. The second page presents fields for snapshot schedules and replication. Use the membership fields to add one or more virtual machines, files, and/or search patterns to a protection group.
4. Enter the protection group name, which can have a maximum of 80 printable characters.

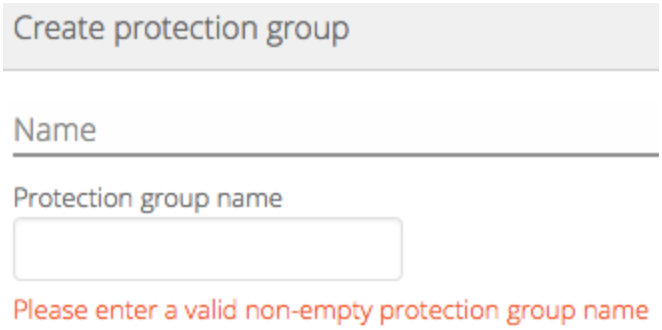

5. **Dynamic VM pattern** – Use the pattern field to add one or more search patterns for virtual machine selection by name. The DVX System will use the pattern(s) to perform a new search each time you use the protection group in a snapshot operation.
6. Enter a pattern string, then press “Preview VMs” to see the results of the search.
7. See VM Search Pattern for information about VM search pattern syntax.

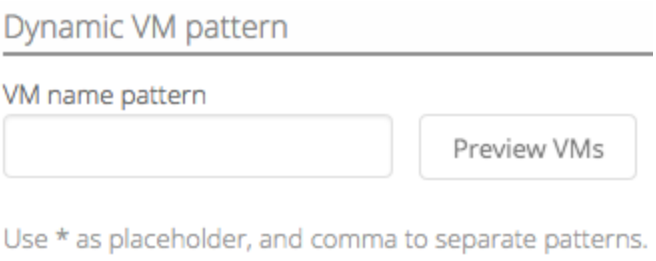

8. **Individual VMs** – Press the “Select individual VMs” button to display the VM selection dialog. The dialog contains a search field, a virtual machine list, and a selected virtual machine list.
9. Enter a virtual machine name expression in the search field to limit the virtual machine list. The expression can include wildcard characters (*, ?, [], !). See VM Search Pattern for information about VM search pattern syntax.
10. Press the magnifying glass icon to execute the search.
11. Select an entry in the virtual machine list to move it to the selected virtual machine list.

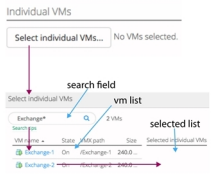

12. **Dynamic file pattern** – Use the pattern field to add one or more search patterns for file or directory selection within the DVX datastore. The DVX System will use the file search pattern(s) each time you use the protection group in a snapshot operation.
13. See File Search Pattern for information about file search pattern syntax.

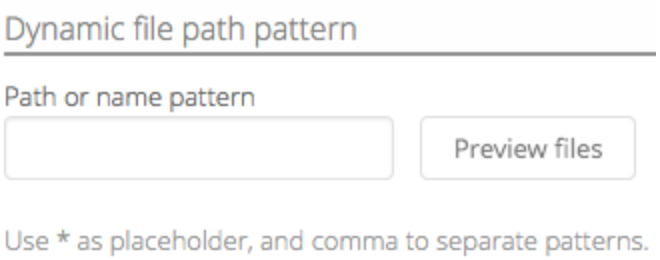

14. Individual files and folders – Press the “Select individual files” button to display the selection dialog. The dialog contains a search field, a file list, and a selected

individual files list.

15. Enter a file path in the search field to limit the file list. A DVX datastore path does not include the ESXi mount point prefix (/vmfs/volumes/*dvx-datastore-label*); it begins with the slash (/) following the DVX datastore label. See File Search Pattern for information about file search pattern syntax.
16. Select an entry in the file list to move it to the selected file list.

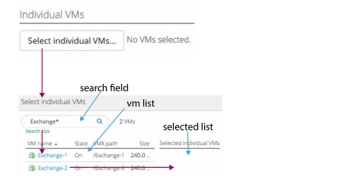

17. After you specify membership, press the “Continue” button to Create a Snapshot Schedule.

VM Search Pattern

> **Note:** The example patterns that are embedded in text use quotes to identify the examples. The quotes are not part of the patterns in these examples. The pattern matching algorithm treats quotes as literal characters that are part of the pattern expression. For example, the pattern “abc” will not match the virtual machine name abc.

A virtual machine search pattern is a string of characters. The DVX System uses Linux-like pathname pattern matching. For example, the pattern “windows[1-5]” matches the set of virtual machines with names in the sequence of “windows1”..”windows5”. The search is case-insensitive.

The following table provides a brief description of the special characters that you can use in a virtual machine name pattern.

Virtual Machine Name Patterns – Special Characters	
*	A single asterisk matches any string of zero or more characters excluding the slash character (/). For example, “windows*” matches all virtual machines with names that start with the string “windows”.

Virtual Machine Name Patterns – Special Characters	
?	A question mark matches any single character. For example, the pattern “linux?” matches any virtual machine with a name that starts with the string “linux” followed by any single character.
[ ]	Square brackets enclose a set of characters. The brackets specify a match of at least one character out of the set. The character set can include one or more ranges of characters. For example, “linux[0-9]” matches anything in the DVX datastore root that begins with the string “linux” and ends with a single digit.
!	The character negation operator. Use this with bracketed expressions to exclude a single character from the pattern. The character negation expression consists of a character set preceded by an exclamation point, enclosed in square brackets: [!*character-set*] The negation expression can include any set of characters and one or more ranges of characters. For example, [!0-9]* matches any virtual machine with a name that does not begin with a number.

There are two types of virtual machine name patterns:

- Inclusion patterns — Specifies virtual machines to be selected. You must specify at least one inclusion pattern.
- Exclusions patterns — Identifies virtual machines to be excluded from the results of inclusion pattern selection. An exclusion pattern begins with the exclamation character (“!”).

The order of patterns in a pattern list does not matter. A pattern list uses a comma (“,”) as a delimiter.

For example, the virtual machine pattern list:

```
Windows*,!Windows7*
```

identifies Windows virtual machines. The combination of inclusion and exclusion patterns matches all virtual machines that begin with the string “Windows”, and excludes virtual machines beginning with the string “Windows7”.

File Search Pattern

> **Note:** The example patterns that are embedded in text use quotes to identify the examples. The quotes are not part of the patterns in these examples. The pattern matching algorithm treats quotes as literal characters that are part of the pattern expression. For example, the pattern “abc” will not match the filename abc.

A file pattern is a string of characters. The DVX System uses Linux-like pathname pattern matching; for example, the pattern “/vm01/*.iso” matches the file “/vm01/rh.iso”. The search is case insensitive.

A pattern that begins with a slash (/) specifies an absolute path, beginning from the root of the DVX datastore. Any other pattern specifies the name of a file or directory that might exist at any location in the datastore; for example, “test.iso” might match any number of test.iso files that occur anywhere in the datastore.

The following table provides a brief description of the most frequently used special characters, plus recursion support.

Filename Patterns – Special Characters and Recursion	
/	Any pattern that includes a forward slash (/) must start with a slash.
*	A single asterisk matches any string of zero or more characters excluding the slash character (/). For example, “/*” matches all files and folders in the root of the DVX datastore. It does not include the contents of the folders.
**	Two asterisks used as a pattern suffix specify a recursive match, including files, folders, and folder contents that match the pattern. For example, “/*_iso/**” displays all root level folders that have the “_iso” suffix, and recursively, all files, folders, and folder content within those root level folders.
?	A question mark matches any single character except the slash character (/).
[ ]	Square brackets enclose a set of characters. The brackets specify a match of one character out of the set. The character set can include one or more ranges of characters. For example, “[0-9]*” matches anything in the DVX datastore root that begins with a number. Bracketed expressions never match “/”, either alone or embedded within a set of characters.

Filename Patterns – Special Characters and Recursion	
!	The character negation operator. Use this with bracketed expressions to exclude a single character from the pattern. The character negation expression consists of a character set preceded by an exclamation point, enclosed in square brackets: [!*character-set*] The negation expression can include any set of characters and one or more ranges of characters. For example, [!0-9]* matches anything in the DVX datastore root that does not begin with a number.

There are two types of file patterns:

- Inclusion patterns — Specifies files to be selected. You must specify at least one inclusion pattern.
- Exclusions patterns — Identifies files to be excluded from the results of inclusion pattern selection. An exclusion pattern begins with the exclamation character (“!”).

The order of patterns in a pattern list does not matter. A pattern list uses a comma (“,”) as a delimiter.

For example, the file pattern list:

```
/iso/*iso,!/iso/Windows7*.iso
```

identifies files in the root directory /iso. The combination of inclusion and exclusion patterns matches all files with the “iso” extension, and excludes iso files beginning with the string “Windows7”.

Creating a Protection Group – CLI

To create a protection group, use this DVX CLI command:

```
protection groups create groupName
```

To add members to a group, use the `protection groups show` command to obtain the group ID, and then use the members add`members add` command to identify virtual machines and/or files, or to specify one or more search patterns.

```
protection groups show
protection groups members add groupID
                                [--vm-path path [path ...]]
                                [--vm-name-pattern pattern [pat-
tern ...]]
                                [--file-path path [path ...]]
                                [--file-name-pattern pattern
[pattern ...]]
```

Specify the protection group ID before any arguments. You can specify any combination of arguments. The final set is the union of the results from the specified arguments. For information about search patterns, see VM Search Pattern and File Search Pattern.

After you create a protection group, create a Snapshot Schedule (below).

Create a Snapshot Schedule

Use a snapshot schedule to define recurring snapshot operation at a particular time based on a recurring interval. A protection group can have a maximum of 10 schedules associated with it.

- The DVX GUI provides schedule fields for a time of day based on the following intervals:
  - Every 30 minutes
  - Hourly
  - Daily
  - Weekly
  - Monthly
- The DVX System supports a minimum schedule frequency of 10 minutes. See Snapshot Schedule – GUI or Snapshot Schedule – CLI .
- The DVX CLI uses cron expressions to define schedules. When you use the DVX GUI to display a schedule that you created or modified with the DVX CLI, the GUI will convert the cron expression into the corresponding interval and time fields in the GUI schedule. If it is not possible to perform the conversion, the GUI displays a read-only version of the cron expression. You must use the CLI to create and modify cron expressions.

- The DVX System default time zone for protection group schedules is UTC. To change the schedule time zone:
- DVX GUI – Use the system menu to display the Configure DVX dialog. The dialog includes the time zone field.
- DVX CLI – Use the command “config time-zone set”.

  ```
  config time-zone set timezone
  ```

  The timezone parameter is one of the three-character time zone codes.

Snapshot Schedule – GUI

The initial schedule display shows the default schedule – take a snapshot every day at 12am (midnight) with a retention period of 4 hours. Use the drop-down menus and text field to modify the interval, time of day, and retention period as needed.

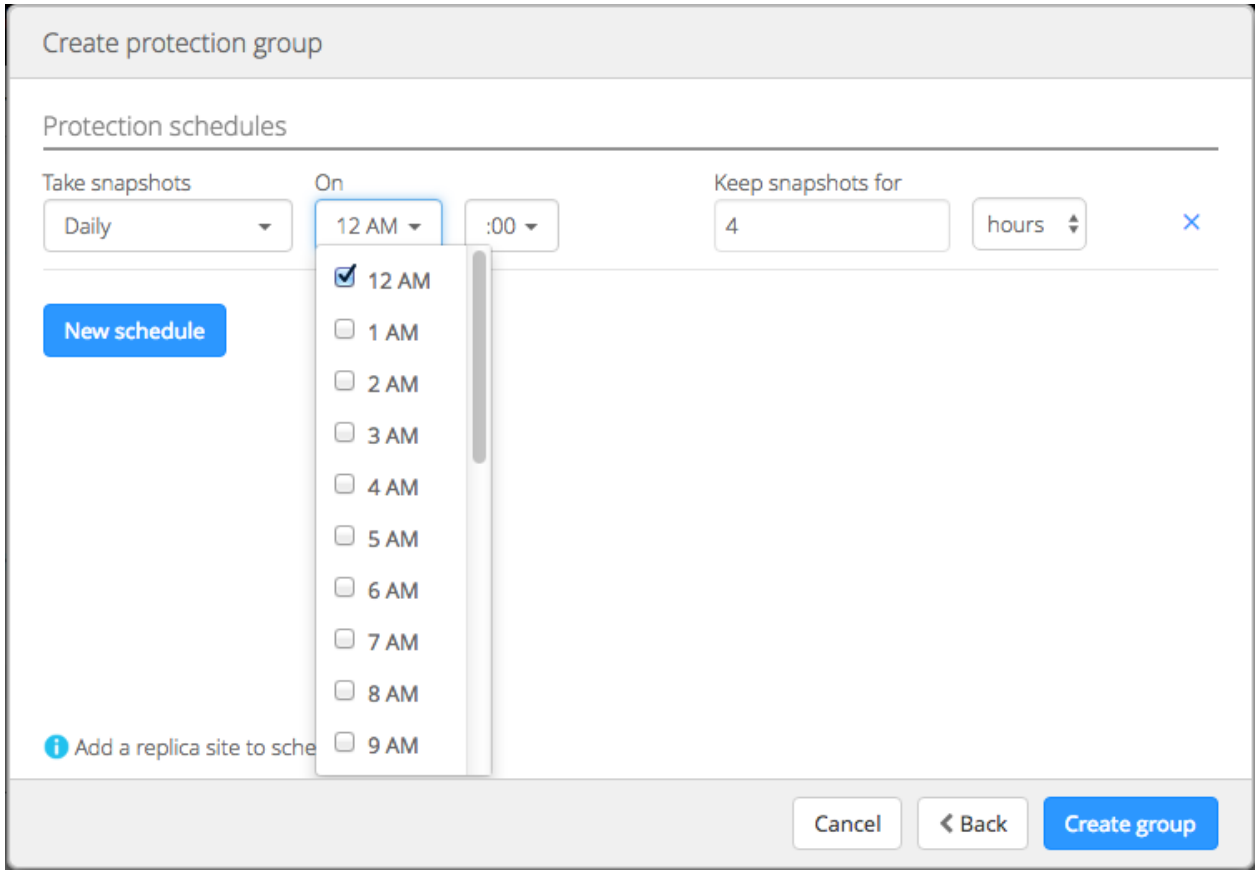

When you finish creating the group, the schedule is automatically enabled. The DVX System will start taking snapshots at the scheduled time.

The DVX System supports a minimum schedule frequency of 10 minutes. To use the GUI to define a 10 minute snapshot schedule, edit the protection group and add a set of schedules with 30 minute intervals, each firing at different time offsets. For example, you can use the settings “0:: and :30”, “:10 and :40”, “:20 and :50”.

Snapshot Schedule – CLI

To add a snapshot schedule to a protection group, use the `protection groups show` command to obtain the group ID, and then use the `protection groups schedules add` command.

```
protection groups show
protection groups schedules add groupID
                                 --schedule-name name
                                 --schedule cronExpression
                                 --retention seconds
                                [--replica-site-names name
[name ...]]
                                [--replica-retention seconds]
                                [--srm-restore-datastore-name
dastorename]
```

- You must specify a protection group ID, a schedule name, a cron expression that determines the frequency of snapshots, and the snapshot retention period. You can also specify Elastic Replication.
- Use the --srm-restore-datastore-name argument only for SRM-enabled protection group schedules. The datastore name identifies the DVX datastore on the recovery site where virtual machines will be restored during an SRM failover operation. See Restore Virtual Machines, Files, and Protection Groups.
- The DVX CLI uses cron expressions to specify the frequency of snapshots. A cron expression is a string that contains 6 or 7 fields that represent a set of times. In the DVX System the schedule describes a sequence of snapshots according to the specified time interval.

```
seconds minutes hours day-of-month month day-of-week
[year]
```

- The fields are separated by spaces. Only the year field is optional. You must provide a non-asterisk value in at least one field. You cannot specify an asterisk in every field. You must specify the value zero (“0”) for the seconds field.

The following tables provide a brief summary of cron expression syntax. For more information, see the cron tutorial.

Cron Field Values and Special Characters							
cron field	seconds	minutes	hours	day of month	month	day of week	[year]
numeric or text	must be zero (0)	0-59	0-23	1-31 (specify “?” if using day of week)	1-12 or JAN-DEC	1-7 or SUN-SAT (specify “?” if using day of month)	empty or 1970-2099
special characters		, - * /	, - * /	, - * ? / L W	, - * /	, - * ? / L #	, - * /

Cron Field Special Character Usage		
Character	Usage	Example
*	Use the asterisk character to select all values for the field.	hours: * selects all hours in the day
?	Use the question mark character in either the day of month or day of week field. The fields are mutually exclusive. Use numeric, textual, and/or special character values for one of the fields and the question mark for the other field.	day of month: 17 day of week : ? day of month: ?

Cron Field Special Character Usage		
		`day of week: 3`
`-`	Use a dash between two values to represent a range of values.	`3-5` `JUN-SEP`
`,`	Use a comma to separate values in a set.	day of week: MON,TUE,WED
`/`	Use a slash to indicate a starting point and increment: start/increment.	minutes: 0/15 selects the following sequence of minutes in an hour: 0,15,30,45
`L`	Use the “L” character to select the last value in the sequence for day of month or day of week.	day of month: 28, 29, 30, 31 (depending on the month) day of week: 7, SUN
`W`	Use the “W” character with a day-of-the-month number to select the weekday nearest to the specified day of the month. (See the cron tutorial for more information.)	day of the month: 10W
`#`	Use the number sign between two values in the day of week field: day-of-week#instance. This represents, within a month, a particular instance of a day of the week.	day of week: 3#1 represents the first Tuesday in the month

The DVX System supports a minimum schedule frequency of 10 minutes. You can use the CLI to define a cron expression that fires every ten minutes. For example, use the cron expression `0/10 * * * *` to set the following sequence of minutes every hour: `0,10,20,30,40,50`.

Manual Snapshots

You can take a manual snapshot of a virtual machine, file, or protection group. A manual snapshot is not associated with a schedule.

- When you take a manual snapshot of a virtual machine or a file, the DVX System puts the snapshot in the Ad Hoc protection group.

- When you take a manual snapshot of a protection group, the snapshot is associated with the protection group.

Manual Snapshot of a Virtual Machine – GUI

To take a manual snapshot of a virtual machine, use the following procedure.

1. Select the “VMs” view. The GUI displays the list of virtual machines in the DVX datastore. When the list is first displayed, the buttons for actions (Clone, Take snapshot, Restore) are disabled.
2. Select a virtual machine row to enable the buttons. There are live links in a virtual machine row (virtual machine name and vmx file path). To select the row, click on white space or plain text. The row will change color to indicate the selection.
3. Click the Take snapshot button.
4. The GUI displays a dialog that contains a field for the snapshot name. The GUI displays the default name – a concatenation of the virtual machine name and the snapshot timestamp. You can edit the field to change the name if necessary.
5. Click on the Take snapshot button in the snapshot name dialog.

A manual snapshot of a virtual machine creates the snapshot in the Ad Hoc protection group.

- To see the Ad Hoc snapshots list, select the Protection view (1) and click on the ad hoc snapshots button (2).
- To see detail about the snapshot, click on the snapshot name in the Ad Hoc snapshots list (3).

Manual Snapshots – Files, Protection Groups – GUI

To take a manual snapshot of a file:

1. Select the Files view.
2. Select a file row.
3. Click on the Take Snapshot button.
4. The resulting snapshot is located in the Ad Hoc protection group.

To take a manual snapshot of a protection group:

1. In the protection view, select a protection group row, or use the protection group's page.
2. Click on the protection group name (shown at right).
3. Click on the take snapshot button.

The resulting snapshot is shown in the snapshot list on the protection group's page. It might be necessary to refresh the page to show the updated list.

Manual Snapshot of a Virtual Machine – CLI

To take a manual snapshot of a virtual machine, use the `vms take-snapshot` command, which requires either the virtual machine ID or the .vmx file path.

1. Use the `vms show` command to display the list of virtual machines, with IDs and file paths.

```
dvx05.controller1>> vms show

------------------------------------------------------------------------------ VM Summary -------------------------------------
VM ID                                    Name               Path                                    Size (GiB)  Power state
---------------------------------------- ------------------ --------------------------------------- ----------- ------------
VM.7.00010c42-fb86-11e6-9197-0050cc7a7a01  n1137b_vdb_23      /n1137b_vdb_23/n1137b_vdb_23.vmx          1200.0  ON
VM.7.01385156-fb86-11e6-9199-0050cc7a7a01  n1497_vdb_69       /n1497_vdb_69/n1497_vdb_69.vmx            1200.0  ON
VM.7.01482080-fb87-11e6-9376-0050cc7a7a01  n1137a_vdb_14      /n1137a_vdb_14/n1137a_vdb_14.vmx          1200.0  ON
VM.7.027e4f10-fb87-11e6-9386-0050cc7a7a01  n1137b_vdb_12      /n1137b_vdb_12/n1137b_vdb_12.vmx          1200.0  ON
VM.7.039d7a2a-fb86-11e6-91a3-0050cc7a7a01  n1137a_vdb_08      /n1137a_vdb_08/n1137a_vdb_08.vmx          1200.0  ON
VM.7.03b50d6a-fb87-11e6-9388-0050cc7a7a01  n1137b_vdb_47      /n1137b_vdb_47/n1137b_vdb_47.vmx          1200.0  ON
```

   You can reduce the size of the list by using the path, pattern, snapstore, or maximum count arguments. See the DVX Command Line Interface manual for descriptions of the vms show arguments.

2. Use either the DVX virtual machine ID or the .vmx file path to identify the virtual machine for the snapshot. You must also specify a snapshot name and either an expiration date or a retention period.

```
dvx05.controller1>> vms take-snapshot --vm-path /n1137b_vdb_23/n1137b_vdb_23.vmx --snapshot-name test-snap --retention 300
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Waiting for operation to finish
Task ID: snapper-008d5216-fd01-11e6-b8a7-0050cc7a84b7

------------------------------------------------------------------------------------------------------------ Task details -------------------
                  Start  ID                                          Kind             State       Details      Progress (
----------------------- ------------------------------------------- ---------------- ----------- ------------ -----------
2017-02-27T15:25:42 UTC  snapper-008d5216-fd01-11e6-b8a7-0050cc7a84b7  Take VM snapshot  SUCCESS     --                      1
------------------------ ------------------------------------------- ---------------- ----------- ------------ -----------
```

- Use the `--vm-path` argument to identify a virtual machine by file path. A virtual machine path expression is the DVX datastore path to the virtual machine configuration file. For vSphere ESXi, the path includes the .vmx file name and extension. The path does not include the ESXi mount point prefix (/vmfs/volumes/*dvx-datastore-label*); it begins with the slash (/) following the the DVX datastore label. For example, the DVX datastore path for "/vmfs/volumes/dvx-Datastore1/vm01/vm01.vmx" is "/vm01/vm01.vmx".
- The snapshot name can have a maximum of 80 printable characters.
- Use the `--retention` argument to specify the amount of time that the DVX System will keep the protection group snapshot. Specify the time value in seconds. To keep the snapshot forever, specify the string “forever”.
- You cannot specify `--retention` and `--expiration` on the same command line.
- By default, the command waits for the snapshot task to complete before it returns to the CLI prompt. If you use the `--no-wait` argument, the CLI will start the task, then display the task ID and return to the CLI prompt. You can use the task ID with the “protection tasks show” command to display information about the snapshot task.

3. To display information about the snapshot, use the `vms snapshots show` command.

   The `vms snapshots show` command requires either a virtual machine ID or a protection group snapshot ID. The example above uses a virtual machine ID. To obtain the virtual machine ID, use the command `vms show`.

Context	Manual Snapshot Command
virtual machine	`vms takeSnapshot`
file	`files takeSnapshot`
protection group	`protection groups takeSnapshot`

Clones

A DVX clone operation creates one or more virtual machines and/or files in the DVX NFS namespace.

You can clone the following entities:

- Live virtual machine (powered on or powered off)
- Virtual machine snapshot
- File located in the DVX Datastore
- File snapshot
- Protection group snapshot.
- Virtual disk

Clones are visible in the DVX Datastore. Virtual machine clones (including clones of virtual machines that are part of a protection group) are not visible in the vSphere inventory. You must add virtual machine clones to the vSphere inventory manually.

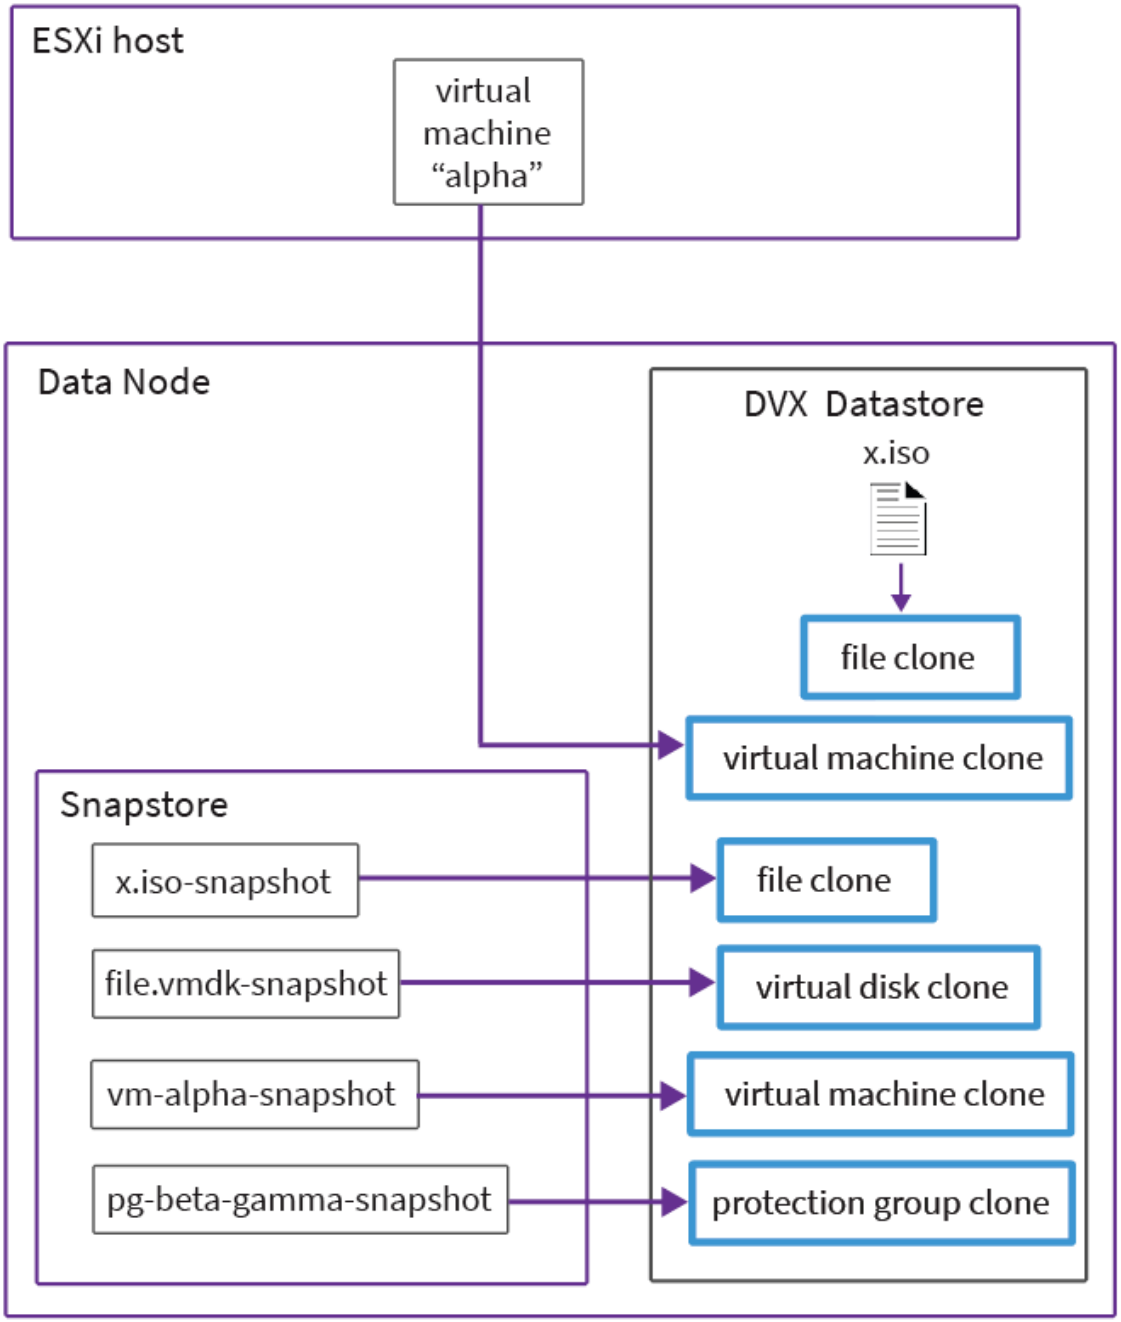

The DVX UI provides the following clone operations:

- Clone a virtual machine from a virtual machine instance or a virtual machine snapshot.
- Clone a file from a file in the datastore or a file snapshot.

- Clone the contents of a protection group from a protection group snapshot.
- Clone a virtual disk from the snapshot of a virtual disk file.

> **Note:** The DVX system does not support cloning between two or more DVX systems.

Clone Operations – GUI

- Clone Virtual Machine (GUI)
- Clone Virtual Machine Snapshot (GUI)
- Clone Protection Group Snapshot (GUI)

> **Note:** The DVX system does not support cloning between two or more DVX systems.

Clone Virtual Machine (GUI)

To make a clone of an instance of a virtual machine:

1. Select the VMs view.
2. Select a virtual machine row.
3. Click the Clone button. The GUI displays the clone dialog.
4. The dialog shows the default clone name, using the "_clone" suffix, and the default location. You can change these values as desired.
5. Click the Clone button. The DVX System creates a snapshot with a retention period of 30 minutes in the Ad Hoc protection group, and it creates the virtual machine clone in the datastore.

You must add the virtual machine clone to the vSphere inventory manually.

Clone Virtual Machine Snapshot (GUI)

To clone a virtual machine from a virtual machine snapshot, select a snapshot from the list of snapshots for the virtual machine:

1. Select the VMs view. The GUI will display the complete list of virtual machines.
2. Select Snapstore view. The GUI will display the list of virtual machines that have snapshots.
3. Select a virtual machine. The GUI will display the virtual machine page.
4. Select Protection. The GUI will display the list of virtual machine snapshots.
5. Select a snapshot row. The GUI enables the Clone button.
6. Click the Clone button.

You must add the virtual machine clone to the vSphere inventory manually.

Clone Protection Group Snapshot (GUI)

1. To make a clone of a protection group snapshot, start in the Protection view
2. In the protection view, click on a protection group name. The GUI will display the protection group page.
3. On the protection group page, click on a snapshot name. The GUI will display the protection group snapshot page.
4. On the protection group snapshot page, click on the “Clone” button.

You must add any virtual machine clones in the protection group to the vSphere inventory manually.

Clone Operations – CLI

- Virtual Machine Clones (CLI)
- File Clones (CLI)
- Protection Group Clones (CLI)

Virtual Machine Clones (CLI)

You can create a clone of a virtual machine from a virtual machine instance or from a snapshot of a virtual machine.

- To create a virtual machine clone from an instance of a virtual machine, use the vms show command to obtain the virtual machine ID or the path to the virtual machine .vmx file.

  ```
  vms show
  vms clone vmID --new-name cloneName
  ```

  or

  ```
  vms show
  vms clone --vm-path vmxFilePath --new-name cloneName
  ```

  The DVX System creates a snapshot with a retention period of 30 minutes in the Ad Hoc protection group, and it creates the virtual machine clone in the datastore.

- To create a virtual machine clone from a virtual machine snapshot, use a virtual machine snapshot ID.
  - To obtain a virtual machine snapshot ID from the list of virtual machine snapshots in the snapstore:

    ```
    vms show --in-snapstore
    vms snapshots show --vm-id vmID
    ```

  - To obtain a virtual machine snapshot ID from the list of virtual machine snapshots associated with a protection group snapshot:

    ```
    protection groups show
    protection groups snapshots show pgID
    vms snapshots show --protection-group-snap-
    shot-id pgSnapID
    ```

- To create the clone:

  ```
  vms snapshots clone vmSnapshotID --new-name clone-name
  ```

By default, the virtual machine clone commands are synchronous. These commands wait for the clone task to complete before returning to the CLI prompt. If you use the --no-wait argument, the CLI will start the task, then display the task ID and return to the CLI prompt. You can use the task ID with the “protection tasks show” command to display information about the clone task.

The virtual machine clone operation creates the clone in the datastore. You must add the virtual machine clone to the vSphere inventory manually.

File Clones (CLI)

You can create a clone of file from a file in the datastore or from a snapshot of a file. You can also use the `file snapshots clone` command to create a clone of a virtual disk. See Clone Operations – CLI

- To create a clone of a file in the datastore, specify the existing file path and a new file path.

  ```
  files clone filePath --new-file-path path [--over-
  write]
  ```

- To create a file clone from a file snapshot, use the file snapshot ID.
  - To obtain a file snapshot ID from the list of file snapshots in the snapstore:

    ```
    files show --in-snapstore  --file-path path |-
    -file-name-pattern patternfiles snapshots show
    --file-path path
    ```

  - To obtain a file snapshot ID from the list of file snapshots associated with a protection group snapshot:

    ```
    protection groups show
    protection groups snapshots show pgID files
    snapshots show --protection-group-snapshot-id
    pgSnapID
    ```

- To create the snapshot:

  ```
  files snapshots clone fileSnapshotID --new-file-path
  path
  ```

By default, the file snapshot clone command is synchronous. The snapshot clone command waits for the clone task to complete before returning to the CLI prompt. If you use the --no-wait argument, the CLI will start the task, then display the task ID and return to the CLI prompt. You can use the task ID with the “protection tasks show” command to display information about the clone task.

Protection Group Clones (CLI)

You can create clones of the contents of a protection group snapshot.

```
protection groups show
protection groups snapshots clone pgSnapshotID
                                  --folder path
                                  --name-prefix prefix
```

By default, the protection group snapshot clone command is synchronous. The snapshot clone command waits for the clone task to complete before it returns to the CLI prompt. If you use the --no-wait argument, the CLI will start the task, then display the task ID and return to the CLI prompt. You can use the task ID with the “protection tasks show” command to display information about the clone task.

Virtual Disk Clones

A virtual disk, for example, “Hard Disk 1” in a virtual machine, is comprised of multiple files. For example, “Hard Disk 1” in virtual machine “vm01” might consist of the files vm01/vm01.vmdk and vm01/vm01-flat.vmdk. As a result, cloning a single file is not sufficient to produce a viable copy of “Hard Disk 1”. The DVX System supports cloning all of a virtual disk’s files in a single step. You do not have to determine the complete set of files and you do not have to clone each file individually.

To clone a virtual disk, you select a virtual disk device from a list of virtual disk devices provided by the DVX UI.

The following image illustrates the process of cloning a virtual disk.

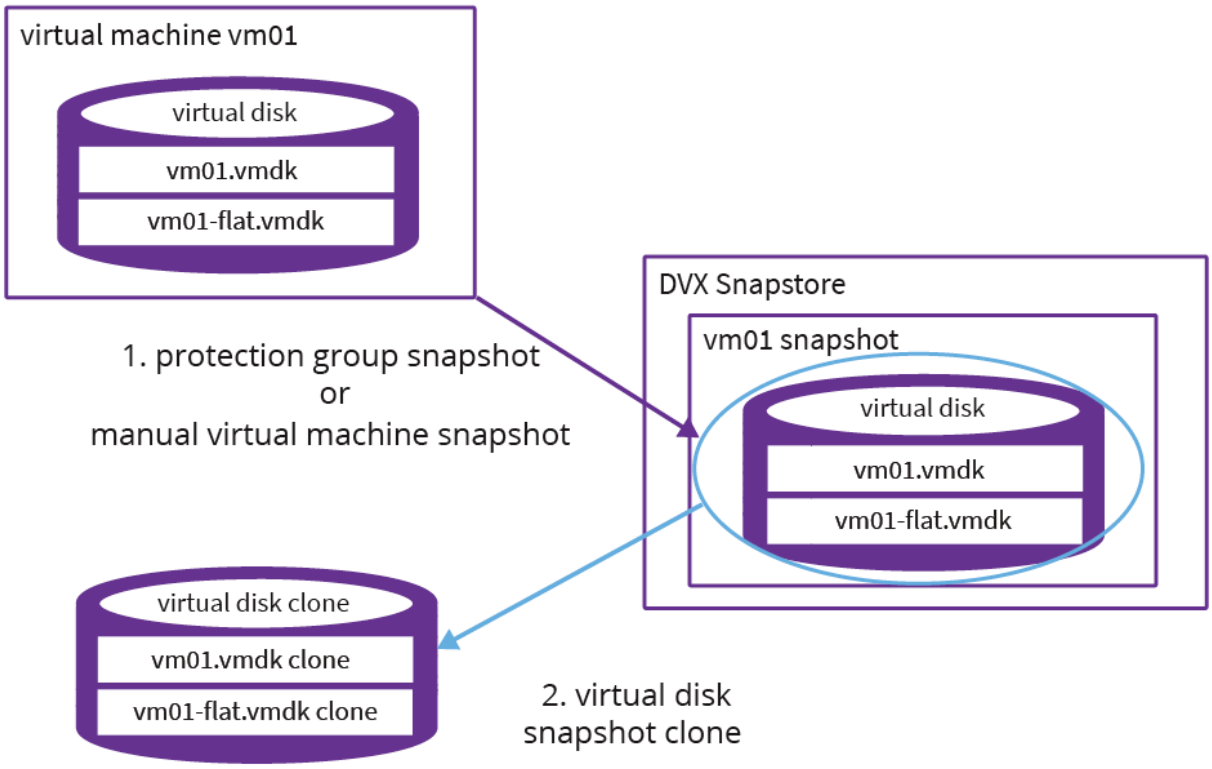

1. A protection group snapshot or a manual virtual machine snapshot operation creates a virtual machine snapshot in the DVX Snapstore. The virtual machine snapshot includes snapshots of all of the files for the virtual disk.
2. A virtual disk clone operation uses the snapshots of the virtual disk files to produce a clone that consists of the entire virtual disk file set.

   For example, a file might be inadvertently deleted in the context of the guest operating system on a virtual machine. To recover the file, use a snapshot of the virtual machine that was taken before the file was deleted. Then attach the disk clone to the virtual machine and copy the file that was deleted from the clone to the permanent disk.

Using the DVX GUI to Clone a Virtual Disk

You can clone a virtual disk from a virtual machine snapshot that is part of a protection group snapshot or you can create the clone from a virtual machine snapshot that you

created manually. The following sequence uses a virtual machine snapshot that is part of a protection group snapshot.

1. Select the Protection view.
2. On the Protection view Summary tab, click on a protection group name.
3. On the protection group page, click on a snapshot name. The snapshot page shows lists of virtual machine snapshots and file snapshots.
4. Select the list entry for a virtual machine snapshot. Note that the virtual machine name is a live link. When you select the list entry, the VM snapshot actions menu will be enabled.
5. Select the “VM virtual disks in snapshot” menu entry.
6. The DVX System displays the primary .vmdk configuration file, virtual disk device, and disk size for each virtual disk in the selected virtual machine snapshot.
7. Select the virtual disk. The Clone virtual disk button is enabled.
8. Click the Clone virtual disk button.
9. In the Clone virtual disk from snapsho” dialog to specify the directory location for the cloned virtual disk files and the base name for these files. The DVX System supplies default names for the directory to contain the cloned disk and for the prefix for the clone files. Change these values if necessary.
10. Click on the “Clone” button to clone the virtual disk.

> **Note:** Use the virtual disk configuration file location when you add the cloned virtual disk to a virtual machine in vCenter. See Add the Cloned Disk to a Virtual Machine.

Using the DVX CLI to Clone a Virtual Disk

To clone a disk file snapshot:

- List the virtual machines that have snapshots in the DVX Snapstore.
- List the snapshots for a particular virtual machine.
- List the virtual disks associated with a particular snapshot.
- Use the virtual machine ID and virtual disk information to clone the virtual disk.

Use the following procedure to clone a virtual disk.

1. Use the `vms show --in-snapstore` command to display the list of virtual machines that have snapshots in the DVX Snapstore.

   The command output shows the virtual machine IDs.

2. Use the virtual machine ID with the `vms snapshots show --vm-id` command.

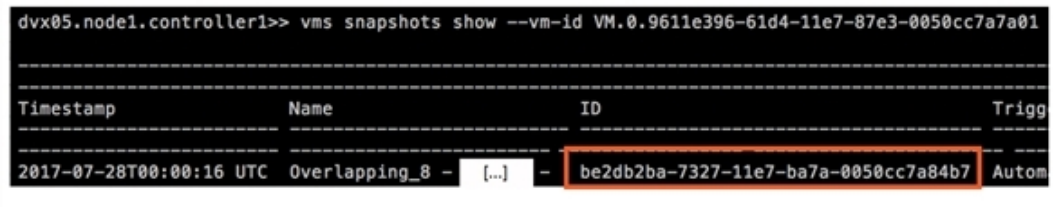

   The command output includes the virtual machine snapshot IDs.

3. Use the virtual machine snapshot ID with the `vms snapshots virtual disks show --vm-snapshot-id` command to display the virtual disk information. The command output shows the identifier for the virtual disk device node.

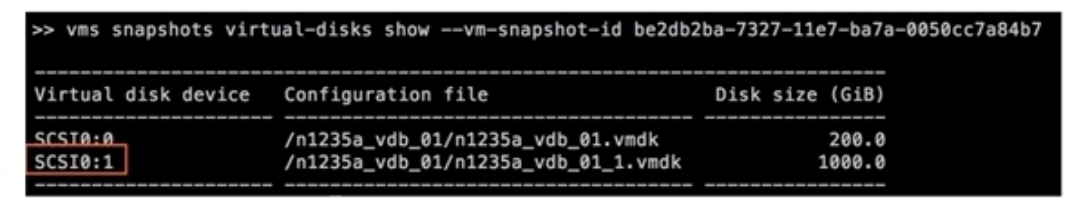

4. Use the `vms snapshots virtual-disks clone` command to create the virtual disk clone.

   The combination of the virtual machine snapshot ID and the virtual disk device node to identify the virtual disk snapshot and generate clones of all of the files in the virtual disk. The command line below has been reformatted for the context of this manual.

```
>> vms snapshots virtual-disks clone --vm-snapshot-id be2db2ba-7327-11e7-ba7a-0050cc7a84b7
>>              --device-node SCSI0:1 --folder-path /vdisk-clones --prefix vdisk-clone-01
Waiting for operation to finish
Task ID: snapper-1f40edda-7330-11e7-b139-0050cc7a84b7

                 Start  ID                                          Kind        State

2017-07-28T01:00:17 UTC  snapper-1f40edda-7330-11e7-b139-0050cc7a84b7  Clone v [...] SUCCESS
```

5. Use the `files show` command to list the clone files. Use the `--file-name-pattern` argument to limit the output. The following example shows the complete clone file set required for the virtual disk.

```
>> files show --file-name-pattern /vdisk-clones/*

------------------------------ Files -------------------------------
File                                      Folder       Size (GiB)
----------------------------------------- ----------- -----------
/vdisk-clones/vdisk-clone-01_1-flat.vmdk  False             1000.0
/vdisk-clones/vdisk-clone-01_1.vmdk       False               ~0.0
----------------------------------------- ----------- -----------
```

   The `files snapshots clone` command specified a new file path of "/vdisk-clones". The DVX System created a directory named "/vdisk-clones" and it used the prefix for the virtual disk clone files ("vdisk-clone-01_1.vmdk" and "vdisk-clones.vmdk").

Add the Cloned Disk to a Virtual Machine

To add the cloned disk to a virtual machine, navigate to the virtual machine page in the vSphere Web Client and select Edit Settings in the VM Hardware frame.

1. Select Existing Hard Disk in the New devices menu.
2. Click Add to display the Select File dialog.
3. Select the virtual disk clone to be added. Use the virtual disk configuration file location provided in the virtual disk clone dialog. Click OK.
4. After you add the disk, use guest operating system procedures to bring the disk online. Both Windows and Linux support hot add/remove of disks, so you can attach/detach a cloned virtual disk without first powering it off.

- Windows — It may be necessary to use the Disk Manager utility to bring a hot added disk online after you attach it to the virtual machine. Similarly, use Disk Manager to take a disk offline before using the vSphere Web Client to detach the disk from the virtual machine.

  If you add the virtual disk clone to the same virtual machine from which it was cloned, you should expect to see "signature collision" warnings in Disk Manager. When you bring the disk online, the Disk Manager resignatures the disk to resolve the collision.
- Linux — See the documentation for your Linux distribution. It may be necessary to scan a device bus to make the newly attached virtual disk visible. Also, you must explicitly mount the disk.

  Before you remove the virtual disk clone from the virtual machine (vSphere Web Client), you must first unmount the file systems associated with the disk and delete the disk device within the guest operating system to prevent the possibility of hanging the operating system.

App-consistent Snapshots and Clones

Datrium provides VSS Agent software that runs on Windows Server 2008, 2012, or 2016 virtual machines, depending on your environment. See the most recent DVX Release Notes for details about VSS compatibility. You can use the Datrium VSS Agent to create application-consistent snapshots and clones.

This section contains the following topics:

- Installing the Datrium VSS Agent
- Creating Application-consistent Snapshots and Clones
- Protection Groups – Application-consistent Snapshots and Clones
- Log Truncation for SQL and Exchange Applications

> **Note:** The Datrium VSS agent is not supported on VMware Virtual NVMe devices.

Installing the Datrium VSS Agent

1. Log in to the Windows virtual machine
2. Use a browser to connect to the management interface of the DVX System. The virtual machine must have network access to the subnet that you are using for the management interface.

   The login page contains an active download link at the bottom – VSS Agent. Click on the link to download the installer to your virtual machine.

   When you run the installer, it will download the Agent from the Data Node and install it on your virtual machine.

3. You must provide the management floating IP address (or fully qualified domain name) and the DVX admin password for authentication.

   The VSS Agent requires network access to the management floating IP address. It also requires network access to the management IP addresses on both controllers (eth1 or eth2, or both eth1 and eth2 for bonded management ports).

4. During the VSS Agent installation, you can assign a port for inbound traffic for the VSS Agent. The default for the inbound traffic is port 8888.

   Datrium VSS operations also require access over port 443 for outbound traffic to the DVX system.

VSS Writers

The Datrium VSS agent recognizes the following VSS application writers.

Task Scheduler Writer VSS Metadata Store Writer Performance Counters Writer System Writer	SqlServerWriter Shadow Copy Optimization Writer ASR Writer COM+ REGDB Writer	BITS Writer Registry Writer WMI Writer Microsoft Exchange Writer

Use the `vssadmin list writers` command to display this list, including details about each entry. If you want a writer that is not in the list, contact Datrium Support and request that the writer be added to the list.

Creating Application-consistent Snapshots and Clones

When you take an app-consistent snapshot or clone of a virtual machine, the Datrium VSS agent uses Microsoft VSS (Volume Shadow Copy Service) to quiesce applications. When activity has ceased, the DVX System pauses I/O to the virtual machine and takes the snapshot or creates the clone.

To create a snapshot or clone a VM:

1. On the VMs view in the DVX GUI, click on the name of a virtual machine to display a virtual machine page.

   The virtual machine page contains Clone and Take snapshot buttons, each of which produces a dialog to accomplish the action.

2. The virtual machine snapshot and clone dialogs contain the Use VSS toggle. If you select the toggle, the DVX System will take an app-consistent snapshot or clone.

You can also use the DVX CLI to create application-consistent snapshots and clones of virtual machines.

```
vms take-snapshot vmID | --vm-path vmx-file-path
                  --snapshot-name name
                  --app-consistent
                  --app-log-truncation

vms clone DvxVmID | --vm-path vmx-file-path
          --new-name clone-name
          --app-consistent
          --app-log-truncation
```

> **Tip:** With the CLI, you can use the `--app-log-truncation` argument to truncate logs for applications with long file names, such as Microsoft SQL and Exchange, when taking an app-consistent snapshot.

For information about virtual machine IDs and file paths, see the description of these commands in *DVX Command Line Interface* guide.

Protection Groups – Application-consistent Snapshots and Clones

You can add VSS-enabled virtual machines to a protection group. When the DVX System takes a protection group snapshot (scheduled or manual) or when you clone the protection group, the snapshot or clone results will include application-consistent snapshots or clones of those virtual machines that are VSS-enabled and which are selected for VSS use.

The DVX System supports the following VSS maximum limits:

- Up to 20 VSS-enabled virtual machines per protection group .
- Up to a total of 100 VSS-enabled virtual machines in a DVX System.

To add VSS-enabled virtual machines to a protection group, use the individual VM selection capability in the DVX GUI.

1. Edit the protection group and click on the “Select Individual VMs” button. The edit dialog shows the virtual machine members that are currently VSS-enabled, next to the selection button (1).
2. The selection dialog shows the list of virtual machines for selection and the list of selected individual virtual machines.
3. To add a virtual machine to the selected list, click on a virtual machine entry (2). (The VM names are active links to the virtual machine pages.)

   You can toggle VSS usage for entries in the selected VM list (3). The selected list also indicates whether or not there is a VSS agent installed on the virtual machines.

You can also use the DVX CLI to add VSS-enabled virtual machines to a protection group.

```
protection groups members add groupID
                                --app-consistent-vm-path path
[path ...]
```

Use the `--app-consistent-vm-path` argument to identify one or more virtual machines on which DVX VSS app-consistent snapshots are enabled. An app-consistent path identifies a single virtual machine by file path. A virtual machine path expression is the DVX datastore path to the virtual machine configuration file. For vSphere ESXi, the path includes the vmx file name and extension. The path does not include the ESXi mount point prefix (/vmfs/volumes/*dvx-datastore-label*); it begins with the slash (/) following the the DVX datastore label.

Log Truncation for SQL and Exchange Applications

Microsoft SQL and Exchange applications provide the capability of continuous logging which can consume substantial disk space. The DVX System provides support for log truncation as part of VSS-enabled virtual machine snapshots. The DVX System supports log truncation for SQL and Exchange applications running on virtual machines in the ESXi environment.

- DVX VSS Log Truncation requires version 1.5 of the Datrium VSS guest agent. See Installing the Datrium VSS Agent.
- Log Truncation uses ODBC (Open DataBase Connectivity) for the SQL Server connection.
- If you are using a 3rd party backup tool to protect your SQL or Exchange application, do not use the Datrium application log truncation feature.

SQL Setup

If you are using log truncation for an SQL application, you must:

- Set permissions for the SQL Login account that will be used for the VSS agent.
- Specify the Login account security context for the Datrium VSS Agent. You must supply the account name and password for the account that has SQL backup permission.

Set permissions for the SQL Login account

In your SQL Database, you need to have a SQL security login user account with the minimum required permissions to allow VSS Agent to take database backups with DVX snapshots.

To set permissions for the SQL security login user account:

1. In the Microsoft SQL Server Management Studio application, add a new SQL Login under Security → Logins.
2. Once this Login is created, right-click the Login and select Properties.
3. In the Login properties dialog, under Select a page on the upper left, select Server Roles.
4. Under the Server roles panel to the right, select the "public" role.
5. Next, under Select a page on the upper left, select User Mapping.
6. On the right, under Users mapped to this Login, select the Login you created for the Datrium VSS Agent.
7. Next, under the Database membership for: panel, select

   `public`— Every user that is set up on the database is part of the public database role.

   `db_backupoperator` — This role will allow the user to take backups of the database.

   The above are the minimum requires permissions to allow the Datrium VSS Agent to make backups of the database.
8. Optional: You can also set two other permissions optionally, to ensure that the VSS Agent cannot make any writes or reads to the database:

   `db_denydatareader` – This role will deny the user access to the database's data, so the user cannot read the data from its tables.

   `db_denydatawriter` – This role will deny the user access to modify the database's data, so the user cannot run any UPDATE or DELETE queries.

9. When you are finished, click OK and then follow the next set of instructions to associate the SQL Login with the Datrium VSS Agent service.

Associate SQL Login account with the Datrium VSS Agent

The following sequence shows how to associate the account with the Datrium VSS Agent.

1. After you install version 1.5 of the Datrium VSS agent on the Windows virtual machine that operates as the SQL Server, log in to that virtual machine.
2. Run `services.msc` to display the service manager listing. Right-click the Datrium Guest Agent entry and select Properties.
3. In the Datrium Guest Agent Properties dialog, select the Log on tab, then click This account.
4. Enter the account name and password for the account that has SQL Server backup permission.
5. Use the backslash ("\") account name format to specify a local account (.*account-name*) or domain account (*domain**account-name*). This example shows a local account (.\administrator).
6. Click OK. Windows displays a message indicating that you must restart the Datrium Guest Agent service.

7. To apply the changes, display the service manager listing and restart the Datrium Guest Agent service.

Log Truncation with VSS-enabled Snapshots

To apply log truncation during VSS-enabled snapshots, in the DVX CLI.

You can configure log truncation for protection group virtual machines that are enabled for application-consistent snapshots a protection group.

```
protection group member add ID --app-consistent-vm-path path
[path ...]
 --app-log-truncation --retention OR --expiration
```

You must specify the `--app-consistent-vm-path` argument to identify one or more virtual machines that have been enabled for application-consistent snapshots. You must also specify the `--app-log-truncation` argument on the same command line.

You can also specify log truncation when taking a manual snapshot of a virtual machine that is VSS-enabled, on which either SQL or Exchange applications are running. Use the `--app-consistent` and`--app-log-truncation` arguments with the `vms take-snapshot` command.

In the following command, specifying either the VM ID **OR** virtual machine path identifies the VSS-enabled virtual machine. If you supply the VM path, then you also need to provide the datastore where the VM lives.

Syntax:

```
vms take-snapshot vmID | --vm-path vmx-file-path --snapshot-
name name
 --app-consistent
 --app-log-truncation
 --datastore (if using --vm-path)
 --retention OR --expiration
```

> **Note:** You cannot specify `--retention` and `--expiration` on the same command line.

For example, using the VM ID with a retention duration of 600 seconds:

```
vms take-snapshot VM.0.45b00110-bff52a-11e9-a285-511q1905577dd
--snapshot-name VSSTest --app-consistent --app-log-truncation
--retention 600
```

For example, using the VM path, the datastore it lives in, and expiration date:

```
vms take-snapshot --vm-path /BO-WinServ-SQL/BO-WinServ-
Test.vmx --datastore Prod --snapshot-name VSSTest2 --app-
consistent --app-log-truncation --expiration 2020-6-6
```

Elastic Replication

The DVX System supports asynchronous replication of DVX snapshots. You can create a single replica of a snapshot or you can create a sequence of protection group replicas based on a snapshot schedule. The following sections describe how to use DVX snapshot replication.

This section contains the following topics:

Protection Group Replication

To use replication, you define one or more replica sites that will store snapshot replicas. After you define a replica site, you can reference the site in protection group schedules or in manual replication operations. If a replica site is referenced in a schedule, every time the DVX System creates a snapshot based on the schedule, it will also send a replica of that snapshot to the replica site.

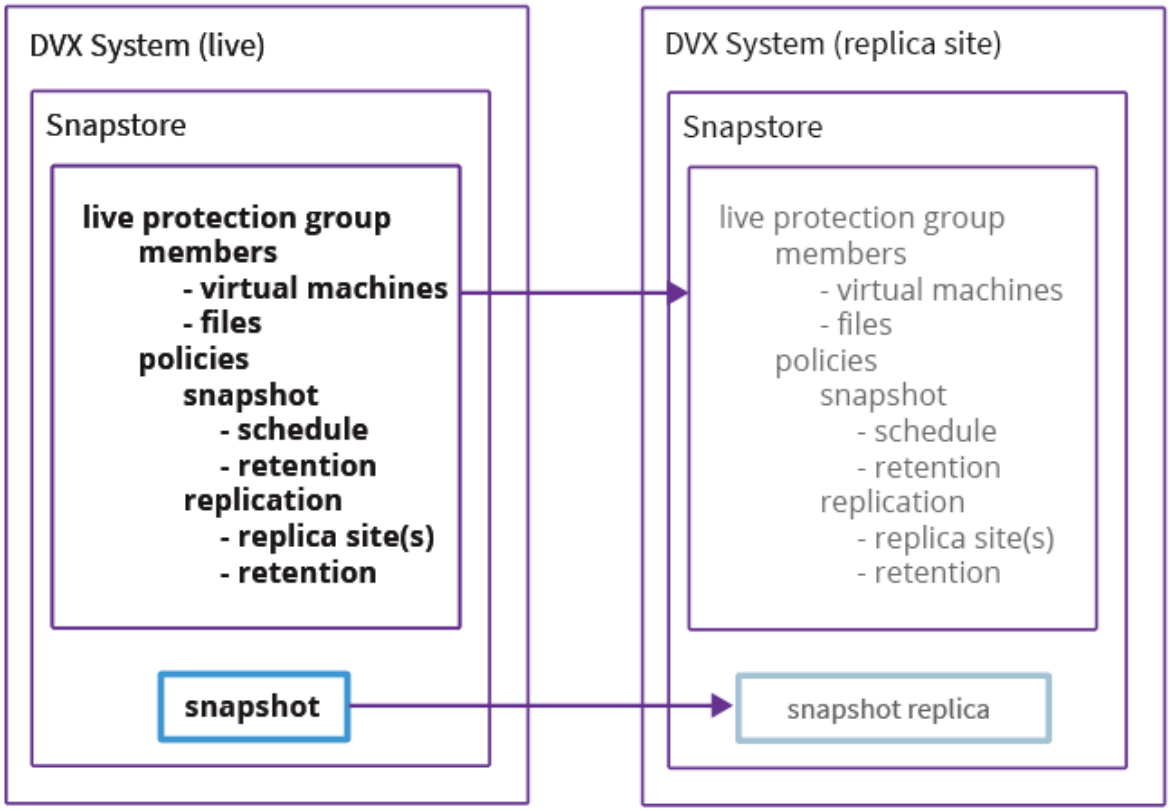

A protection group consists of members (virtual machines and/or files) and policies for snapshots (schedule, snapshot retention), and replication (replica site, replica retention).

When the source DVX System replicates the first snapshot of a protection group, the replica site creates a replica of the protection group. A replica protection group can accept incoming replica snapshots. The replica protection group schedule is inactive. An inactive schedule will not generate snapshot or replication operations. You cannot modify the contents or policies of a replica protection group and you cannot take snapshots of a replica protection group. You can have a maximum of 100 replica protection groups on a site.

You can perform only the following actions on a replica protection group

Clone

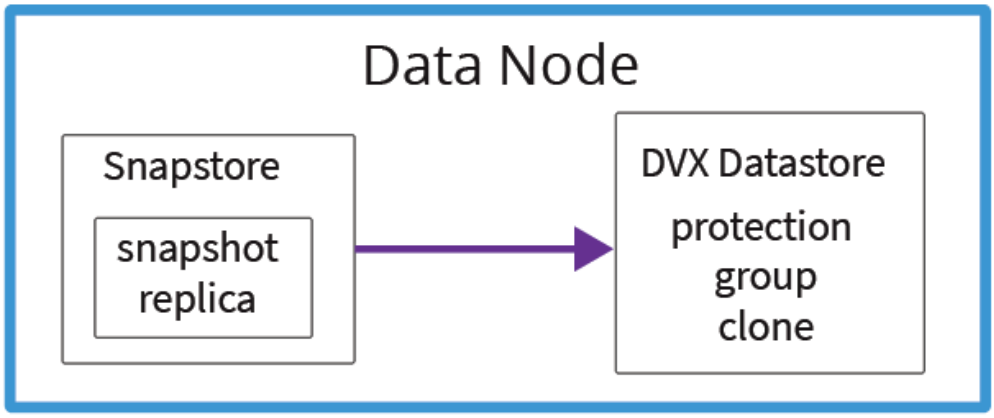

You can clone the contents of a protection group snapshot replica. The virtual machine and/or file clones have new names and are located in the DVX datastore. You must add any virtual machine clones to the vSphere inventory manually.

See Clones.

Restore

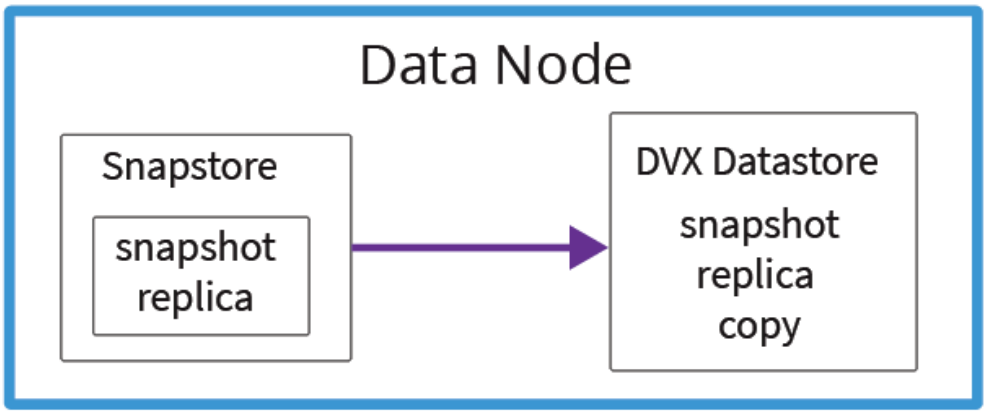

You can restore a protection group from the contents of a protection group snapshot replica. When you restore a protection group, the snapshot contents become visible in the datastore. Restore preserves the file path(s) and name(s) of the snapshot origin. On a vSphere system that does not have previous versions of the protection group virtual

machines in its inventory, you must add any restored virtual machines to the vSphere inventory manually.

See Restore Virtual Machines, Files, and Protection Groups.

Promote

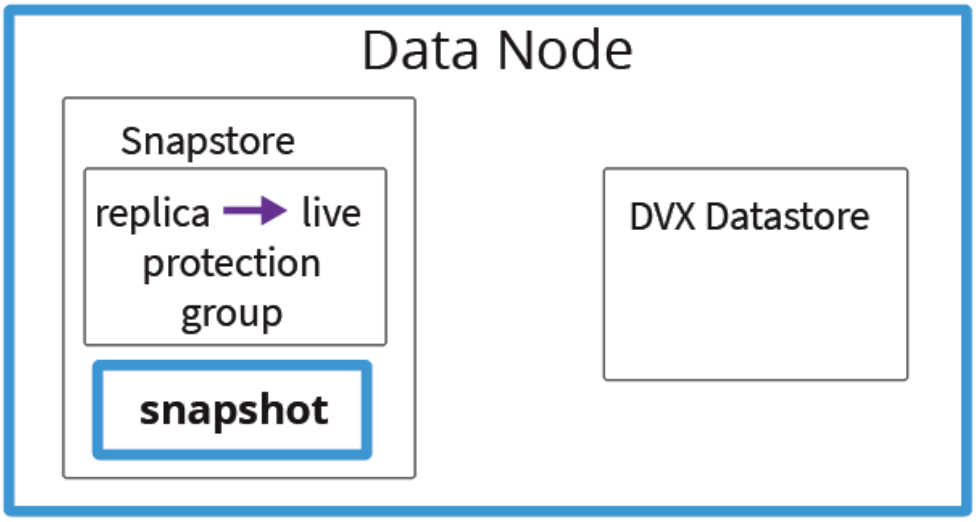

You can promote a replica protection group. Promotion changes a replica protection group into a live protection group. After the group has been promoted, you must enable the schedule(s), add replica sites, and populate the datastore. You can use all of the protection group operations and the group acts as the new origin for future replication operations.

It is important to maintain only one live protection group at a time in a set of associated protection groups – a live group and its replicas. When you promote a protection group, you should demote the other live group to change it to a replica group.

See Disaster Recovery for more information.

The following actions on the live source protection group will be mirrored on the related replica group(s):

- Change the protection group name.
- Change the replica retention.

You can perform the following actions on replica snapshots:

- Replicate a snapshot from a replica protection group.
- Edit the properties of a snapshot in a replica protection group.
- Delete a snapshot from a replica protection group.

Replica Site Definitions

To use DVX snapshot replication, you must create one or more replica site definitions. A replica site definition identifies the destination for *outgoing* replication traffic. After you have defined a replica site, you can reference it in snapshot schedules and replication commands.

To add a replica site, the DVX System must verify that the destination will accept replication traffic and that there is connectivity between the source and destination sites.

- The DVX System supports a maximum of 16 replica sites for both *outgoing* and *incoming* replication traffic. When you successfully add a replica site to a source DVX System, the destination site will accept incoming replication traffic from the source. If the destination site has already accepted the maximum number of incoming source sites, the destination site will reject the request and the command will fail.
- Snapshot replica traffic – Source site hosts use snapshot data from local flash to send snapshot replicas to destination hosts. The destination DVX System stores the replicas in the destination site Snapstore. A source host will use a snapshot from the source Snapstore when the snapshot is not in local flash, for example, in the case of manual replication of an old snapshot.

  The DVX System uses different techniques to reduce the amount of replication traffic between sites. The DVX System applies global deduplication to the data before replication, so that it replicates only those changes that are not already on the destination site. The source site sends only the differences between the snapshot that it is replicating and the latest snapshot that both sites already have in common. The DVX System replicates data in its native compressed form without rehydration.

- Snapshot metadata traffic – The source site Data Node sends snapshot metadata directly to the destination site Data Node. The size of the snapshot metadata is roughly equivalent to one percent of the changed data size.

The following picture shows the data paths used in DVX snapshot replication.

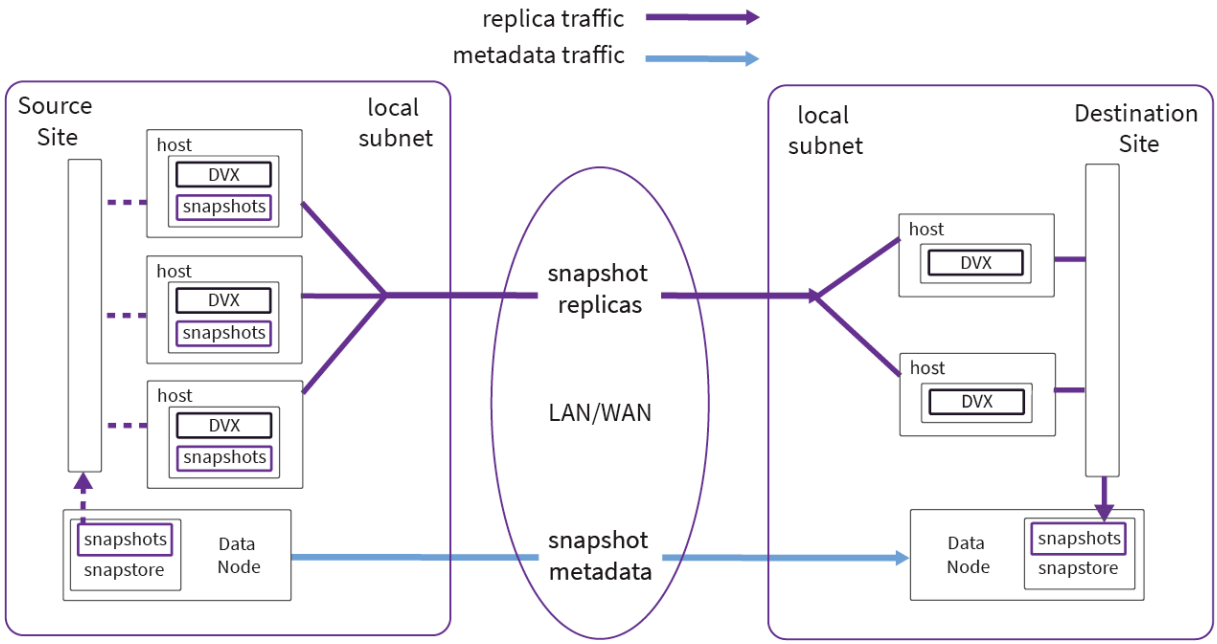

When you add a replica site, the DVX System verifies connectivity between the source and destination systems. DVX replication communication requires direct connections between source and destination hosts, using the actual IP addresses of the hosts. To support replication between remote sites, you might use a VPN; in other situations, a WAN can support the direct access. The DVX System does not currently support replication to destinations behind a NAT firewall.

- The DVX Systems use host resources to transfer snapshot replicas. These resources include host flash, CPU, memory, and NICs.
  - Each source site host must have network access to at least one IP address at every destination site host.
  - You can designate Host Ports for Replication Traffic.
  - Replication tasks are scheduled on hosts that have snapshot contents in host flash. This allows the DVX System to transfer snapshot replicas directly from flash without accessing the Data Node.
- The DVX System uses the Data-Node-to-Data-Node connection to transfer snapshot metadata.
- The source site Data Node must have network access to the replica site Data Node.

Specify the management floating IP address when you define the replica site. The DVX System will use the management interfaces for Data Node communication between the source and destination sites.

In an individual DVX System site, source or destination, the Data Node and hosts are on the same subnet. Datrium recommends that data traffic between the host(s) and Data Node in a single DVX System should be on the same subnet.

Network Ports for On-Premises Replication

To support replication traffic between on-premises source and destination sites, the DVX System uses ports 1525 and 4105. For information about network ports for Cloud DVX Cloud traffic, see Replica Site Definitions.

Host Ports for Replication Traffic

The DVX System can automatically determine the source and destination ports for replication traffic, or it can use port groups that you select for replication traffic to limit the traffic to a specific path. The methodology for choice of ports is based on the replication topology.

There are three replication types supports by DVX:

- In *single stream replication*, where a single source sends replication traffic to a single destination, Datrium recommends that you create replication port groups on both the source and destination hosts.

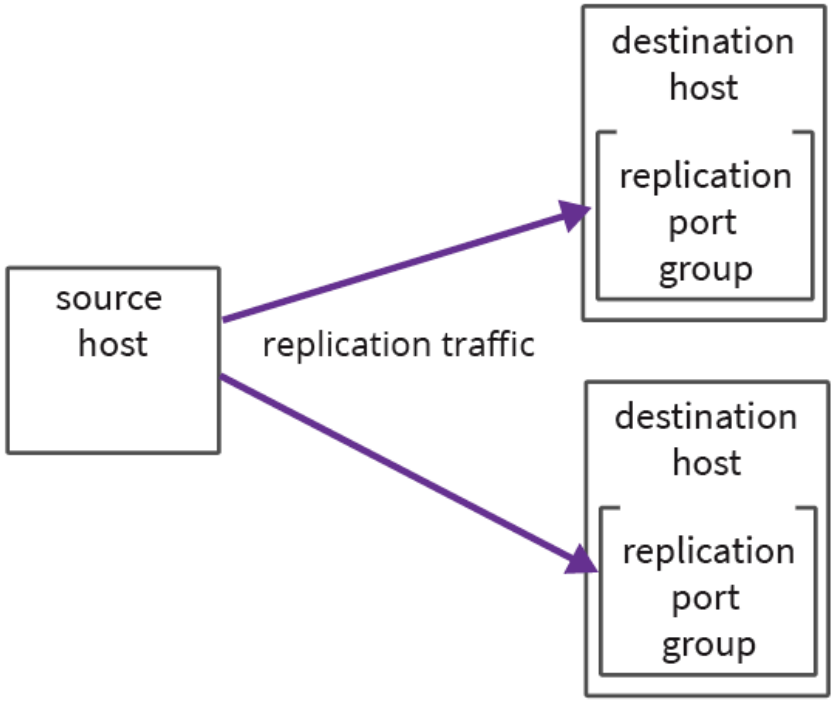

- In *one-to-many replication*, a single source sends replication traffic to more than one destination. In this case, create a port group for each of the destination hosts, but do not use a replication port group for the source host. The DVX System will automatically choose one or more source host ports based on reachability determined for each destination.

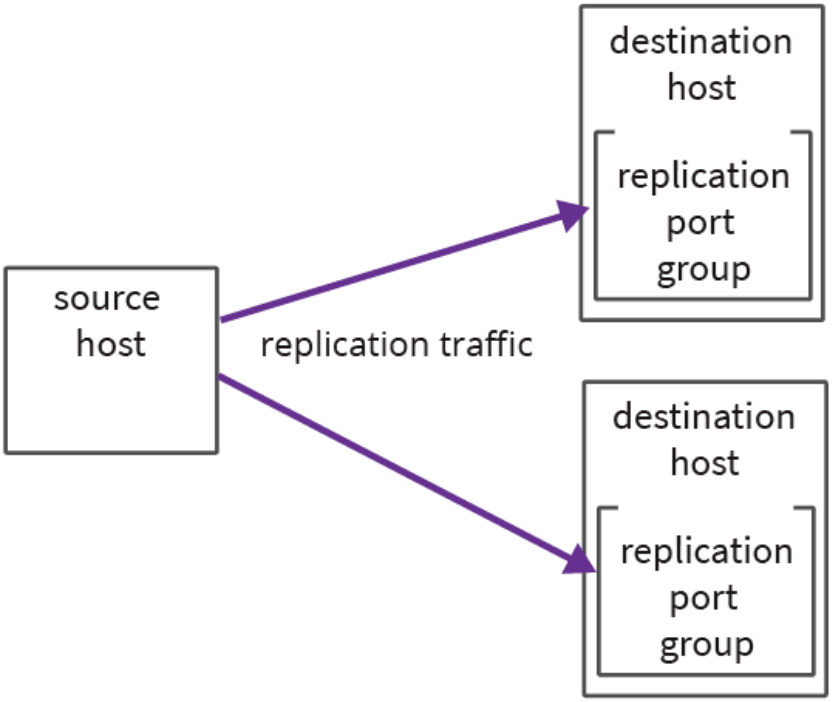

- In *many-to-one replication*, multiple sources send replication traffic to a single destination. In this case, create port groups for the source hosts, but do not use a replication port group for the destination host. The DVX System will automatically choose one or more destination host ports based on reachability determined for each source.

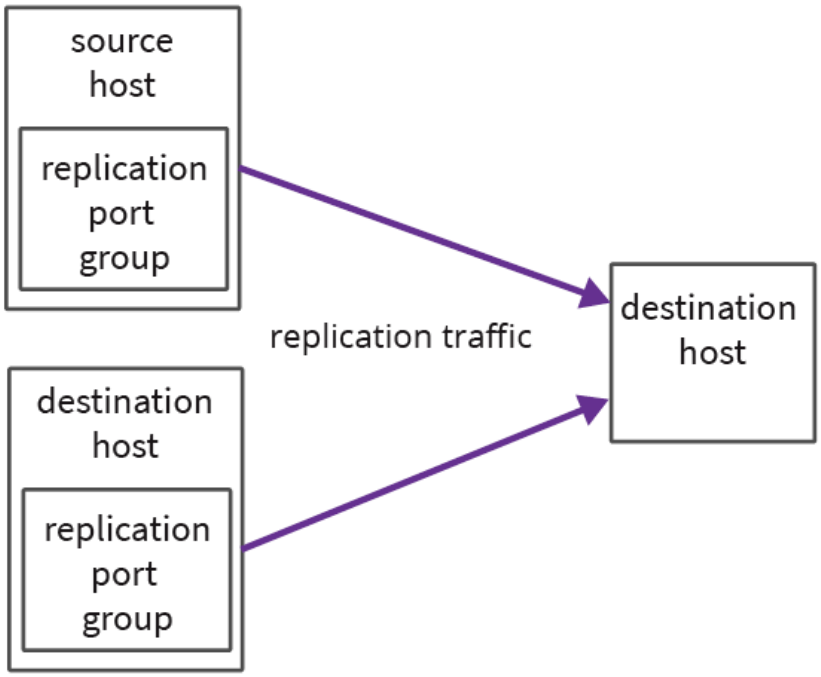

Host Port Group Selection for Replication Traffic

The DVX System provides host port group selection on a per-host basis. You can designate any port group on host(s) where it is needed.

Use the DVX GUI plug-in to the vSphere Web Client to select a destination host port group for replication traffic.

1. On the Datrium DVX view of the vSphere host Monitor tab, click on the gear menu to display the DVX host menu.
2. Select “Host settings” to display the list of replication port groups.
3. Select a port group and click OK. You have the option of using the default “any port group”, or you can use a specific port group. Datrium recommends that you use symmetrical settings for source and destination replication port groups. For source and destination, use any-port-group / any-port group or specific port groups for both sites of the replication pair.

When you create a replica site definition, the DVX System verifies the replication traffic paths.

Replica Site Definition – GUI

In the “Replica sites” section on the Summary tab of the Protection view, click on the “Add” button. The DVX GUI displays the “Add replica destination” dialog.

You must supply the following information:

- The IP address or fully-qualified domain name for the replica site. Use the floating IP address for the Data Node management interface.
- The password for the admin account on the DVX replica site.
- A unique site label. The DVX System will use this label to identify the replica site for DVX operations and differentiate between DVX Systems that have the same name.

Replica Site Definition – CLI

Use the following command to identify a DVX System that will store snapshot replicas.

```
protection replica-sites add floatingIPAddress | domainName
                              --replica-site-name siteName
                             [--password password]
                             [--is-on-cloud]
```

- Use the floating IP address or corresponding fully-qualified domain name for the replica site.
- The DVX System will use the replica site name for various CLI commands and GUI operations.
- The password argument is the admin account password on the DVX System replica site. If you do not specify a password, the DVX CLI will display a password prompt.
- Use the `--is-on-cloud` argument to indicate an AWS Cloud DVX replica site. See DVX CLI Commands for Cloud DVX Storage.

Schedule Replication

To schedule snapshot replication, add a replica site to a protection group snapshot schedule. If a replica site is referenced in a schedule, every time the DVX System creates a snapshot based on the schedule, it will also create a snapshot replica on the replica site(s).

A replication operation is a separate, asynchronous task. Schedules and/or manual snapshot operations might produce multiple snapshots of a protection group before scheduled replication occurs. If there is more than one scheduled snapshot available at the time of the replication operation, the DVX System will replicate the most recent scheduled snapshot.

To add a replica site to a snapshot schedule:

1. Edit the snapshot schedule. Note that it is also possible to set up replication when you create the group.
2. On the Summary tab of the Protection Group view, select a protection group.
3. On the protection group page, click on the Edit group and schedule button.
4. Click on Continue. The first screen in the Edit protection group dialog contains fields for protection group membership. This example procedure assumes there is no change to group membership.
5. For the appropriate schedule:
    a. Select the Replicate snapshots check box.
    b. Select a replication destination.
    c. Enter a number for the time units.
    d. Select a time unit for the retention period.
6. Click on Finish to save the replication settings.

Manual Replication

You can use the DVX GUI or the DVX CLI to create a single replica of a snapshot for a virtual machine, file, or protection group.

Manual Virtual Machine Replication – GUI

To configure manual VM replication:

1. Select the “Protection” view.

2. Click on the “ad hoc snapshots” button. The Ad Hoc protection group contains manual snapshots of independent virtual machines and files. The GUI displays the Ad Hoc virtual machine and files list.
3. Select the row for the virtual machine snapshot to be replicated.
4. Click on the “Replicate” button. The GUI displays the “Replicate Snapshot” dialog.
5. Choose a replication destination. Set a retention period if necessary. If you do not set a retention period, replicas inherit the snapshot retention period from the source site. The replica retention period begins after the snapshot is fully replicated to the destination.
6. Click on the “Replicate” button. You can also replicate a snapshot of a virtual machine that is part of a protection group.

Manual Virtual Machine Replication – CLI

To create a single snapshot replica of a virtual machine snapshot, use the following commands:

```
vms snapshots replicate snapshotID --replica-site siteName
                                    [--retention seconds |
forever]
```

You must specify a virtual machine snapshot ID and a replica site.

To obtain a virtual machine snapshot ID, use the following commands:

```
vms show --in-snapstore
vms snapshots show --vm-ID ID
```

or

```
protection groups show
protection groups snapshots show groupID
vms snapshots show --protection-group-snapshot-id pgSnapID
```

The replica-site argument identifies a remote DVX System that will store the snapshot replica. To obtain a replica site name, use the `protection replica-sites show` command to display the replica-site definitions.

Use the `--retention` argument to specify the number of seconds to retain the replica, after it has been stored on the replica site. To retain the replica forever, specify the string “forever”.

Managing Replication Traffic

By default, protection groups and protection group schedules are enabled when you create them. When you add a replica site reference to a schedule, replication is automatically enabled. Every time the DVX System takes a snapshot according to that schedule, it will send a replica to the replica site.

Replication Progress

When the DVX system is replicating protection groups, the replication task moves through these three general phases:

- From **0 to 14%**, The source DVX is computing the difference between the current and the previously-replicated snapshot(s). The destination DVX then checks if has any data from the difference, to further reduce the amount of data to be sent.
- From **14 to 75%**, DVX is sending replication data over the network.
- After **75%**, DVX is applying the received data deltas at the destination. In some cases, the progress of the replication task might seem like it has stopped, but it is still doing work until it completes.

Testing Network Access to Replication Sites

To test network access to replication sites, use the CLI command “protection replica-sites test”.

```
>> protection replica-sites show

-------------------------------------------------------------------------------------------
DVX name    Address                 Type       Replica site name   Enabled        Health
----------- ----------------------- ---------- ------------------- ---------- ----------
dvx26       dvx26-mgmt.datrium.com  DST        dvx26               True               OK
----------- ----------------------- ---------- ------------------- ---------- ----------
>>
>>
>> protection replica-sites test dvx26

---------------- Origin DVX -----------------
Name        Mgmt IP      Throttle(MiB/second)
----------- ------------ --------------------
dvx24       10.2.0.234                    0.0
----------- ------------ --------------------

----------------------- Target DVX ------------------------
Name        Mgmt IP                  Throttle(MiB/second)
----------- ------------------------ --------------------
dvx26       dvx26-mgmt.datrium.com                    0.0
----------- ------------------------ --------------------
Success: True

-------------- Network status ---------------
Host           Reachable   Unreachable hosts
-------------- ----------- -----------------
192.168.47.21  True        All reachable
192.168.47.23  True        All reachable
```

You can control replication traffic by using the DVX UI to manage replication at different levels:

Protection Group Replication Traffic	Enable or disable protection group schedules.
Scheduled Replication Traffic	Add or remove replica site references to control replication for an individual schedule.
Replica Site Replication Traffic	Enable or disable replication traffic associated with a particular site.

Protection Group Replication Traffic

A protection group can generate multiple streams of scheduled replication traffic. You can define one or more snapshot schedules for a protection group. Each schedule can reference one or more replica sites. Each time a schedule generates a snapshot, it also sends a replica of the snapshot to each replica site.

Important: Do not turn off replication for an SRM-enabled protection group.

Disable protection group schedules

Stop snapshot and replication operations for all schedules defined for the group

CLI

Use the protection groups show command to obtain the protection group ID:

```
protection groups show
protection groups disable groupID
```

GUI

In the “Live groups” frame of the Protection view, select a row in the protection group list and click on the menu icon to display the live group operations. Click on “Disable schedule”.

Enable protection group schedules

Start snapshot and replication operations for all schedules defined for the group.

CLI

Use the protection groups show command to obtain the protection group ID:

```
protection groups show
protection groups enable groupID
```

GUI

In the “Live groups” frame of the Protection view, select a row in the protection group list and click on the menu icon to display the live group operations. Click Enable schedule.

Scheduled Replication Traffic

A protection group schedule can generate replication traffic to one or more replica sites. You can manage replication traffic for a single schedule by adding and removing replica sites:

> **Important:** Do not turn off replication for an SRM-enabled protection group.

Stop replication for a protection group schedule

Remove replication site references from the schedule. The schedule will still produce snapshots

CLI

Use the show commands to obtain the protection group ID and schedule name. Use the `--stop-replication` argument with the schedules edit command to stop replication for the specified schedule.

```
protection groups show
protection groups schedules show groupID
protection groups schedules edit groupID
                                --schedule-name name
                                --stop-replication
```

GUI

Edit the protection group schedules and click the Replicate snapshots check box to turn off replication for the schedule. The GUI will hide the replication fields.

Start replication for a protection group schedule

Add a replica site to the schedule. Replication to that site will start with the next scheduled snapshot. You can add multiple sites to a schedule.

CLI

Use the show commands to obtain the protection group ID and schedule name. Use the --replica-site-names argument with the schedules edit command to add one or more replica sites to the specified schedule.

```
protection groups show
protection groups schedules show groupID
protection groups schedules edit groupID
                              --schedule-name name
                              --replica-site-names name [name
...]
```

GUI

Edit the protection group schedules. Click the Replicate snapshots check box for the appropriate schedule, then specify a destination and replica retention. The default retention period is 4 hours.

Replica Site Replication Traffic

You can use the following operations to manage replication traffic at the replica site level.

- Stop outgoing replication to a specific site
- Start outgoing replication to a specific site.
- Stop incoming replication from a specific sit
- Throttle replication traffic

> **Important:** Do not turn off replication for an SRM-enabled protection group

Stop outgoing replication to a specific site

Stop future replication traffic to a replica site. All current replication traffic to the site is paused.

CLI

Use the show command to obtain the replica site name.

```
protection replica-sites show
protection replica-sites disable replica-site-name
```

GUI

In the Replica Sites frame of the Protection view, select a replica site row in the destination list and click on the menu icon to display the replica site operations. Click disable.

Start outgoing replication to a specific site.

Start replication to the site, according to the schedule. Any paused replication tasks are resumed.

CLI

Use the show command to obtain the replica site name.

```
protection replica-sites show
protection replica-sites enable replica-site-name
```

GUI

In the Replica Sites frame of the Protection view, select a replica site row in the Destination DVX list and click on the menu icon to display the replica site operations. Click enable.

Stop incoming replication from a specific sit

Site does not have permission to create snapshot replicas on the local DVX System. There are two situations for stopping incoming replication from a site.

Remove the destination replica site on the source system to stop incoming replication from a site as a part of normal operation. In this case both the source and destination are operating normally and are in communication with each other.

Permanently revoke permission for incoming replication because the local source site is no longer available or is not reachable. The revoke operation is an unusual circumstance. To reverse this operation, you must delete the revoked replica site on the source system and add it again.

CLI

Use the show command to obtain the replica site IP address or fully-qualified domain name

To remove the destination replica site:

```
protection replica-sites show
protection replica-sites remove replica-site-name
```

To revoke permission:

```
protection replica-sites show
protection replica-sites revoke
        --ipAddress replica-site-IPAddress
```

GUI

To stop incoming replication traffic, use the operations in the Replica Sites menu

If you are using the same replica site as a destination for replication and as a source of replication for a local replica group, you can remove the destination replica site. Select a row representing the outgoing replication traffic, then click on the menu icon to display the replica site operations. Click Remove.

If you are only using the replica site for incoming replication traffic, you must revoke the site. Select a row in the representing the incoming replication traffic from the site and click on the menu icon to display the replica site operations. Click Revoke and remove.

Throttle replication traffic

Reduce or increase the bandwidth for replication traffic between sites. Bandwidth is specified in mebibytes per second. The suggested setting is 100 Mbps. The minimum setting is 20 Mbps. Use a single bandwidth throttle setting, or you can set a schedule to throttle traffic. You can also turn off the throttle setting.

CLI

When you set throttling for replication traffic, you must specify either a destination site to throttle outgoing traffic or a source site address to throttle incoming traffic. You must specify all desired settings each time you execute the command.

To turn off throttling for a site, you can specify the source or destination site and omit the `--bandwidth` argument, specify zero for the bandwidth, or you can use the `protection replica-sites throttle unset` command.

```
protection replica-sites throttle set
      --dst-site-name destinationSiteName | --src-site-address
srcIPAddress
     [--bandwidth bandwidth]
     [--reduced-bandwidth bandwidth

    --reduced-start-time time

    --reduced-end-time time
    --reduced-days days [days]]
     [--use-dest-timezone]
     [--use-src-timezone]
```

- Use the `--dst-site-name` argument to specify throttling for outgoing traffic to the specified destination site. The argument value is the name of the destination site. The output of the "protection replica-sites show" command includes destination site names.
- Use the `--src-site-address` argument to specify throttling for incoming traffic from the specified source site. The source site address is the floating IP address for the management interface on the source site. The output of the protection replica-sites show command includes the IP address for remote source replication sites.
- Specify the bandwidth as a positive integer representing MiB per second. If you do not specify a bandwidth, or you specify zero, the DVX System removes throttling for the specified site.
- Use the reduction arguments to schedule throttling (bandwidth, start and end times, day(s) for reduced bandwidth). If you use any of the reduction arguments, you must specify all of the reduction arguments.
    - Time value is specified as a 24 hour clock value (hh:mm).
    - Day value is a three-character abbreviation for a day of the week, one of:
    - mon tue wed thu fri sat sun
    - Use space characters to separate day values.
- Use the time zone arguments (`--use-dest-timezone` and `--use-src-timezone`) to indicate the timezone for scheduled throttling.

GUI

In the Summary tab of the Protection view:

1. Select a source or destination site in the Replica sites frame.
2. In the frame menu, select Throttle.
3. Set the replication bandwidth in the Throttle replication dialog.
4. Click OK.

Monitoring Replication Tasks

To monitor replication tasks, use either the DVX GUI or the DVX CLI.

- DVX GUI: Select the “Tasks” tab in the “Protection” view.
- DVX CLI: Use the protection tasks show command.

When there are multiple schedules generating replicas for a protection group, the DVX System checks the status for all of the protection group schedules and displays the health status based on the combined results. If a scheduled snapshot replication fails, the DVX GUI displays a banner indicating the error. The GUI continues to display the banner until a subsequent task for that schedule successfully replicates all of the protection group snapshots. The banner duration is governed by the schedule interval, so the schedule could keep the banner up for a considerable amount of time before its next task executes. To remove the banner, use one of the following procedures:

- If the task is still available, retry the same task. In the task list, the GUI indicates the failed task by a red icon and by its Error status. To retry the task, select the task row. The GUI will enable the Retry task button. Click the button.
- Use manual replication to replicate the snapshot for which replication previously failed. In the task list, select the snapshot name in the entry for the failed task. The GUI will display the snapshot details. Click on the “Replicate” button.

Restore Virtual Machines, Files, and Protection Groups

You can use snapshots to restore virtual machines and files to the time corresponding to a snapshot.

- The restore operation preserves the location (file paths and names) of the snapshot origin.
- The DVX System sets the time of the restored elements to the time of the restore operation.

The following sections provide information about restore operations:

Restore a Virtual Machine	If a virtual machine is corrupted or you want to revert to the state of the virtual machine at an earlier time, use a local snapshot to restore the virtual machine.
Restore a Protection Group	Use a protection group snapshot to revert the state of the entire contents of the group.
Restore Guest Files	Restore one or more individual files from the guest operating system context of a virtual machine snapshot.
Disaster Recovery	If you are using replication, you can use a replica site for recovery after a source site goes down. If you are using VMware SRM, you can use the DVX System as an array-based solution to provide replication support for SRM operations.

Restore a Virtual Machine

You can restore a virtual machine to the time corresponding to a snapshot. You must power off the virtual machine before you restore it. The DVX System uses the snapshot to overwrite the existing virtual machine and it uses the same virtual machine name.

When you restore a virtual machine, the DVX System performs the following actions:

1. The DVX System creates a new snapshot of the existing virtual machine. The new snapshot has a retention period of 30 minutes. You can change the retention period of this snapshot.

   You can undo the restore transaction by restoring to this newly created snapshot.

2. The DVX System uses the specified pre-existing snapshot to restore the virtual machine to an earlier point in time. The DVX System uses the pre-existing snapshot to overwrite the virtual machine data. The restored virtual machine has a timestamp based on the time of the restore operation.
3. You can also use the restore operation to restore a virtual machine that has been deleted. The DVX System will restore the virtual machine at the site where you execute the restore operation, using the virtual machine file path from the site of the snapshot origin. In this case, you must register the virtual machine manually.

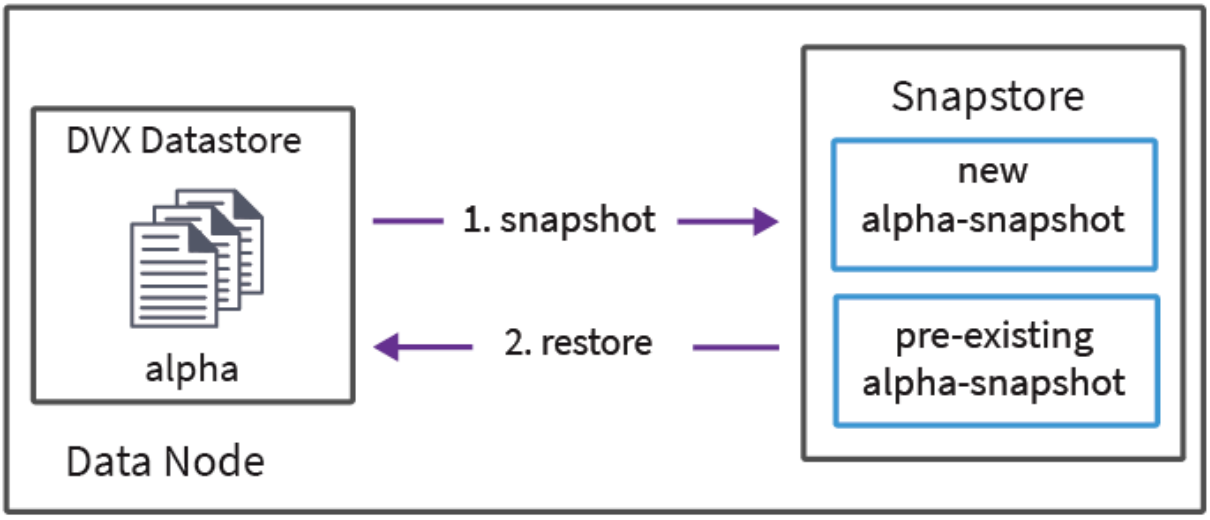

Virtual Machine Restoration – GUI

You must power off the virtual machine before you start the restore operation.

To restore the virtual machine:

1. Select the VMs view, then select Snapstore view.
2. Select the row for the virtual machine to be restored.
3. Click the Restore button. The GUI displays the Restore VM dialog.
4. Use the default snapshot selection or select another. By default, the GUI will use the latest snapshot.
5. Click the Restore button. The DVX System creates a task. You can monitor the task on the Protection tasks sidebar or on the Tasks tab of the Protection view.

Virtual Machine Restoration – CLI

You must power off the virtual machine before you start the restore operation.

```
vms restore vmID | --vm-path vmx-file-path
             --snapshot-id vm-snapshot-ID
            [--no-wait]
            [--force]
```

You must specify either a virtual machine ID or the path to the virtual machine .vmx file. Use the vms show command to display the list of virtual machines. The output includes both IDs and paths. You must also specify the snapshot ID that will be the basis for the restore operation.

- To obtain the DVX virtual machine ID, use the `vms show --in-snapstore` command.
- The .vmx file path is the DVX datastore path to the virtual machine .vmx file in the DVX datastore. For vSphere ESXi, the path includes the vmx file name and extension. The path does not include the ESXi mount point prefix (/vmfs/volumes/*dvx-datastore-label*); it begins with the slash (/) following the the DVX datastore label. For example, the DVX datastore path for "/vmfs/volumes/dvx-Datastore1/vm01/vm01.vmx" is "/vm01/vm01.vmx".
- To obtain the virtual machine snapshot ID, use the `vms snapshots show` command.
- By default, the command waits for the restore task to complete before it returns to the CLI prompt. If you use the `--no-wait` argument, the CLI will start the task, then display the task ID and return to the CLI prompt. You can use the task ID with the `protection tasks show` command to display information about the restore task.
- Use the `--force` argument to suppress confirmation prompting when overwriting an existing virtual machine.

Restore a Protection Group

You must power off any virtual machines in the protection group before you start the restore operation.

To restore a protection group:

1. Select the Protection view.
2. Select the row for the group to be restored.
3. Click Restore. The GUI displays the Restore group dialog.
4. Use the default snapshot selection or select another. By default, the GUI will use the latest snapshot.
5. Click the Restore button. The DVX System creates a task. You can monitor the task on the Protection tasks sidebar or on the Tasks tab of the Protection view.

Before it restores the protection group, the DVX System will take a snapshot of the group based on the current group membership. The retention period of this new snapshot is 30 minutes. You can change the retention period if needed.

If any of the virtual machines in the protection group had been deleted, the DVX System will restore the virtual machine(s) at the site where you execute the restore operation, using the virtual machine file path from the site of the snapshot origin. You must register these virtual machines manually.

Restore Guest Files

The DVX System supports the ability to retrieve one or more files contained within the guest file system of a DVX virtual machine snapshot.

- You can download a single file directly to your computer.
- You can place one or more files or folders in an ISO file. You can then mount the ISO file in the guest, and access the recovered file(s).

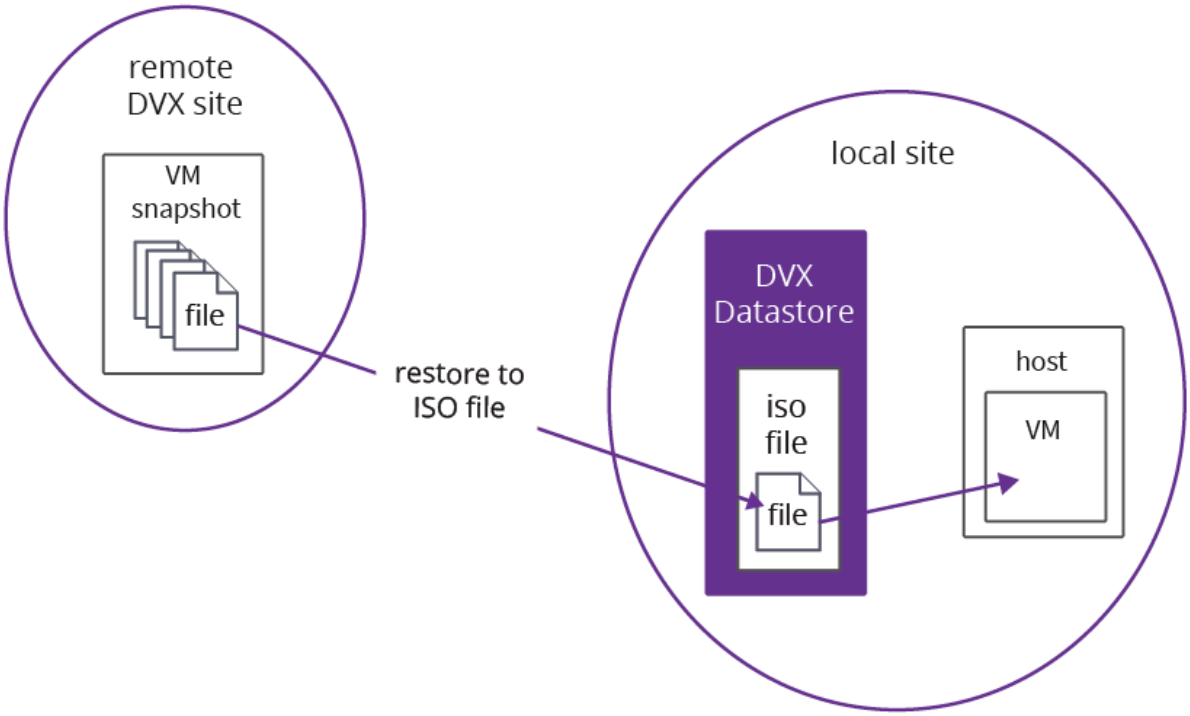

- The virtual machine snapshot can be on your local site, on a remote on-premises site, or in a Cloud DVX site. The figure above shows the use of an ISO file to recover a guest file from a remote site.

> **Note:** The DVX System supports the following guest file systems for browsing and file retrieval:

- NTFS
- FAT32
- EXT3, EXT4

If you have FAT32 partitions in your guest file system, it can take extra time for the DVX System to read these partitions. Guest file recovery supports a maximum of 4 FAT32 partitions in a guest file system.

Guest file recovery size limits:

File	Maximum size
ISO file.	10GB
File inside an ISO file.	4GB - 1 byte
Single file for download.	2GB
Virtual disk with VMware snapshot.	2TB
Paths (folders or files) which can be added to ISO.	25
Guest volumes.	100
Entries in a single guest directory.	10,000
Files that can be restored in one ISO file.	5,000

Further Guest file caveats and restrictions:

- The DVX System does not support guest file recovery operations on dynamic volumes in Windows guest file systems or logical volumes in Linux guest file systems.
- The entire path of the guest file is preserved within an ISO file unless the path meets any of the following conditions:
- Longer than 156 characters.
- More than 7 levels deep.
- Contains non-ASCII characters.
- In these cases, the DVX System creates the file using a different name under root within the ISO file. A separate README file is created to indicate the actual full path of the file.
- If you create a partition name, do not use slash characters in the name. The DVX guest file recovery capability does not support slash characters in partition names.
- The DVX System attempts to determine Windows drive letters and displays drive letters when browsing files. There might be situations where the DVX System displays an incorrect drive letter.
- The DVX System does not support guest file recovery of a Windows encrypted file.

- The DVX GUI shows BitLocker-encrypted partitions as unsupported. You cannot navigate unsupported links, or perform guest file recovery operations on these partitions.
- You can browse NTFS encrypted folders, but the DVX System does not support guest file recovery of NTFS encrypted files.
- The DVX System shows an empty disk in a guest file system as an unsupported volume.
- Files inside a recovery ISO file have the same size and modification time as the original files. File permissions are not persisted. The files are set to read access only. You must have permission to read the drive to browse the files. To use the files, copy them to the desired location.

Restore Guest Files – GUI

The DVX GUI provides the Restore Guest Files capability within the context of a virtual machine or a protection group. In either case, you browse the guest file system of a virtual machine snapshot, and then download one or more files either directly or use the files to create an ISO file on your local datastore.

You can download a single file from your local system, an on-premises remote system, and from a Cloud DVX. You can also create an ISO file on your local datastore using guest files from your local system, an on-premises remote system, or from a Cloud DVX. The following sections provide some examples of guest file recovery.

Snapshot Selection

Guest file recovery requires a virtual machine snapshot. The snapshot can be on your local DVX site or it can be on a remote site (on-premises or Cloud DVX). If the snapshot is on both the local and remote sites, the DVX System will use the local snapshot, even if you select the remote snapshot.

- **Start the recovery process from a virtual machine on your local system**. When you work from the VMs view, select a local virtual machine to start guest file browsing. By default, the DVX System uses the most recent snapshot in the local SnapStore.

- **Start the recovery process from a protection group**. When you work from a protection group, first you select a protection group snapshot, then a virtual machine snapshot from the group snapshot to start guest file browsing. If you have replicated the protection group to a remote site, you have the choice of browsing local or remote snapshots, otherwise you can browse only local snapshots.

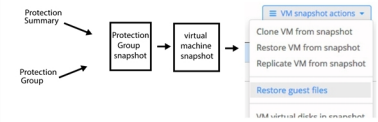

- The Restore guest files dialog displays the guest file content of the virtual machine snapshot. The dialog also gives you the option of selecting a different snapshot, including any snapshots that are available on remote sites.

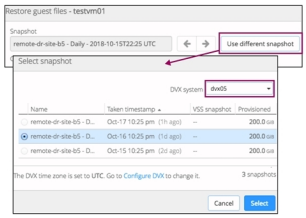

- When you select a snapshot, the DVX GUI displays the root level content of the guest file system.

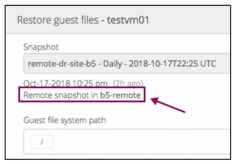

The remote snapshot selection is indicated in the Restore guest files dialog.

- For an example of guest file recovery based on a virtual machine, see Download a File from a Local VM snapshot.
- For an example of guest file recovery based on a protection group, see Recovery from a Remote Protection Group Snapshot (ISO File).

Browsing Guest Files

The Restore guest files dialog indicates the snapshot context and the files and folders that are available in the virtual machine guest file system(s). The dialog provides a selection mechanism to browse the file system.

The “Guest file system path” field shown in the image below is a dynamic text field, the contents of which are updated as you select from the list of available files and folders shown beneath it. The dialog starts at the root comprising the set of all volumes in the virtual machine.

- When you select a file name entry, the DVX GUI appends the selection to the guest file system path and displays the corresponding list of files and folders.
- The elements in the path are active. Click on an item to navigate to that level of the file system.

You cannot navigate symbolic link references in the guest file system. The DVX GUI displays them as inactive links.

This image shows the progression from the root directory to the C: drive to the “Users” folder.

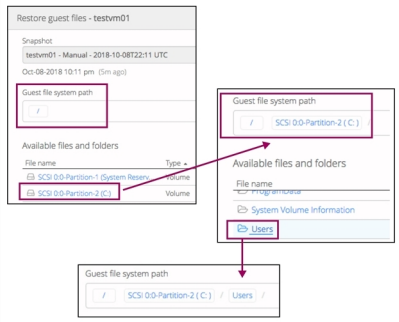

The DVX GUI provides the simple navigation paradigm of selecting a folder to extend the path. You can also edit the path field manually.

- To edit the path, click on blank space in the path text field.
- Press the Return key to return to the active element path representation.

The path specification starts with the virtual disk and a volume on that virtual disk. The path syntax is required for both GUI and CLI.

- The path begins with a forward slash ('/').
- The path must contain a space between the device type and the bus number.

The partition number prefix ("-Partition-") is required.

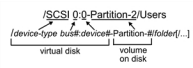

Download a File from a Local VM snapshot

To download a file from a local VM snapshot:

1. Click on the VMs view.
2. Click the Snapstore view.
3. Select a virtual machine. (Either selecting the toggle for the VM, or clicking on the VM name to navigate to the virtual machine page will provide the "Restore Guest Files" capability).
4. Click Restore guest files.

The following image shows the extended path – /SCSI0:0 Partition 2 (C:)/Users – and the available files and folders in the Users directory. The DVX GUI displays a downward-pointing arrow besides individual files that can be downloaded.

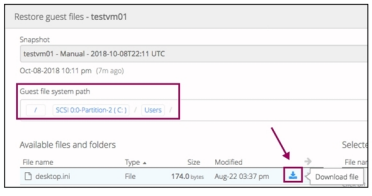

Depending on the browser you are using, when you click on the download icon, the DVX GUI might display a confirmation dialog. To download the file, click on “Save File”. The DVX System creates a copy of the file in the download location established for your browser.

A single file download operation is shown as a “Preparing file for download” operation in the Protection tasks list.

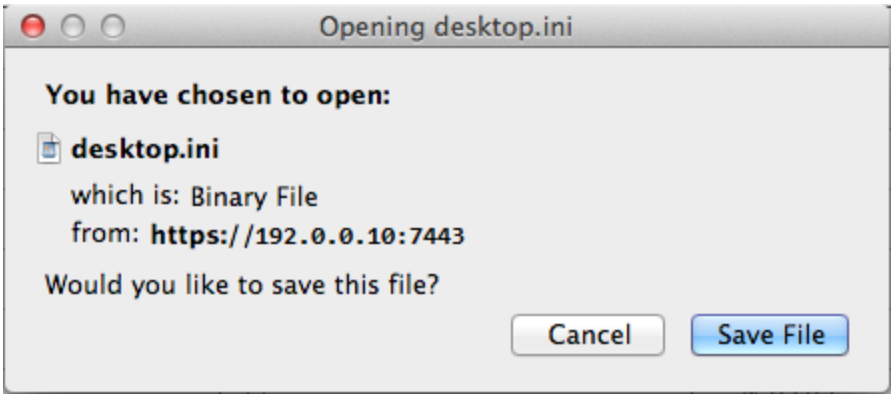

Recovery from a Remote Protection Group Snapshot (ISO File)

The following steps show the navigation sequence to display the “Restore Guest Files” dialog from a protection group snapshot.

1. Click on the Protection view.
2. Click on a protection group that replicates to your remote site. The figure below shows a protection group that replicates to the remote site “dvx05”.

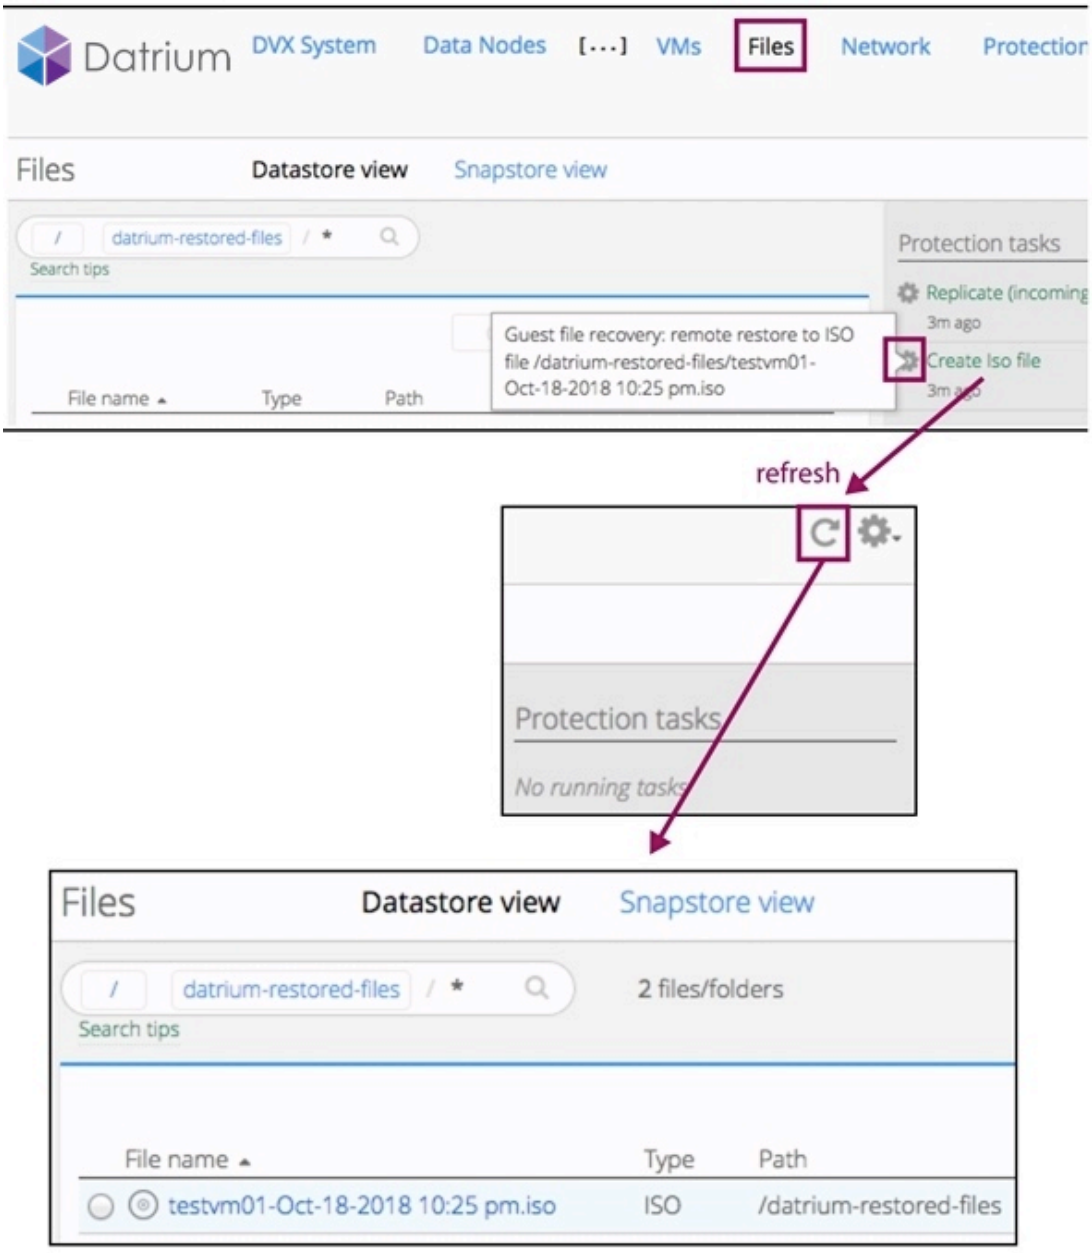

3. In the “Snapshots” frame of the protection group page, click on the checkbox for “Show remote snapshots”.
4. Select a remote snapshot. The figure below shows a situation in which the latest snapshot exists only on the remote site (dvx05). The DVX System will use this remote snapshot; if your local system has a copy of this snapshot, it will use the local copy.

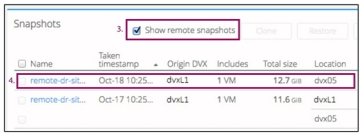

5. The snapshot page indicates that you are using the remote snapshot. In the Virtual machines section of the snapshot page, select a virtual machine, click on the VMsnapshot actions menu, and click Restore guest files”

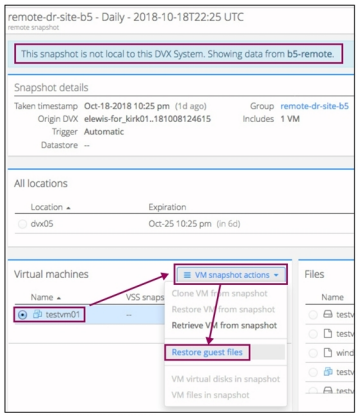

6. The Restore guest files dialog indicates the selected snapshot. This example uses a remote snapshot. The GUI displays a right-pointing arrow for files and folders that you can add to the ISO file. When you click on the arrow, the DVX System adds the selection to the ISO bundle. The maximum number of elements (folders or files) that you can add to an ISO bundle is 25. To continue, click on “Create ISO”.

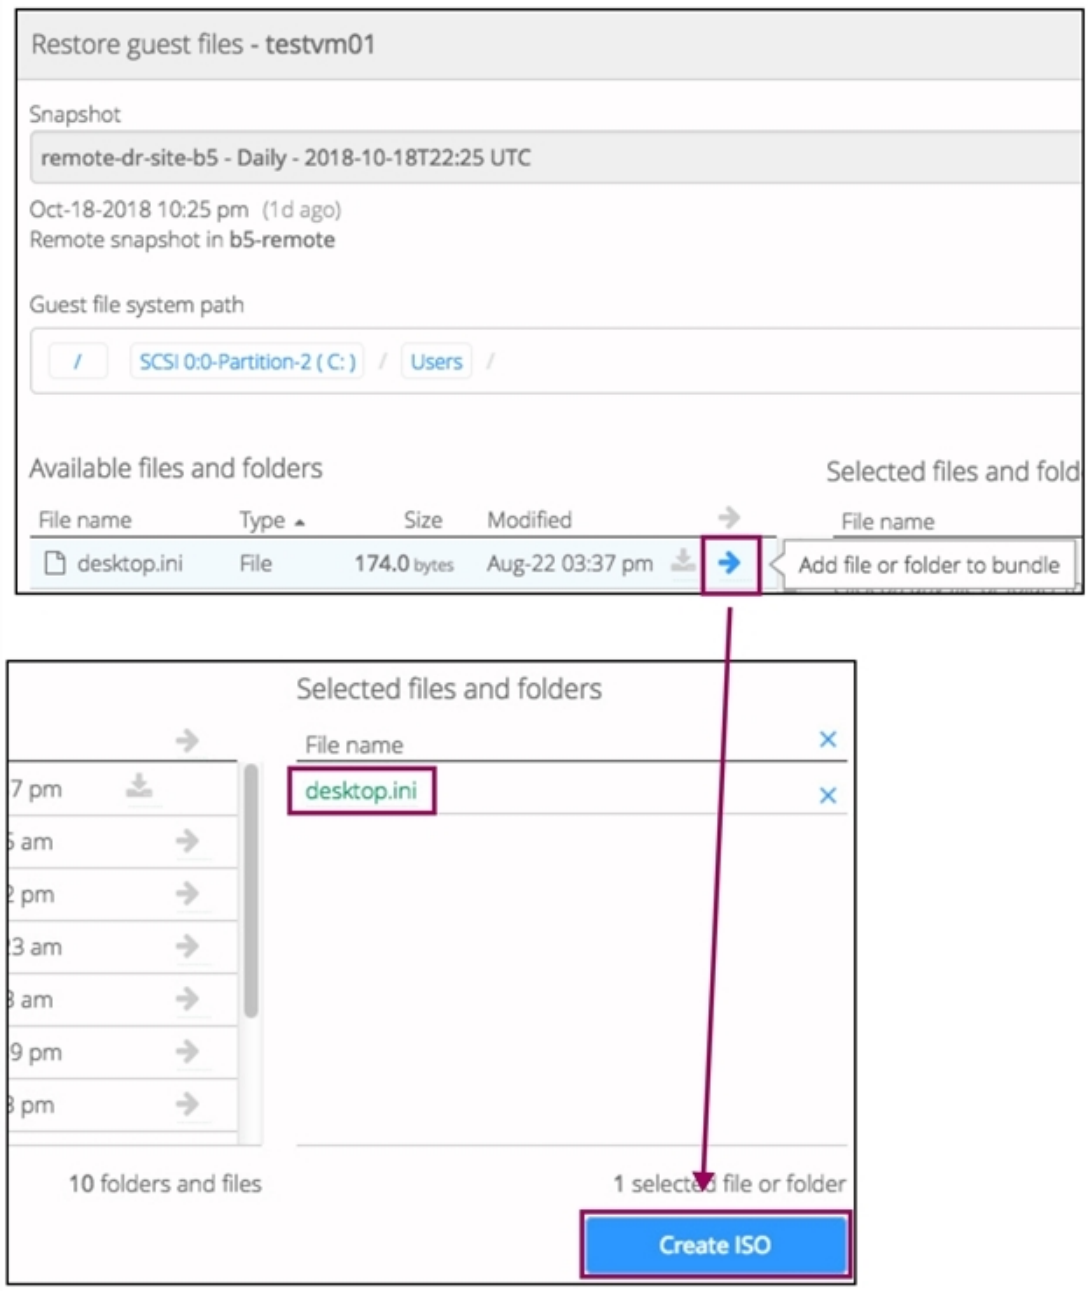

7. The Create ISO dialog indicates the ISO file name, location, and expiration. If your DVX System has more than one datastore, you will have the option of selecting a datastore. The suggested location is the directory “/datrium-restored-files”. You can change the file name and location. The DVX System will create the directory path in the ISO file name if it does not exist. To continue, click on “Create ISO”.

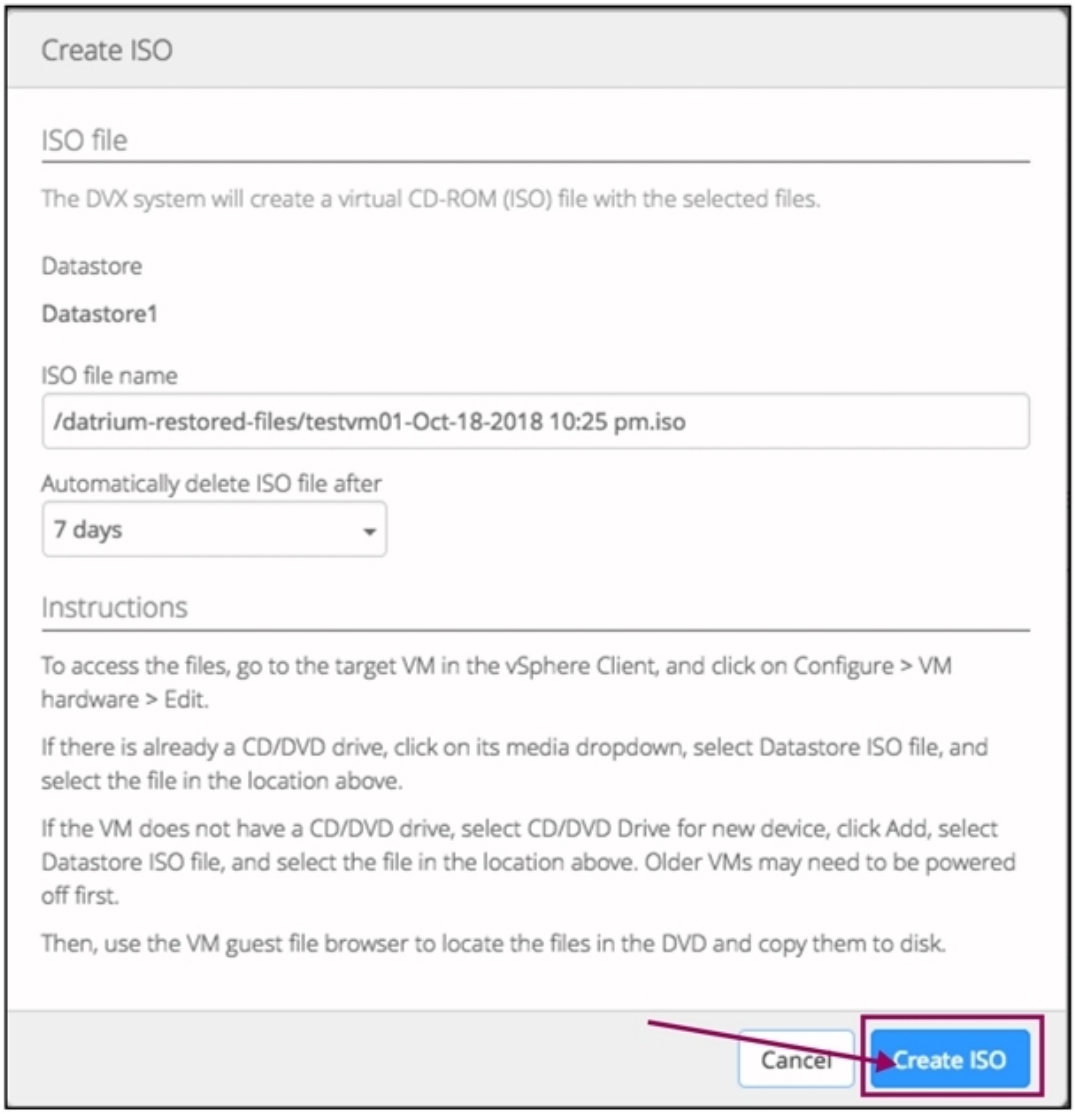

8. The “Restore guest files” dialog shows the state of the ISO file download.

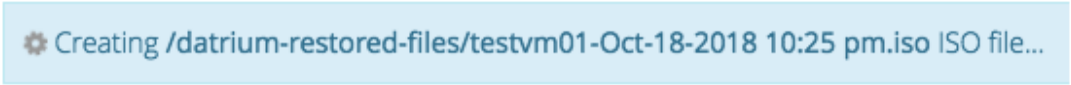

Use the “Files” view to monitor the ISO file download. The “Files” view provides a DVX datastore browser and it shows the current protection tasks. The GUI shows a rotating gear by the task as long as it is executing. After the task has completed, refresh the display to see the ISO file in the DVX datastore.

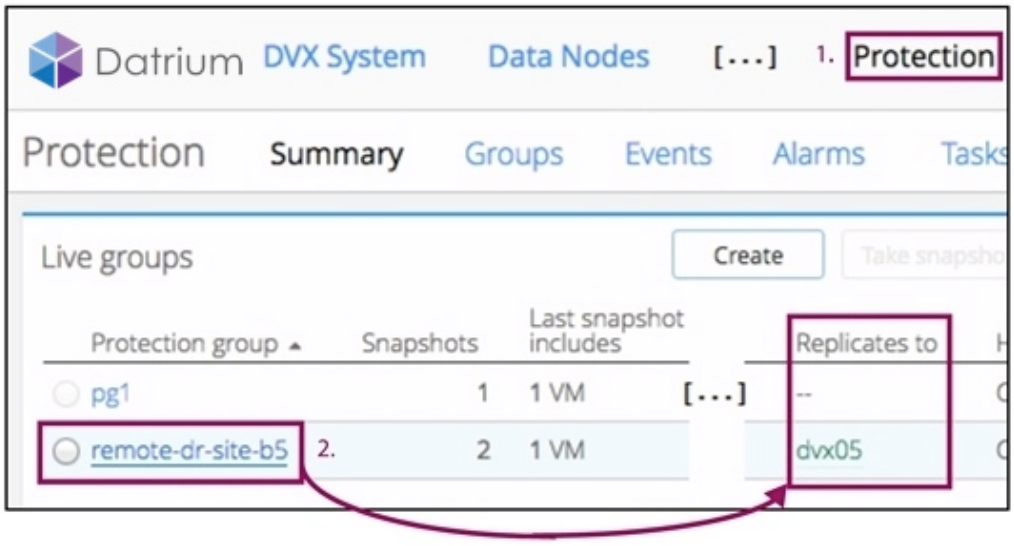

Restore Guest Files – CLI

Use a virtual machine snapshot to restore a file to a virtual machine. You can restore a file from a snapshot on your local DVX System or you can restore from a snapshot on a remote DVX System (on-premises or DVX Cloud).

To restore a file:

1. Identify a virtual machine snapshot. The snapshot can be on your local DVX System or on a remote DVX System.
2. Search for a file contained in the guest file system in the snapshot.
3. Retrieve the file by downloading it to your local desktop or by restoring it to an ISO file.

Recovery from a Local Snapshot – CLI

This example demonstrates how to use a local snapshot to recover a file.

1. Display the virtual machine snapshots in the local DVX Snapstore. The output data includes the virtual machine ID.

   ```
   vms show --in-snapstore
   ```

2. Display the virtual machine snapshots. The output data includes the virtual machine snapshot ID.

   ```
   vms snapshots show --vm-id
   ```

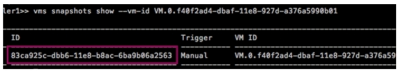

3. Display the guest files in the virtual machine snapshot.

   ```
   vms guest-files show --vm-snapshot-id --path
   ```

   The image below shows the sequence of `vms guest-files show` commands to traverse the directory tree. This example shows navigation to the “Users” directory, containing the desktop.ini file to be downloaded.

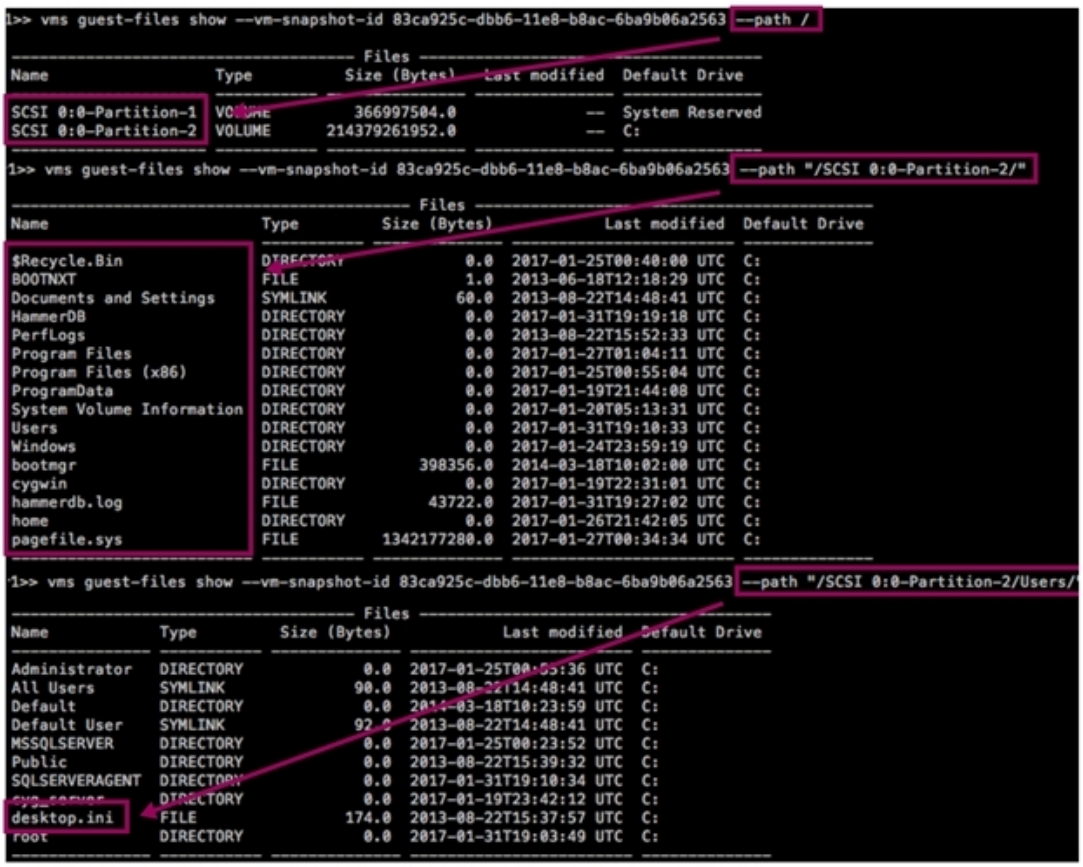

The path specification starts with the virtual disk and a volume on that virtual disk.

- The path begins with a forward slash (‘/’).
- The path must contain a space between the device type and the bus number.
- The partition number prefix (“-Partition-”) is required.

4. Create an ISO file containing the desired file(s). The `vms guest-files restore-to-iso` command requires the virtual machine snapshot ID, guest file path, and ISO file information.

```
vms guest-files restore-to-iso --vm-snapshot-id
                                    --paths
                                    --iso-path
                                    --iso-datastore
```

- The `restore-to-iso` command uses the same values for snapshot ID and guest file path that were used with the vms guest-files show command that identified the desired file.
- The `--iso-path` argument identifies the ISO file location in the local DVX datastore. The DVX System will create the directory path in the ISO path specification if it does not exist.
- The `--iso-datastore` argument identifies the local DVX datastore.
- The `restore-to-iso` command starts an asynchronous task to create the ISO file. This example uses the suggested directory “/datrium-restored-files” in the ISO path. The command line in the figure below has been formatted for readability. This image shows the `restore-to-iso` command, followed by a `files show` command to display the resulting ISO file.

Recovery from a Remote Snapshot – CLI

In the following example, there is no local copy of the desired virtual machine snapshot for guest file recovery. There is a copy of the snapshot on a remote site. To retrieve the file, you need the following elements:

- The protection group that replicates to the remote site. You will use a protection group snapshot on the remote site.
- The virtual machine snapshot that contains the desired file for recovery. When you do not have a local instance of the virtual machine, you must be able to identify the virtual machine from the protection group membership.

The following example recovers a desktop.ini file from the virtual machine snapshot.

To use a snapshot on a remote site:

1. Display the protection group that contains the desired virtual machine and which replicates to the remote site. The output data includes the protection group ID.

   ```
   protection groups show
   ```

   

2. Display the protection group snapshots that are on the remote site. The command output includes the protection group snapshot ID.

   ```
   protection groups snapshots show --site-names
   ```

   

3. Find the virtual machine ID.

   ```
   protection groups members show
   vms show --vm-name-pattern
   ```

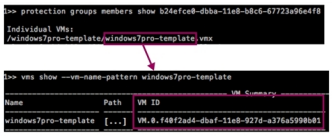

4. Display the virtual machine snapshots in the protection group snapshot on the remote site. The command output includes the virtual machine snapshot ID.

   ```
   ms snapshots show --protection-group-snapshot-id
                                        --site-name
   ```

5. Display the guest files in the virtual machine snapshot.

   ```
   vms guest-files show --vm-snapshot-id --site-name --
   path
   ```

   The following figure shows the sequence of guest-files-show commands that navigate the directory tree. The first command starts at the root level ("--path /").

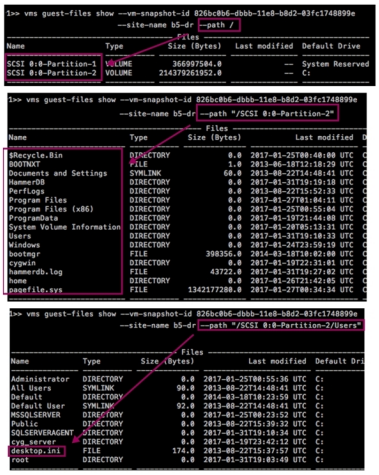

   The path specification starts with the virtual disk and a volume on that virtual disk.

   - The path begins with a forward slash ('/').
   - The path must contain a space between the device type and the bus number.
   - The partition number prefix ("-Partition-") is required.

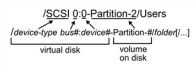

6. Restore the guest file to an ISO file. This example uses the virtual machine snapshot ID. As an alternative, you can use the combination of protection group snapshot ID and virtual machine ID.

   ```
   vms guest-files restore-to-iso --vm-snapshot-id
                                                  --site-
   name
                                                  --paths
                                                  --iso-path
                                                  --iso-data-
   store
   ```

   The snapshot, site name, and paths arguments apply to the remote site. The ISO arguments identify DVX datastore locations on the local site.

   - The `restore-to-iso` command uses the same values for snapshot ID and guest file path that were used with the vms guest-files show" command that identified the desired file.
   - The `--iso-path` argument identifies the ISO file location in the local DVX datastore. The DVX System will create the directory path in the ISO path specification if it does not exist.
   - The `--iso-datastore` argument identifies the local DVX datastore.
   - The `restore-to-iso command` starts an asynchronous task to create the ISO file. The command output includes the task ID. Use this task ID to monitor the task. This example uses the suggested directory "/datrium-restored-files" in the ISO path. The command line in the figure below has been formatted for readability.

7. Monitor the recovery task.

   ```
   protection tasks show
   ```

8. Use the `protection tasks show` command to monitor the recovery task. Do not use the ISO file until the “State” value is “SUCCESS”.

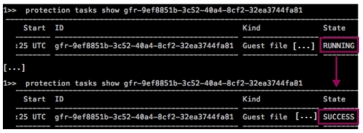

9. Display the ISO file location. This example shows the files in the “datrium-restored-files” directory.

   `files show --file-path`

```
1>> files show --file-path /datrium-restored-files/*

Details:
------------------------------------ Summary ------------------------------------ ----------
File                                                    Folder  Size (GiB) Datastore | Snapshots
------------------------------------------------------- ------- ---------- ----------|----------
/datrium-restored-files/win-7-desktop-ini-from-b5-dr         0        ~0.0 Datastore1|         0
------------------------------------------------------- ------- ---------- ----------|----------
```

Non-Ascii Characters in Recovery File Names

The DVX System does not support non-ascii characters in guest file names. If you recover a file with a name that includes non-ascii characters, or with a path that is too long (more than 156 characters) or too deep in the directory tree (more than 7 levels deep), the DVX System creates the file using a different name under root. A separate README file is created to indicate the actual full path of the file.

The recovered file has a name using the following format:

renamed_file_*nnnnnn.ext*

The file name starts with the prefix “renamed_file_”, followed by a 6 digit number. The file extension is the same extension as the source file.

The corresponding readme filename has the following format:

renamed_file_*nnnnnn.ext*.README.TXT

The readme file name contains the same “renamed_file...” file name. It contains the original file name.

Disaster Recovery

The DVX System supports the following methods of disaster recovery:

- ControlShift
- DVX Integration with vCenter SRM
- DVX Disaster Recovery Example

DVX Integration with vCenter SRM

Datrium provides a Storage Replication Adapter (SRA) for integration with VMware Site Recovery Manager (SRM). SRM uses the storage and replication capabilities of the DVX System to automate disaster recovery operations in the vSphere environment.

Datrium provides SRA Agent software that runs on VMware SRM servers. The Datrium SRA Agent supports SRM operations in a DVX on-premise environment. The DVX System works with SRM that is configured to use array-based replication.

An SRM protection group corresponds to one DVX protection group. The DVX protection group is created explicitly for SRM use and it is associated with a single DVX datastore.

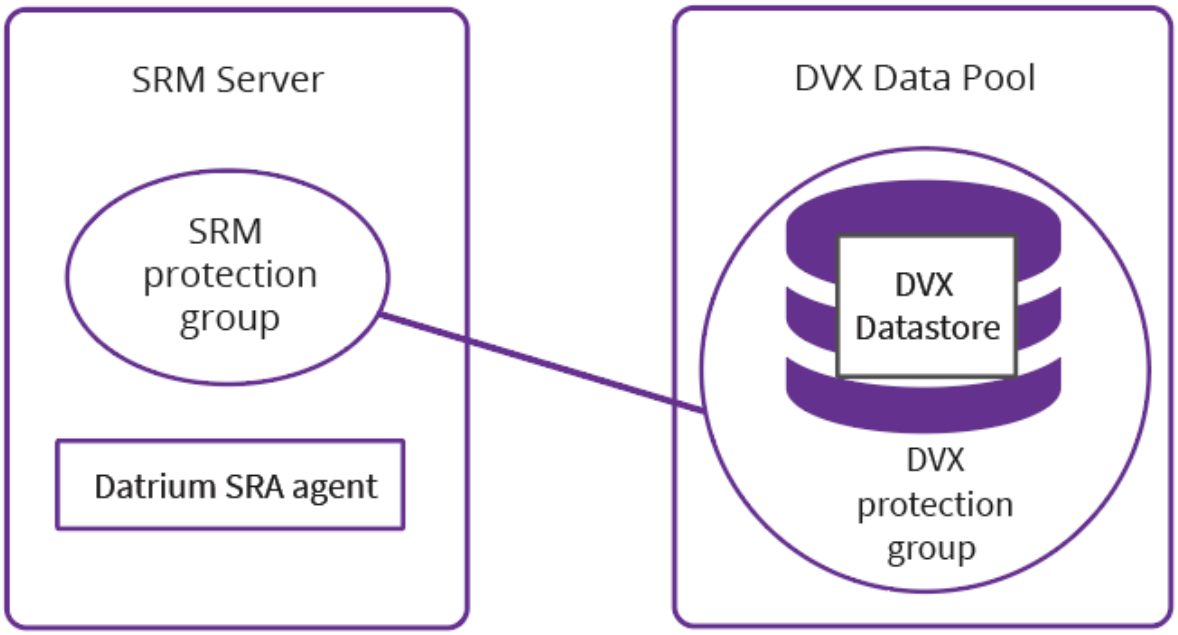

Important: After you have created the association between SRM and DVX protection groups, do not perform any DVX restore operations using virtual machine snapshots or file snapshots that are derived from SRM protection group content.

Site Configuration for SRM Operations with DVX Storage

To set up SRM sites to use DVX storage and replication, you will perform the following operations:

- Install the Datrium SRA Agent on the SRM Servers.
- Create DVX datastores for virtual machine storage and for SRM use.
- Create DVX replica site definitions.
- Create DVX and SRM protection groups, and DVX snapshot/replication schedules.
- Add an SRM array manager for the DVX Systems.
- Configure placeholder datastores and set up the SRM mappings.

Prerequisites

1. Before you set up DVX Systems for SRM operations, you must have the SRM protected site and recovery site already paired. See VMware documentation.

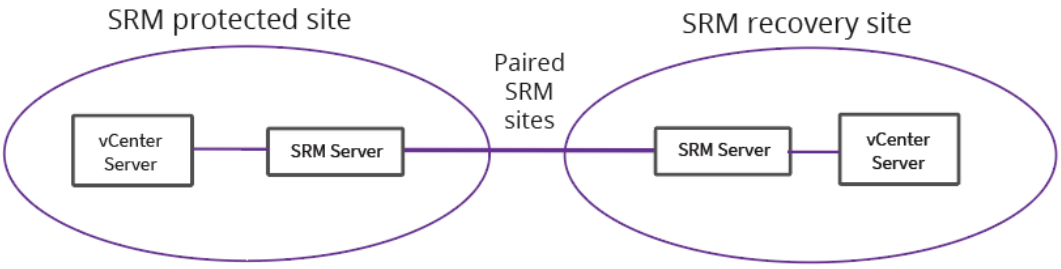

2. If you have not already done so, install DVX Systems on both protected and recovery sites. On each site:

   - Install the DVX System, including registering the DVX System with the vCenter Server. See the *Datrium DVX Data Node Hardware Installation* and *Datrium DVX Software Configuration* manuals.
   - Add hosts to the DVX System; these are the hosts on which the protected virtual machines will run. See Adding a Host to a DVX System.
   - Add the hosts to the vCenter Server. See VMware documentation.

3. If you have existing DVX Systems and you are creating a larger Data Pool, add the additional Data Node(s) to the Data Pool. See Expanding the Data Pool.

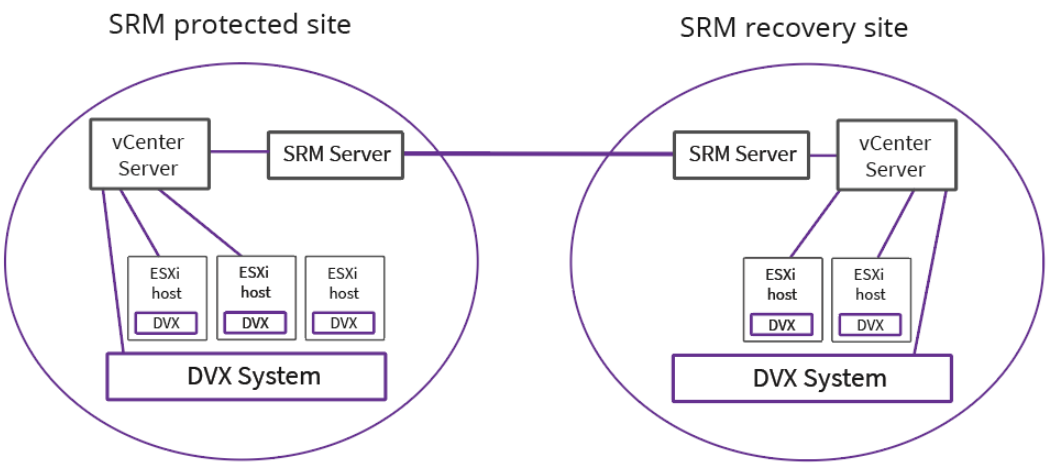

4. Add replica site definitions to each DVX System to create replica pairs. Snapshot/replication schedules will refer to these site definitions. See Replica Site Definitions.

    - On the SRM protected site, the DVX replica site definition must use the SRM recovery site DVX System as the destination.
    - On the SRM recovery site, the DVX replica site definition must use the SRM protection site DVX System as the destination.

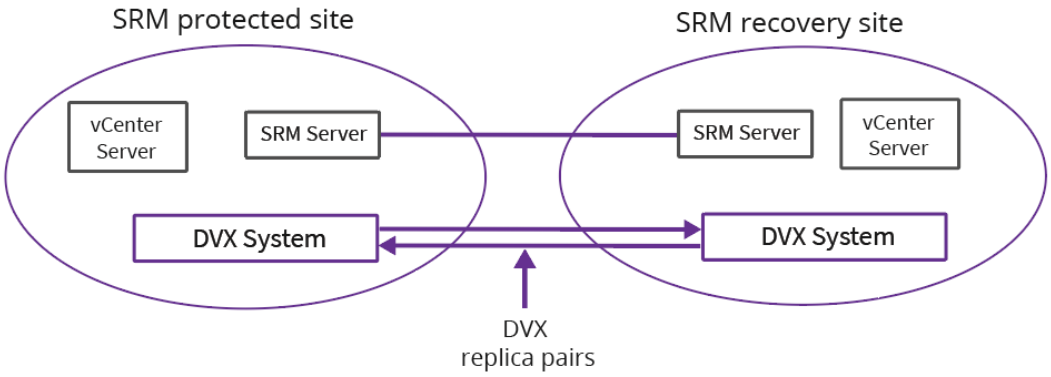

DVX/SRM Setup

1. Install the Datrium SRA Agent on the protected and recovery SRM Servers.

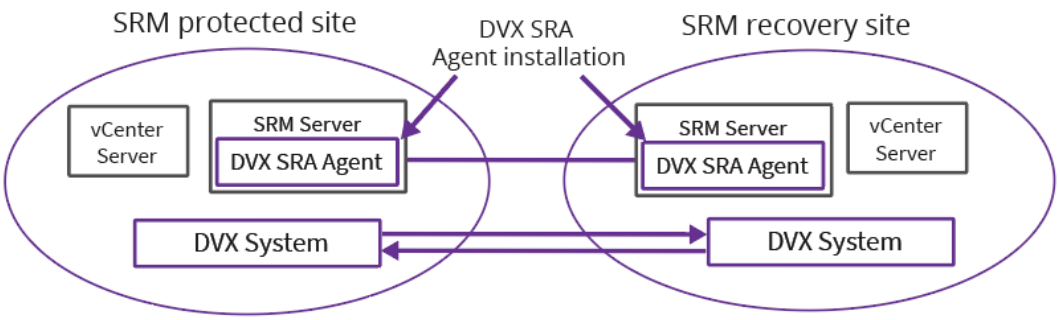

2. Download the Datrium SRA Agent MSI file.
3. Log in to the Windows virtual machine on which the SRM Server runs.

   Use a browser to connect to the management interface of the DVX System. The virtual machine must have network access to the subnet that you are using for the management interface.

   The login page contains the active download link at the bottom – “SRM Agent”. Click on the link to download the installer to your virtual machine.

   When you run the installer, it will download the Agent from the Data Node and install it on your virtual machine.

4. Double-click on the file to run the Windows installer. The installer presents a short sequence of dialogs. You have the opportunity to change the location of the Agent, if desired.
5. After finishing the setup, restart the SRM Server.

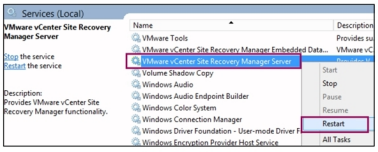

6. Log out of the vCenter Servers and then log back in again.
7. Create DVX datastores on both protected and recovery sites. (See Creating a Datastore.)
   On each site, create two datastores:

- 1 datastore for VM storage.
- 1 placeholder datastore for SRM use (do not use this datastore for anything else).

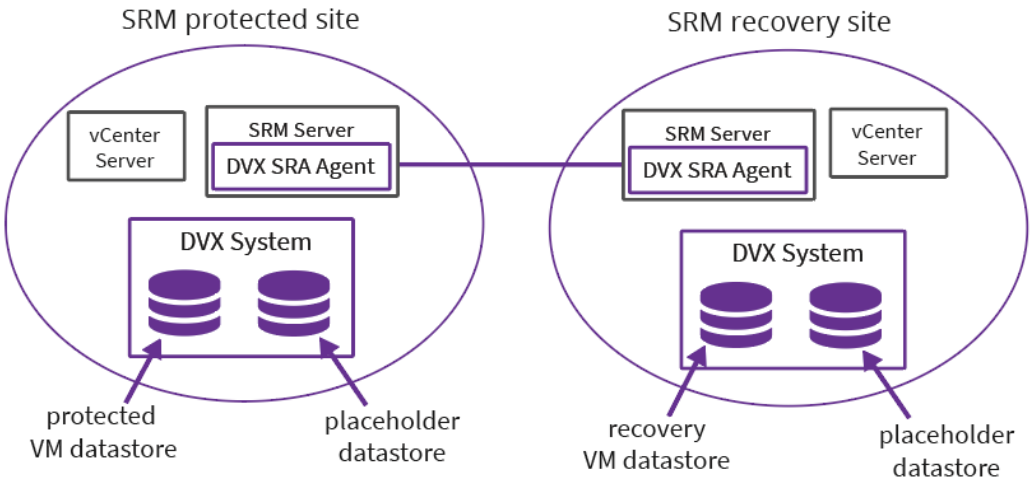

Notes:

- Virtual machines that are to be protected must be stored in the DVX System on the protected site. You can create virtual machines to use the DVX protected VM datastore, and if you are using other storage for virtual machines, use storage vmotion to transfer them to the DVX datastore.
- The protected virtual machines must be in synchronous time with the vCenter Server. Use NTP to support synchronization.
- You should plan on using the recovery datastore paired with the protected datastore for the life of this SRM protection group pairing. SRM will use the recovery datastore for failover to recover protected virtual machines. If you change the recovery datastore, you might encounter errors.

8. Add an SRM array manager for the protected and recovery site DVX Systems.

> **Note:** This procedure requires the password of the DVX System account “vmwaresrmuser” on the DVX System on both sites. The default password is “datrium#1”. To use a different password, log in to the DVX Systems and

> use the CLI command `config password` to change the password before you add the array manager. For security, we recommend that you create a password that contains at least 10 characters, including a combination of uppercase and lowercase letters, numbers, and special characters.

`config password --user-name vmwaresrmuser --password` *`new-password`*

This is an internal account that exists only to support DVX/SRM operations. You cannot display information about the vmwaresrmuser account.

9. To add an array manager, you will perform the following steps:

   - Identify vCenter Servers on the protected and recovery sites.
   - Select the Datrium Storage Adapter.
   - For each site: provide the DNS name or IP address of the floating IP management interface of the DVX System and enter the password for the DVX System account “vmwaresrmuser”.
   - Enable the array pairs and complete the process.

   a. Start the procedure by selecting “Add a pair of array managers” in the “Options” screen of the “Add Array Manager” dialog, then click Next.
   b. Select the paired SRM sites, then click Next.
   c. Select the SRA type “Datrium Storage Adapter”, then click Next.
   d. Enter a display name and credentials for the protected array manager. Specify the DNS name or IP address of the floating IP management interface on the Data Node.
   e. Enter the password for the “vmwaresrmuser” account. Click Next.
   f. Enter a display name and the credentials for the recovery array manager.
   g. Specify the DNS name or IP address of the floating IP management interface on the Data Node.
   h. Enter the password for the “vmwaresrmuser” account.
   i. Click Next.
   j. Select the array pair for enablement.

k. Click Next.

l. Verify the result and click Finish.

10. Configure mappings and the recovery placeholder datastore to support failover and recovery operations on the recovery site. The following is a brief overview of the mapping and placeholder datastore tasks. For information about these SRM tasks, see VMware documentation.

    As part of SRM site configuration, you must map system components (network, folder, and resource) on the protected site to the corresponding components on the recovery site.

    - Network mapping – Identify the sites, select test networks, and prepare reverse mappings.
    - Folder mapping – Prepare mappings and reverse mappings.
    - Resource mapping – vCenter and host mapping.

    You must also configure the placeholder datastore on both the protected and recovery sites. SRM uses the placeholder datastores for planned migration and reprotection. The following figure shows the selection of the recovery placeholder datastore.

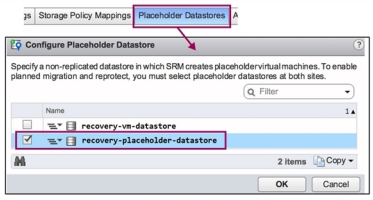

11. Log in to each host on both protected and recovery sites to mount the datastores (virtual machine storage and placeholder). Use the DVX GUI to mount the datastores. The DVX System does not support the use of the esxcli to mount DVX datastores.

    In the Datrium DVX tab, click on the DVX menu icon (gear icon) and select “Mount/unmount datastores”. See Mounting/Unmounting a DVX Datastore.

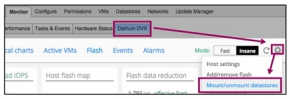

12. Create an SRM-enabled DVX protection group on the protected site.

    - This protection group is associated with a single datastore for virtual machine storage on the protected site. This protection group also identifies the SRM recovery site; you must select the destination DVX System and the recovery datastore.
    - You cannot manually promote or demote an SRM-enabled protection group.
    - You cannot turn off replication for an SRM-enabled protection group.
    - You cannot disable or delete the last snapshot schedule associated with an SRM-enabled protection group.
    - If an SRM-enabled protection group on the protected site is deleted, you will not be able to restore the protection group by replication from the recovery site.

    a. Enter the SRM protection group name and select the protected datastore to be associated with the protection group.
    b. Select the recovery site DVX System and the corresponding recovery datastore.
    c. You have the opportunity to choose VSS-enabled virtual machines in the protected datastore for app-consistent snapshots. Otherwise, the DVX System will take crash-consistent snapshots of the VSS-enabled virtual machines.
    d. Create a schedule for snapshots and replication to the recovery site. This schedule should replicate within a reasonable time frame, so there will be snapshots on the recovery site to support test, migration, or recovery operations. See Scheduled Snapshots.
    e. Create the SRM-enabled protection group.

    As an alternative, you can use the DVX CLI to create an SRM-enabled DVX protection group and snapshot/replication schedule. This schedule should replicate within a reasonable time frame, so there will be snapshots on the recovery site to support test, migration, or recovery operations. For information about commands described below, see Creating a Protection Group – CLI and Snapshots.

```
protection groups create groupName --srm-group
                                   --datastore pro-
tectedVMDatastore
```

- The `--datastore` argument identifies the protected site DVX datastore for virtual machine storage.

(Optional) If you have VSS-enabled virtual machines for app-consistent snapshots, you can add members to the SRM-enabled protection group to identify those virtual machines. VSS-enabled virtual machines are the only type of member that you can add to an SRM-enabled protection group. When the DVX System takes snapshots of the SRM-enabled protection group, it will take app-consistent snapshots of the VSS-enabled group members. Otherwise, the DVX System will take crash-consistent snapshots of those virtual machines.

```
protection groups show
protection groups members add groupID
                              --app-consistent-vm-path
vss-vm-path [vss-vm-path ...]
```

- Use the `show` commands to obtain the protection group ID.
- The `--app-consistent-vm-path` argument identifies a single virtual machine by file path. A virtual machine path expression is the DVX datastore path to the virtual machine configuration file. For vSphere ESXi, the path includes the vmx file name and extension. The path does not include the ESXi mount point prefix (/vmfs/volumes/*dvx-datastore-label*); it begins with the slash (/) following the the DVX datastore label. For example, the DVX datastore path for "/vmfs/volumes/dvx-Datastore1/vm01/vm01.vmx" is "/vm01/vm01.vmx".

13. Create a schedule for snapshots and replication to the recovery site.

```
protection groups show
protection groups replica-sites show
protection groups schedules add groupID
                  --schedule-name name
                  --schedule cronExpression
                  --retention seconds
```

```
                    --replica-site-names recoveryDVXname
                    --replica-retention seconds
                    --srm-restore-datastore-name dstore-
name
```

- Use the `show` commands to obtain the protection group ID and the recovery replica site name.
- Use the `--replica-site-names` argument with the schedules add command to specify the replica site name for the recovery DVX System. You can specify only one recovery replica site for a particular schedule.
- The `--srm-restore-datastore-name` datastore name identifies the DVX datastore on the recovery site where virtual machines will be restored during an SRM failover operation.

Datrium recommends that the initial recovery datastore persists as the recovery datastore for the life of this protection group pairing. SRM will use the recovery datastore for failover to recover protected virtual machines. If you change the recovery datastore, you might encounter errors. The following sequence demonstrates this.

a. Failover to the recovery site.
b. Failback to the protected site.
c. Change the recovery datastore associated with the protected site protection group.
d. Failover to the recovery site.

During the second failover, with the attempt to use the second recovery datastore, the DVX System will fail to restore the virtual machines on the recovery site. This is because copies of those virtual machines still exist in the original datastore even though those virtual machines have been removed from the vCenter registry on the recovery site. The DVX System does not allow more than one instance of a virtual machine. You will see the following SRM errors:

Change recovery site storage to writable.
Failed to promote replica devices. Failed to promote replica devices *deviceID*. SRA command 'failover' failed for device *deviceID*. Failed to execute failover command DVX errCode=5900.

At this point, the recovery protection group is live, and you cannot change the recovery datastore setting back to the original recovery datastore.

To resolve this situation, delete the virtual machines from the original recovery datastore. After this, the SRM failover to the recovery site can succeed.

14. Perform an SRM array scan to detect the DVX protection group.

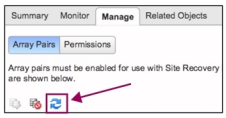

15. Create an SRM protection group by entering a name for the SRM protection group and select the site pair.

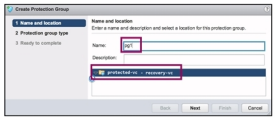

16. Select protection in the direction of the recovery site, then select the “Datastore groups” protection group type.

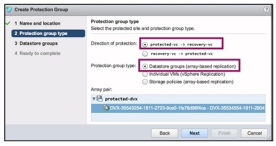

17. Select the DVX datastore group. The SRM datastore group name is the DVX protection group name.

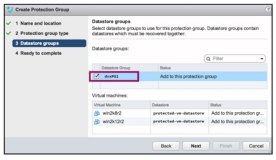

18. After you have finished the procedure, you can create an SRM recovery plan for this protection group.

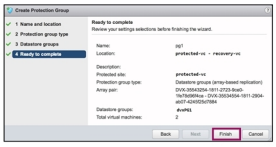

Notes on DVX/SRM Operations

1. The DVX System monitors the protected and recovery sites during planned migration or disaster recovery that is treated as planned migration when both sites are operative. The DVX System will verify that the health status of the DVX protection groups and DVX Snapstores on both sites is OK.

    - The Snapstore must have sufficient space for planned migration operations.
    - The DVX protection groups cannot be in a Critical state.

    If these conditions are not satisfied, the DVX System will prevent SRM from proceeding with the planned migration operation. To resolve this situation, use the DVX UI on both sites to determine the error and resolve the situation.

2. You can change the virtual machine membership of an SRM-enabled protection group by deploying or removing virtual machines and/or changing the VSS-enabled virtual machine membership. Normal replication operations will propagate any membership changes to the recovery site. If the protected site fails before the changes can be propagated, disaster recovery operations will prevent these changes from being applied in the future.

    To avoid this situation, Datrium recommends that whenever the virtual machine membership of an SRM-enabled protection group changes, you perform a manual snapshot and replication to the recovery site.

3. Datrium recommends that you do not use the “Force cleanup” option when performing a reprotect operation.
4. When SRM performs a fail-over, it will unmount the datastore on the protected site. When SRM fails back to the protected site, it re-mounts the datastore, and constructs a datastore name that is the concatenation of the protected and recovery datastore names. To avoid possible confusion, whenever your system executes fail-over / fail-

back operations, you should manually remount the datastore on the primary site before fail-back occurs to specify the appropriate datastore label.

DVX Disaster Recovery Example

If there is some event that causes a source site to go down or become unavailable, you can use replica snapshots on a destination site to resume normal operation. To recover virtual machines and files from replica snapshots, you promote a replica protection group. Promotion changes a replica protection group into a live protection group. A live protection group is capable of supporting source snapshot and replication operations.

It is important to maintain only one live protection group at a time in a set of associated protection groups – a live group and its replicas. The DVX System does not support replication between two live protection groups. When you promote a protection group, you must demote the other live group to change it to a replica group.

The following sequence shows an overview of DR operations.

1. During normal operations, system A is the source system in a replication pair. When the DVX System produces snapshots, it replicates snapshots to the destination DVX System B.

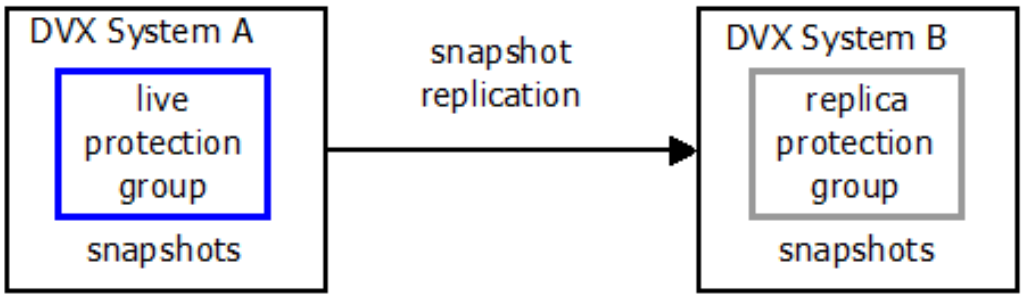

2. If system A goes down, to start the recovery, promote the replica protection group on system B.

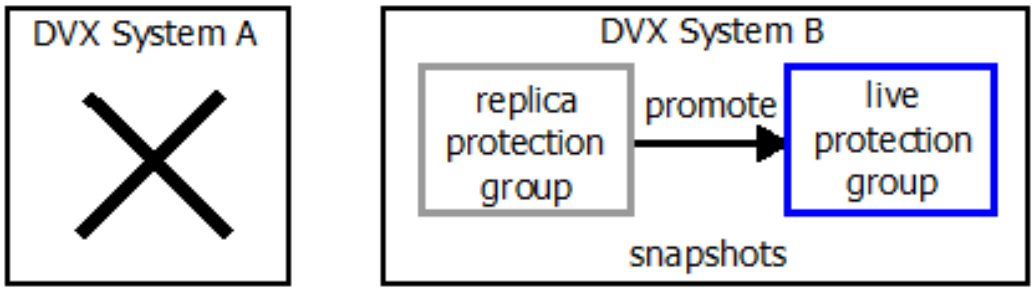

3. On system B, use a snapshot of the promoted protection group to restore the protection group contents to the DVX datastore. You must add any restored virtual machines to the vSphere inventory manually.

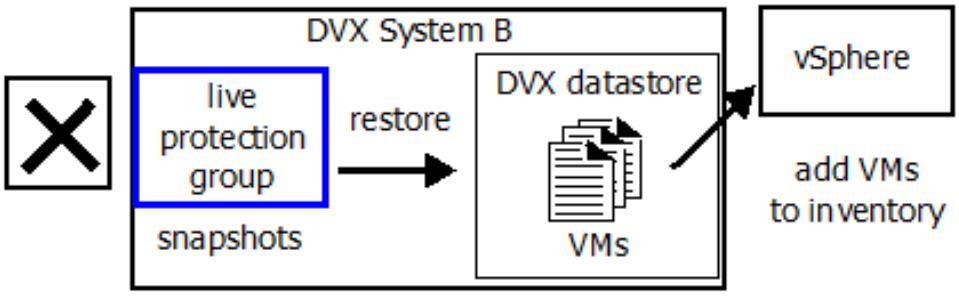

4. Replica protection group schedules are disabled, and replica site definitions have been removed. To complete the recovery on the new live protection group:

   - Add the original source system as a replica site.
   - Add a replica site reference to the snapshot schedule(s).

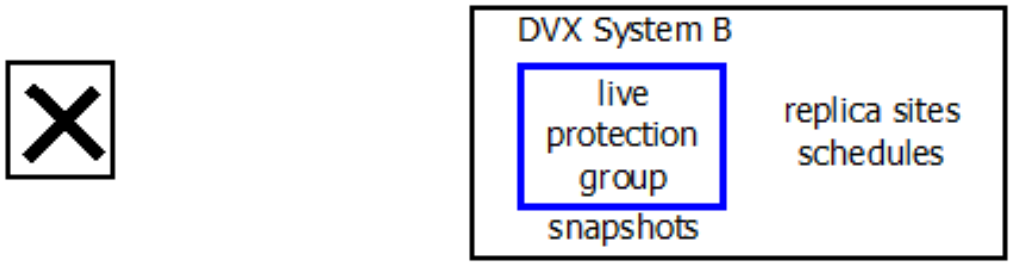

5. When the original source system is operational, change the original source protection group to a replica group. There must be only one live group in a set of associated live and replica protection groups in your environment.

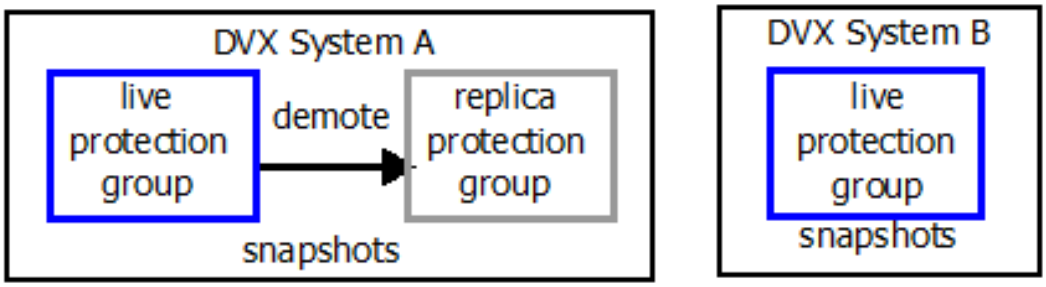

6. Enable the schedule(s) for the live protection group. When snapshot operations are performed on DVX System B, it will send snapshot replicas to DVX System A.

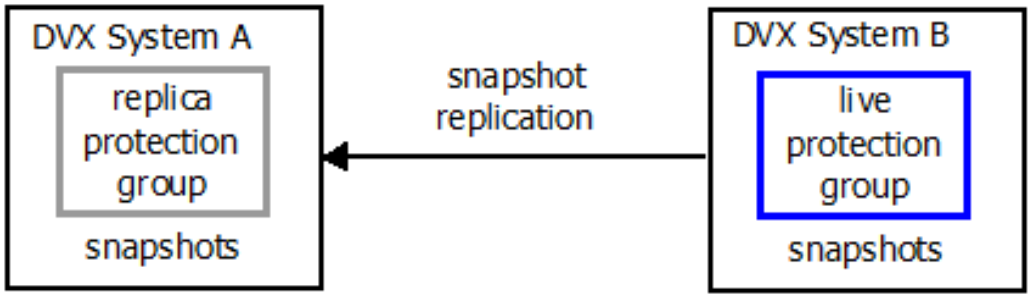

The following sections provide more detail about the procedure to recover protection group contents from a replica site.

Recovering Protection Group Content

1. **Promote**

    a. Select the Protection view.

    b. In the Protection view, select a protection group row in the “Replica groups” frame. The GUI will enable the “Promote” button.

    c. Click the “Promote” button.

2. **Restore**

a. Promotion moves the replica group into the live group frame. The group schedule is disabled.
b. Select the row for the promoted group. The GUI will enable the “Restore” button.
c. Click the “Restore” button. The GUI displays the “Restore group” dialog.
d. Select a snapshot and click the “Restore” button.
e. After you restore the group contents, you must add any virtual machines to the vSphere inventory manually. For more information, see Restore a Virtual Machine.

3. **Replica site references**

Add the original source system as a replica site.

In the “Replica sites” section on the Summary tab of the Protection view, click on the “Add” button. The DVX GUI displays the “Add replica destination” dialog.

You must supply the following information:

- The floating IP address for the Data Node management interface.
- The password for the admin account on the DVX replica site.
- A unique site label.

4. **Replica site references**

a. Add a replica site reference to the snapshot schedule(s).
b. Select the “Live groups” row for the promoted group.
c. Select the “Edit group and schedule” menu entry. The GUI displays the “Edit Protection group” dialog.
d. Replication for the promoted protection group is disabled. To establish replication for the protection group, click on the “Replicate snapshots” check box and fill in the replication fields.

> **Note:** A live protection group will not send replication traffic to a live protection group. A live protection group will not accept replica snapshots.

> The system that you use as a replication destination must be able to accept snapshots for that protection group.

- If the system does not already have that protection group, when replication starts, it can create a replica group to accept the snapshots.
- If the system already has that protection group, it must be a replica protection group. If it is live, you must demote the live group to turn it into a replica group.

5. **Demote**

   Demote the original source protection group. On the Summary tab of the Protection view:

   a. Select the row for the group to be demoted.
   b. Select the “Demote to replica group” entry on the “Edit group and schedule” menu.

6. **Enable schedule**

   a. In the “Live groups” frame of the Protection view, select the row for the promoted protection group and click on the menu icon to display the live group operations. Click on “Enable schedule”.

Recovering Protection Group Content – CLI

Use the following sequence of commands to promote a protection group.

1. `protection groups promote`

   Change the replica protection group into a live protection group.

2. `protection groups restore`

   Restore the protection group contents in the DVX datastore. You must add any restored virtual machines to the vSphere inventory manually.

3. `protection replica-sites add`
   `protection groups schedules edit --replica-site-names`

   Add a replica site definition and the corresponding replica site reference to the schedule(s)

4. `protection replica-sites demote`

   On the source system, change the source protection group to a replica group. There should be only one live group in a set of associated live and replica protection groups in your environment.

   A live protection group will not send replication traffic to a replica protection group. A live protection group will not accept replica snapshots. The system that you use as a replication destination must be able to accept snapshots for that protection group.

   If the system does not already have that protection group, when replication starts, it can create a replica group to accept the snapshots.

   If the system already has that protection group, it must be a replica protection group. If it is live, you must demote the live group to turn it into a replica group.

5. `protection groups enable`

   Start the schedule(s) for snapshots and replication.

Blanket Encryption

The DVX System provides Blanket Encryption for data in use, at rest, and in transit. Blanket Encryption includes both DVX System and replication contexts.

DVX System encryption

- Virtual machine data on host flash drives.
- Data traffic between hosts and the Data Node.

- All new data that is stored on the Data Node, including durable virtual machine data, snapshots, NVRAM data, and files stored in the DVX NFS namespace.
- The DVX System software uses a cryptographic module that is validated to FIPS 140-2. You can use either FIPS 140-2 approved or validated mode of operation. Approved mode is the default mode of operation.

Replication encryption

- The DVX System uses tunnel encryption (TLSv1.2) for replication traffic between source and destination DVX Systems. Replication encryption is enabled separately using the DVX CLI or UI.

DVX Cloud encryption

The DVX System uses SSL to encrypt all data in-transit to and from Cloud DVX. Cloud DVX uses AWS SSE-S3 (Server-Side Encryption) to encrypt the data at rest in the cloud. DVX Cloud encryption is always enabled.

DVX encryption is integrated with the DVX deduplication and compression features so that the DVX System performance and economy of scale are preserved.

> **Tip:** When using Cloud DVX, if you are using if blanket encryption with your on-prem DVX, we strongly recommend you enable replication encryption as well.

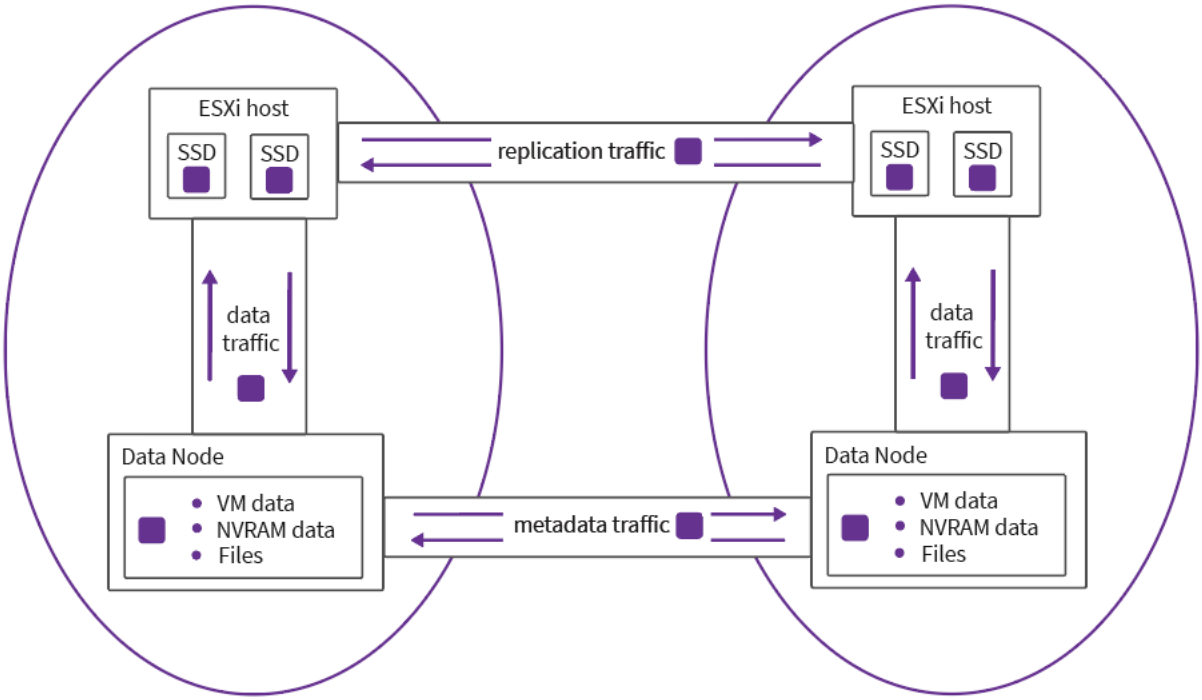

Important: The DVX System uses the AES XTS 256 (Advanced Encryption Standard) algorithm with a key size of 512 bits, and it uses AES-NI (AES New Instruction set). In some environments, you can disable AES-NI in the bios. The DVX System can encrypt data without the use of AES-NI, but it is considerably slower. You should not disable AES-NI.

vSphere 6.5 includes a virtual machine encryption capability. vSphere encryption will have a negative effect on DVX performance and might prevent you from creating DVX snapshots of virtual machines. Datrium recommends that you disable vSphere encryption

Using Blanket Encryption

The following table shows the DVX CLI commands that you use to perform encryption operations.

Encryption Operation	DVX CLI Commands
Enable and disable encryption on the local DVX System.	`datastores encryption enable`

Encryption Operation	DVX CLI Commands
	`datastores encryption disable`
Set FIPS mode (approved or validated).	`datastores encryption set --fips-mode`

Enable and disable encryption of replication traffic.	`protection replica-sites encryption enable`
	
	`protection replica-sites encryption disable`

Change the key that encrypts DVX data.	`datastores encryption rotate-key`

Set the startup mode for access to data.	`datastores set --startup-mode`
	
	`datastores unlock`

Set the encryption passwords. The DVX System supports two encryption passwords. You can use either password to access the encryption sub-system. When creatinig a new password, we recommend that the password contains at least 10 characters, including a combination of uppercase and lowercase letters, numbers, and special characters.	`datastores encryption password set`
Display encryption status.	`datastores encryption show`

Enabling DVX System Encryption

To enable encryption on a DVX System, use the `datastores encryption enable` DVX CLI command. The CLI will prompt for two encryption passwords and password confirmation.

An encryption password should have a minimum of 10 printable characters, and should including a combination of uppercase and lowercase letters, numbers, and special characters within the ASCII 7-bit character set.

The CLI echoes the passwords so that you can verify your input visually. You can use the `--hide-password` argument to prevent password display on the screen.

```
datrium1.node1.controller1>> datastores encryption
enable
New data will be encrypted. Existing data may remain
unencrypted.

You must provide an encryption password and a backup
encryption password.

IMPORTANT! Datrium cannot recover the encryption pass-
word. If you lose both passwords, you will lose access
to all data in the DVX System.

The password must be a minimum of 8 characters long.
By default, the password shows in plain text as you
type.To hide the password, run the command with the --
hide-password argument.

Are you sure you want to enable encryption? {yes|no}
[no] :  yes

REMINDER: Datrium cannot recover the encryption pass-
words. If you lose both passwords, you will lose
access to all data in the DVX System. You will be
prompted to enter and confirm both the encryption pass-
word and the backup password.

            Encryption password : password-example
               Confirm password : password-example

     Backup encryption password : password-backup-ex
        Confirm backup password : password-backup-ex

Encryption enabled.
Startup mode is unlocked.

datrium1.node1.controller1>>
```

An encryption access password provides two levels of secure access:

- When encryption is enabled, you must enter an encryption access password to execute commands that change encryption settings.
- You can use the `datastores encryption set` command to increase access security by setting the startup mode to lock the system. A locked system requires the encryption access password for access to data after you restart or upgrade the Data Node. Use locked startup mode to increase the security of your data during shipping. If you set the startup mode to locked, you must use the `datastores encryption unlock` command when you restart or upgrade the Data Node.

> **Important:** Save the encryption access passwords. The DVX System secures the passwords but you cannot recover the passwords. If you set the startup mode to locked and then lose the encryption access passwords, you will not be able to access your data after restarting or upgrading the Data Node, and you will not be able to change the encryption settings.

When you enable encryption, all new data is encrypted and a password is required to use the encryption commands. After you disable encryption, new data is not encrypted. Previously encrypted data remains encrypted until it is modified or until DVX Space Reclamation processes the data.

Enabling Encryption for a DVX System with Existing Data

When encryption is enabled, the DVX System encrypts new data. If you enable encryption on a new DVX System, all data will be encrypted. If you enable encryption on a DVX System with existing data that is not encrypted, use the following steps to encrypt all of the data.

1. Move the existing data off of the Data Node.
2. Re-initialize the Data Node to its factory settings. Contact Datrium Support for information about performing the re-initialization.
3. Enable encryption.
4. Move the data back onto the Data Node.

Enabling Replication Encryption

To enable replication encryption, use the `protection replica-sites encryption enable` DVX CLI command. This command enables encryption for all

outgoing replication traffic.

```
protection replica-sites encryption enable
```

The DVX System uses tunnel encryption for replication traffic between source and destination sites.

You can also use the DVX GUI to enable replication encryption. In the Summary tab of the Protection view:

1. Click the Replication encryption menu entry in the Replication settings menu.
2. In the Replication encryption dialog, set encryption for outgoing replication traffic.

Encryption Status – DVX CLI

To view the DVX System encryption status, use the datastore encryption show command:

```
dvx05.node1.controller2>> datastores encryption show

------------------------------------------------------------------------------
Name        Encryption   Algorithm    Current state   Startup mode   FIPS 140-2
----------- ------------ ------------ --------------- -------------- -----------
Datastore1  Enabled      AES-XTS-256  Unlocked        Unlocked       Approved
----------- ------------ ------------ --------------- -------------- -----------
```

The show command output includes the following information:

DVX System encryption setting	Indicates whether encryption is enabled or disabled.
Algorithm	If encryption is enabled, the command shows the algorithm used for data encryption.
Current state	Indicates whether the system is currently locked. If you have set the startup mode to “locked”, you must enter the encryption password with the `datastores encryption unlock` command when you restart or upgrade the Data Node. Before you unlock the system, the `datastore encryption show` command will show both the current status and the startup mode as locked.

DVX System encryption setting	Indicates whether encryption is enabled or disabled.
Startup mode	Indicates the current startup mode setting (locked or unlocked). If you set startup mode to “locked”, you must enter the encryption password with the `datastore encryption unlock` command when you restart or upgrade the Data Node.
FIPS 140-2	Indicates the cryptographic mode (approved or validated). The DVX System software uses a cryptographic module that is FIPS 140-2 certified. Approved mode is the default mode of operation.

To view the replication encryption status, use the `protection replica-sites encryption show` CLI command.

Encryption State – DVX GUI

DVX System Encryption

In the DVX GUI, you can see the encryption state in the durable capacity chart on the Summary tab of the Real Time view. When encryption is enabled, the shield icon is shown in green.

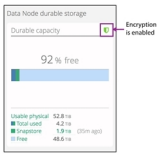

Replication Encryption

The encryption status is shown in the Summary tab of the Protection view. When replication encryption is enabled, the shield icon is shown in green.


Datrium
Protection
Protection
Summary
Replica sites
Source DVX
dvx05
Encryption
is enabled

DVX System Monitoring

The DVX system monitors the following in order to ensure smooth operation in your environment:

- Thermal Threshold
- Fan Speed
- Non Transparent Bridge (NTB)
- Cyclic Redundancy Check (CRC) Error Monitoring for NIC
- DNS Server Reachability
- Duplicate IP Address
- Jumbo Frames
- Backup Battery Unit
- Boot Drive Partition
- PCI Express (PCIe) Device Width and Speed
- Memory Usage
- Enclosure Manager
- Fan speed
- DVX System Monitoring
- Boot Drive Health and SMART Status

Events from DVX monitoring can be views using the CLI using the `events show` command. In some cases, these alerts will also be presented in the DVX GUI.

Thermal Threshold

The Data Node provides multi-stage thermal threshold monitoring and event generation so you can be aware of the device's temperature at all times, and can take active measures if temperatures become too high.

Thermal monitoring events will appear in the DVX GUI under the Data Nodes tab, under general health under Events and Alarms, and also will affect the Data Node hardware health.

Events →	SensorFailedEvent	TemperatureLowWarnEvent	TemperatureNormalEvent	TemperatureHighWarnEvent	TemperatureHighWarnEvent	TemperatureHighWarnEvent
Temperature Range in Celsius	<-0 or >80	1 to 4	5 to 37	38 to 41	42 to 45	46 to 80
Event Severity	WARNING	WARNING	INFO	WARNING	ERROR	EMERGENCY
Health Status	UNKNOWN	OK	OK	OK	CRITICAL	Critical
Alert?	No	No	No	No	Yes	Yes

Fan Speed

The Data Node monitors its cooling fan speed to indicate if the fan speed is low, normal, or high.

If the fan speed is low, it could indicate a malfunction of the fan, and if the fan is too high, the Data Node could be increasing in temperature, which could affect availability and cause other potentially serious damage.

This feature generates the following events:

Events →	SensorFailedEvent	FanSpeedLowEvent	FanSpeedNormalEvent	FanSpeedHighEvent
Fan speed	< = 0 or > = 2000	1 to 2000	2000 to 10800	10801 to 21999

Events →	SensorFailedEvent	FanSpeedLowEvent	FanSpeedNormalEvent	FanSpeedHighEvent
RPM				
Event Severity	WARNING	WARNING	INFO	INFO
Health Status	UNKNOWN	UNKNOWN	OK	OK

Non Transparent Bridge (NTB)

The DVX system monitors the NTB bridge between the two Data Node controllers for "bogus" values from the peer controller and generate events. Some of the alerts are bogus and are not cause for concern, while other alerts are important. The DVX system will alert you regarding this connection and indicate if the issue is serious or not.

Two new events are logged for this type of monitoring:

- `ControllerNTBPeerQueueValuesOKEvent`. Communication over the PCI bridge between controllers is OK and no serious errors are being reported.
- `ControllerNTBPeerQueueValuesInvalidEvent`. An NTB error occurred over the PCI bridge and you should contact support to resolve.

In the DVX GUI, if you see the following alert in the DVX GUI (under the Data Node tab), you should call support immediately for assistance:

"Controller nodeX.controllerX encountered a problem (NTB error). Contact support"

Cyclic Redundancy Check (CRC) Error Monitoring for NIC

The DVX controller monitors for CRC errors on DVX controllers Network Interfaces periodically, potential indicating incompatible or bad cable/tranceiver and would require cable replacement.

This feature generates the following CRC alerts:

- `EthernetPortCRCOKEvent`. This Info even indicates that the network interface has recovered from CRC errors and is good now as there are no CRC errors since 3 hours
- `EthernetPortCRCInfoEvent`. This Info event indicates a low nu.mber of CRC events, within the range of 3001 – 19999
- `EthernetPortCRCWarningEvent`. This Warning event indicate a high number of CRC errors, greater than 20000 .

DNS Server Reachability

The DVX controller monitors DNS server reachability by doing a ping test periodically and generates the following event if a DNS server is unreachable:

- UnreachableDnsServerEvent. A Warning event that indicates that the DNS server configured in your network is unreachable.

Duplicate IP Address

The DVX controller monitors for duplicate IP of all configured hosts and controllers IP addresses and generates the following event when duplicate IP addresses are detected:

- `DuplicateIpEvent`. Indicates a duplicate IP address assigned to controllers.

Jumbo Frames

The DVX system now monitors if a Data Node controller cannot ping its host with Jumbo frames enabled. In this situation, the system will generate events to help narrow any issues on the intermediate switches.

Backup Battery Unit

The Data Node monitors the following for each nodes Backup Battery Unit (BBU):

BBU End Of Life status

The DVX controller software now monitors the Data Node BBU End Of Life status and generates events that alert customers to contact support for BBU replacement.

BBU Charging

The DVX System monitors if BBUs are not charging and arming properly for extended duration of time.

Boot Drive Partition

The DVX System now monitors boot drive usage and generates these two events to indicate normal or high usage:

- `ControllerSystemPartitionUsageOKEvent`
- `ControllerSystemPartitionUsageHighEvent`

PCI Express (PCIe) Device Width and Speed

The Data Node proactively monitors PCIe device width and speed for degradation, and logs events so you can detect hardware failures early.

Memory Usage

The DVX System now periodically monitors Data Node memory usage.

Serial Attached SCSI (SAS) controller reset

The Data Node proactively monitors for unexpected resets of SAS HBA controller and generates events for diagnostic and troubleshooting purposes.

Enclosure Manager

The system monitors the Data Node enclosure management software module, and generates events if it is not responding.

Fan speed

Te system monitors the Date Node fan speed, including a fan's raw RPM value in system events. “ to “Fan speed. This feature monitors the Date Node fan speed, includes fan's raw RPM value in system events.

Baseboard Management Controller (BMC)

The system monitors the BMC during boot and run-time and generates events when it is not responding. The software will attempt to auto-recover by issuing reset to BMC during boot for all the Data Node models. Runtime recovery of BMC is applicable to Data Node models D12X4B, D12X4C, D12X10D, F24X2B, and F24X2D.

Boot Drive Health and SMART Status

The DVX System monitors Data Node boot drive health and Self-Monitoring, Analysis and Reporting Technology (SMART) status, and will generate an event when a drive's SMART status is not OK.

DVX System Monitoring

The DVX system monitors the following in order to ensure smooth operation in your environment:

- Thermal Threshold
- Fan Speed
- Non Transparent Bridge (NTB)
- Cyclic Redundancy Check (CRC) Error Monitoring for NIC
- DNS Server Reachability
- Duplicate IP Address
- Jumbo Frames
- Backup Battery Unit
- Boot Drive Partition
- PCI Express (PCIe) Device Width and Speed
- Memory Usage
- Enclosure Manager
- Fan speed
- DVX System Monitoring
- Boot Drive Health and SMART Status

Events from DVX monitoring can be views using the CLI using the `events show` command. In some cases, these alerts will also be presented in the DVX GUI.

Thermal Threshold

The Data Node provides multi-stage thermal threshold monitoring and event generation so you can be aware of the device's temperature at all times, and can take active measures if temperatures become too high.

Thermal monitoring events will appear in the DVX GUI under the Data Nodes tab, under general health under Events and Alarms, and also will affect the Data Node hardware health.

Events →	SensorFailedEvent	TemperatureLowWarnEvent	TemperatureNormalEvent	TemperatureHighWarnEvent	TemperatureHighWarnEvent	TemperatureHighWarnEvent
Temperature Range in Celsius	<-0 or >80	1 to 4	5 to 37	38 to 41	42 to 45	46 to 80
Event Severity	WARNING	WARNING	INFO	WARNING	ERROR	EMERGENCY
Health Status	UNKNOWN	OK	OK	OK	CRITICAL	Critical
Alert?	No	No	No	No	Yes	Yes

Fan Speed

The Data Node monitors its cooling fan speed to indicate if the fan speed is low, normal, or high.

If the fan speed is low, it could indicate a malfunction of the fan, and if the fan is too high, the Data Node could be increasing in temperature, which could affect availability and cause other potentially serious damage.

This feature generates the following events:

Events →	SensorFailedEvent	FanSpeedLowEvent	FanSpeedNormalEvent	FanSpeedHighEvent
Fan speed	< = 0 or > = 2000	1 to 2000	2000 to 10800	10801 to 21999

Events →	SensorFailedEvent	FanSpeedLowEvent	FanSpeedNormalEvent	FanSpeedHighEvent
RPM				
Event Severity	WARNING	WARNING	INFO	INFO
Health Status	UNKNOWN	UNKNOWN	OK	OK

Non Transparent Bridge (NTB)

The DVX system monitors the NTB bridge between the two Data Node controllers for "bogus" values from the peer controller and generate events. Some of the alerts are bogus and are not cause for concern, while other alerts are important. The DVX system will alert you regarding this connection and indicate if the issue is serious or not.

Two new events are logged for this type of monitoring:

- `ControllerNTBPeerQueueValuesOKEvent`. Communication over the PCI bridge between controllers is OK and no serious errors are being reported.
- `ControllerNTBPeerQueueValuesInvalidEvent`. An NTB error occurred over the PCI bridge and you should contact support to resolve.

In the DVX GUI, if you see the following alert in the DVX GUI (under the Data Node tab), you should call support immediately for assistance:

"Controller nodeX.controllerX encountered a problem (NTB error). Contact support"

Cyclic Redundancy Check (CRC) Error Monitoring for NIC

The DVX controller monitors for CRC errors on DVX controllers Network Interfaces periodically, potential indicating incompatible or bad cable/tranceiver and would require cable replacement.

This feature generates the following CRC alerts:

- `EthernetPortCRCOKEvent`. This Info even indicates that the network interface has recovered from CRC errors and is good now as there are no CRC errors since 3 hours
- `EthernetPortCRCInfoEvent`. This Info event indicates a low nu.mber of CRC events, within the range of 3001 – 19999
- `EthernetPortCRCWarningEvent`. This Warning event indicate a high number of CRC errors, greater than 20000 .

DNS Server Reachability

The DVX controller monitors DNS server reachability by doing a ping test periodically and generates the following event if a DNS server is unreachable:

- UnreachableDnsServerEvent. A Warning event that indicates that the DNS server configured in your network is unreachable.

Duplicate IP Address

The DVX controller monitors for duplicate IP of all configured hosts and controllers IP addresses and generates the following event when duplicate IP addresses are detected:

- `DuplicateIpEvent`. Indicates a duplicate IP address assigned to controllers.

Jumbo Frames

The DVX system now monitors if a Data Node controller cannot ping its host with Jumbo frames enabled. In this situation, the system will generate events to help narrow any issues on the intermediate switches.

Backup Battery Unit

The Data Node monitors the following for each nodes Backup Battery Unit (BBU):

BBU End Of Life status

The DVX controller software now monitors the Data Node BBU End Of Life status and generates events that alert customers to contact support for BBU replacement.

BBU Charging

The DVX System monitors if BBUs are not charging and arming properly for extended duration of time.

Boot Drive Partition

The DVX System now monitors boot drive usage and generates these two events to indicate normal or high usage:

- `ControllerSystemPartitionUsageOKEvent`
- `ControllerSystemPartitionUsageHighEvent`

PCI Express (PCIe) Device Width and Speed

The Data Node proactively monitors PCIe device width and speed for degradation, and logs events so you can detect hardware failures early.

Memory Usage

The DVX System now periodically monitors Data Node memory usage.

Serial Attached SCSI (SAS) controller reset

The Data Node proactively monitors for unexpected resets of SAS HBA controller and generates events for diagnostic and troubleshooting purposes.

Enclosure Manager

The system monitors the Data Node enclosure management software module, and generates events if it is not responding.

Fan speed

Te system monitors the Date Node fan speed, including a fan's raw RPM value in system events. “ to “Fan speed. This feature monitors the Date Node fan speed, includes fan's raw RPM value in system events.

Baseboard Management Controller (BMC)

The system monitors the BMC during boot and run-time and generates events when it is not responding. The software will attempt to auto-recover by issuing reset to BMC during boot for all the Data Node models. Runtime recovery of BMC is applicable to Data Node models D12X4B, D12X4C, D12X10D, F24X2B, and F24X2D.

Boot Drive Health and SMART Status

The DVX System monitors Data Node boot drive health and Self-Monitoring, Analysis and Reporting Technology (SMART) status, and will generate an event when a drive's SMART status is not OK.

Datrium Support for the DVX System

Datrium provides software-based support for the DVX System:

- Datrium Support Portal
- Network Access for Datrium Support
- DVX Autosupport
- Manual Support Submission
- DVX Remote Support
- DVX System Upgrade
- VAAI VIB Upgrade
- DVX Software Releases

Datrium also provides online support through the Datrium Support Portal.

Datrium Support Portal

Datrium provides online support through the Datrium Support Portal at support.datrium.com. Your Datrium account team sets up a support account for you. You can use the Portal to ask a question, report a problem, or request a new feature. It provides access to product documentation, parts replacement documentation, Knowledge Base articles, and software downloads. The Portal also supports a search capability that gives you access to all of the information on the Portal.

Network Access for Datrium Support

All DVX support capabilities require access to Datrium Support servers. When you configure the network interface(s) on the Data Node, you provide the IP address for a gateway to the Internet

To see the gateway IP address in the DVX GUI:

1. Select the Data Nodes view.
2. Select the Network tab to display the Data Node network configuration.

3. To use the DVX CLI to display the gateway IP address, use the `network show` command.

Data Node gateway configuration is necessary to support DVX System access to the DVX Autosupport, DVX Remote Support, and DVX System Upgrade servers. The following table shows the network interface information for this access.

DVX Capability	Network Configuration	Data Node Interface	Datrium Server, Port, and Protocol
Auto Support	Combined data and management	Data interface: Floating IP address and both controller IP addresses	server: autosupport.datrium.com port: 443 protocol: HTTPS
	Separate data and management	Management interface: Floating IP address and both controller IP addresses	server: autosupport.datrium.com port: 443 protocol: HTTPS
Remote Support	Combined data and management	Data interface: Floating IP address and both controller IP addresses	servers: autosupport-tunnel.datrium.com autosupport-tunnel-https.datrium.com port: 443 protocol: HTTPS
	Separate data and management	Management interface: Floating IP address and both controller IP addresses	servers: autosupport-tunnel.datrium.com autosupport-tunnel-https.datrium.com port: 443 protocol: HTTPS
Software	Combined	Data interface:	server: upgrade-center-01.datrium.com

DVX Capability	Network Configuration	Data Node Interface	Datrium Server, Port, and Protocol
Upgrade	data and management	Floating IP address and both controller IP addresses	port: 443 protocol: HTTPS
	Separate data and management	Management interface: Floating IP address and both controller IP addresses	server: upgrade-center-01.datrium.com port: 443 protocol: HTTPS

You can use the DVX CLI command `support test` to verify access to the Auto Support server at `autosupport.datrium.com:80`.

If there is a problem, you might see the following HTTP error:

```
HTTP connection to support server failed due to bad
HTTP status code 404
```

- If you are using a firewall, make sure that you have configured access appropriately.
- If you are using an HTTP proxy in your environment, use the config web-proxy set command to identify the proxy server and port for HTTP access.
- The DVX System uses the management interface for access to the gateway. Make sure that your management subnet provides access to the gateway.

DVX Autosupport

The DVX System includes support software that monitors the storage system and sends log data to Datrium Support. The System sends two kinds of log data to Datrium. This support capability is automatic.

- The DVX System sends a small amount of heartbeat data to Datrium every five minutes.
- The DVX System sends accumulated statistics and log data once a day.

Manual Support Submission

You can use either the GUI or CLI to perform a manual support submission. When you submit a support bundle, the DVX System collects support data, creates a support bundle, and sends it to autosupport.datrium.com. You would normally submit a support bundle after consultation with Datrium Support personnel.

Support Submission (GUI

To submit a support bundle from the DVX GUI, use the “Submit support bundle” entry on DVX system menu. To display this menu, click on the gear icon in the upper right corner of the DVX window.

When you use the GUI to submit a support bundle, the DVX System creates the support bundle using data for all hosts and Data Node controllers in the System.

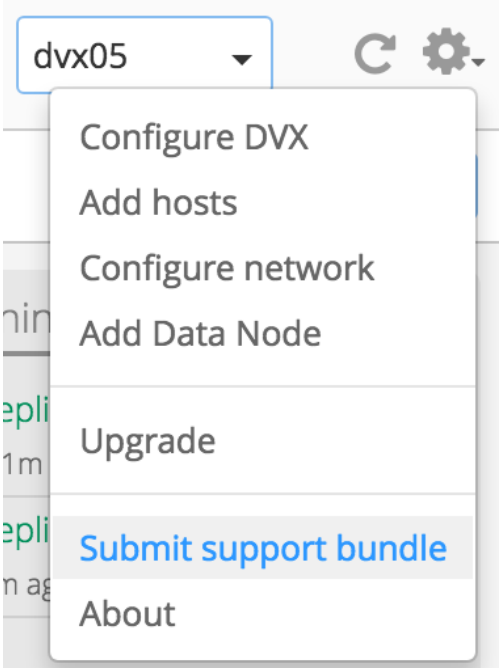

Support Submission (CLI)

Use the support submit command to perform a manual support submission. By default, the command sends data for all hosts and Data Node controllers in the DVX System. You can

use command line arguments to limit the amount of data that is submitted. By default, the System collects all support data that was generated since the last manual support submission.

```
support submit [--entity-id entityId [entityId ...]]
                        [--all-files]
                        [--timespan-hours hours]
```

- Use the `--entity-id` argument to limit support data to one or more hosts and/or controllers. Use tab completion to see the valid entity identifiers. If you do not specify this argument, the DVX System submits data for all entities.
- Use the `--all-files` argument to direct the DVX System to collect data for all available support files associated with the specified entities, even if some were submitted previously.
- Use the `--timespan-hours` argument to indicate the number of hours of data for collection. If you also specify `--all-files`, the `--all-files` argument overrides the timespan specification.
- If the command completes successfully, it displays an identifier for the support bundle. Use this identifier with the `support show` command.

The following table provides brief descriptions of the CLI support context commands for support submission.

`support test`	Verify access to the Datrium Support server.
`support submit`	Manually submit a support bundle to Datrium Support.
`support show`	Monitor the status of a manually submitted support bundle.

DVX Remote Support

In the context of a support call, it may be useful for Datrium Support personnel to have network access to the Data Node. When you enable remote support, Datrium Support personnel can login to the Data Node for the purpose of running diagnostics and collecting data. You must explicitly enable remote access.

The DVX CLI provides the following commands to manage access to the Data Node for Datrium Support operations.

`support test`	Verify the path between the Data Node and the Datrium support server.
`support remote-access show`	Display the current setting for remote access.
`support remote-access enable`	Allow remote access to the Data Node.
`support remote-access disable`	Disable remote access to the Data Node.

DVX System Upgrade

The DVX System provides a simple upgrade procedure. You can use either the GUI or the CLI to download software and install it. The DVX System manages the procedure so that DVX system operation is minimally impacted during the upgrade. If you are monitoring real-time activity, you might notice a short interruption of approximately one minute. There is no need to take any additional action.

The following sections provide information about DVX System upgrade:

- Upgrade Overview
- Before You Upgrade
- Upgrade (DVX GUI)
- Upgrade (CLI)
- DVX System Upgrade

Upgrade Overview

The following table provides a brief overview of DVX software installation when you perform the upgrade.

Data Node	The DVX System upgrades the Data Nodes in the Data Pool immediately. As part of this process, it upgrades the software on

	both controllers of each node in the Pool. DVX System upgrade might also perform automatic firmware upgrade. If there is a firmware upgrade, the complete system upgrade can take up to two hours.
DVX Hyperdriver	Hyperdriver software on each host detects the software version change and retrieves DVX system software from the Data Node to bring the Hyperdriver into sync with the Data Node.
DVX GUI Plug-in	The DVX System updates the DVX GUI plug-in registration on the vCenter Server. The vCenter Server will update the Web Client with the new DVX GUI software within a short period of time.
AWS Cloud DVX	If you are using Cloud DVX storage, Cloud DVX performs automatic upgrades.

The DVX System also includes VAAI integration software that runs on the ESXi hosts. DVX VAAI VIB upgrades are infrequent. This software is distributed as an independent VIB that is a separate software download. You can use the VMware vSphere Update Manager or the esxcli to install the VAAI VIB. See VAAI VIB Upgrade.

> **Note:** You should not remove or replace any Data Node components while you are upgrading the DVX System.

Before You Upgrade

Before you upgrade the DVX System, verify network and host access and then run a pre-upgrade check.

Note that the DVX System upgrade requires access to your DVX registered vCenter Servers on ports 80 and 443.

> **Important:** During the DVX software upgrade process, the DVX Data Node will continue to serve I/Os. However, for a short duration during upgrade, I/O services are stopped to switch to the newer version of DVX software. This short duration falls within VMware's default APD timeout window of 140 seconds. During this period, some I/Os from the host connected to the DVX might experience high latency. If you have any applications

that can not handle such IO delays during upgrade, you might need to shut those applications down for maintenance during the upgrade process.

Network Access for Upgrade

To verify network access to the Data Node configuration, perform the following steps for each floating IP address.

1. Ping the Data Node floating IP address(es) from a remote system.
2. Log in to the Data Node and use the DVX CLI `nodes failover` command to change the active controller on the Data Node.
3. Ping the Data Node floating IP address(es) again.

   If any of the ping operations fail, verify your switch configuration.

Host Access for Upgrade

Verify that none of the hosts that use Datrium storage are in lockdown mode. If any of the hosts are in lockdown mode, you cannot upgrade the DVX System.

Pre-Upgrade Check

Before you upgrade the DVX System, you can run an upgrade check to determine if you can perform an upgrade. You can use the DVX GUI or the DVX CLI to perform an upgrade check.

The upgrade check process determines the pre-upgrade status according to the following criteria:

- All Data Node processes are running.
- All host Hyperdriver processes are running.
- Host SSDs have enough space for the upgrade.
- Host SSDs have the appropriate DVX partitions.
- Host has enough temporary space.

If you see an error related to these checks, resolve the situation and then perform the upgrade-check again.

Using the DVX GUI to Perform an Upgrade Check

To perform an upgrade check, download the available software from the Datrium Upgrade Center and then use the DVX GUI to run the upgrade check. If the upgrade check process completes successfully, the Data Node and the hosts that use it are in a state that can be upgraded to the latest version.

1. Use the following procedure to perform an upgrade check:
2. Click Upgrade in the DVX GUI system menu. The GUI displays the upgrade dialog.
3. Download the available software, but do not upgrade. The upgrade dialog will show the software as downloaded.

Click on the Pre-check button.

Using the DVX CLI to Perform an Upgrade Check

To perform an upgrade check, download the available software from the Datrium Upgrade Center and then use the DVX CLI to run the upgrade check. If the upgrade check process completes successfully, the Data Node and the hosts that use it are in a state that can be upgraded to the latest version.

Use the following procedure to perform an upgrade check:

1. Use the `dvx software show` command to display the available DVX software version, if there is one.
2. Use the `dvx software download` command to download the latest version.
3. Use the `dvx software upgrade-check` command to perform an upgrade check.

   ```
   dvx software upgrade-check version [--output-format
   json]
   ```

Pre-Upgrade Check Events

The following table shows upgrade check events. The upgrade process also runs a pre-upgrade check, so these events might be displayed during an upgrade.

Event Type / Message Text	Description
UpgradeCheckStartEvent Upgrade check of *version* started	The upgrade check process generates this event at the beginning of the upgrade check.
UpgradeCheckSuccessEvent Upgrade check of *version* passed	The upgrade check process completed successfully.
UpgradeCheckFailEvent Upgrade check of *version* failed with n error(s)	The upgrade check process failed. Possible causes: • Data Node or host Hyperdriver processes might not be running. • Host SSD(s) might not have enough space for the upgrade. • DVX partitions on host SSD(s) might not be configured correctly.
UpgradeCheckInternalErrorEvent Upgrade check error: *msg-text*	The upgrade check process encountered an internal error. Contact Datrium Support.
UpgradeCheckErrorDetailEvent Upgrade check error: *msg-text*	The upgrade check process encountered an error. The message text provides an explanation.

Upgrade (DVX GUI)

To upgrade the DVX software from the GUI, select the “Upgrade” menu entry on the DVX system menu. To display this menu, click on the gear icon in the upper right corner of the DVX window. This menu is available on the Datrium Storage home page, DVX hosts page, and the Data Node page. (The navigation pane on the left of the Web Client window contains links to these pages.)

The following picture shows the DVX menu and the dialog that is produced when you select Upgrade.

GUI-based Upgrade

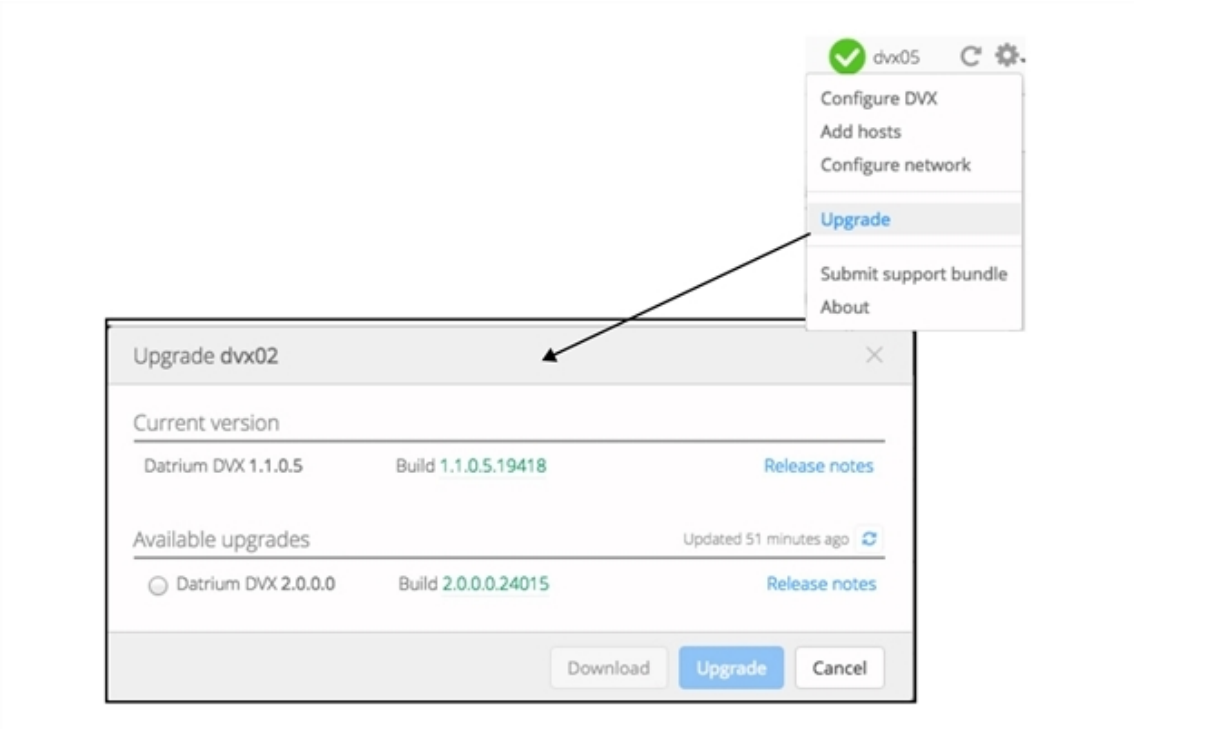

The upgrade dialog indicates the current DVX software version that is installed on your system and it indicates the version that is available for download and installation. The available upgrade entry will indicate whether you have downloaded the upgrade.

- Select the version for the upgrade.
- Choose Download to retrieve the software for installation at a later time.

- Choose Upgrade to download and install the software now.

The upgrade bundle includes Data Node software, Hyperdriver software, and DVX GUI plug-in software. When you perform the upgrade, the DVX System automatically updates all three components.

The upgrade task performs the following actions:

1. Download the upgrade bundle.
2. Upgrade the standby controller.
3. Resync with the standby controller.

After the DVX System has completed these actions, it displays a message indicating the task has been completed. During the upgrade, you can monitor the progress by moving your cursor over the Upgrade task label in the task list on the right hand side of the display.

The following picture shows the sequence of events during the upgrade.

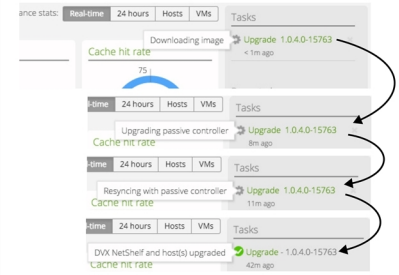

Upgrade Events

The event log will show the results of the upgrade. The table below shows upgrade events. The upgrade process also runs a pre-upgrade check. For information about upgrade-check events that might be displayed during upgrade, see Pre-Upgrade Check Events.

Event Type / Message Text	Description
UpgradeStartEvent Upgrade from *version* to *version* started	Every upgrade procedure generates this event at the beginning of the upgrade.

Event Type / Message Text	Description
UpgradeSuccessEvent Upgrade from *version* to *version* completed successfully	The upgrade process was successful.
UpgradeFailToStartEvent Upgrade from *version* to *version* failed to start	The upgrade process did not start. Possible causes: • Not enough disk space for the download. • Download failure. • Pre-upgrade check error.
UpgradeFailAndRollbackEvent Upgrade from *version* to *version* failed and rolled back	Upgrade started and failed. The upgrade process rolled back the DVX System to the state that it was in prior to the upgrade.
UpgradeRollbackFailEvent Upgrade from *version* to *version* failed and rollback failed	The upgrade process failed; the DVX System tried to rollback to the previous version but failed.
UpgradeCompleteWithErrorEvent Upgrade from *version* to *version* completed with error	The DVX System was upgraded to the new software version, but some error occurred afterwards. The DVX System cannot be rolled back to the earlier version. Possible causes: • Failed to upgrade the former active controller. • Post-upgrade HA (High Availability) re-sync error.

Upgrade Verification

To verify the results after you have upgraded the System, use the DVX GUI plugin to display the “About” dialog and compare the version number.

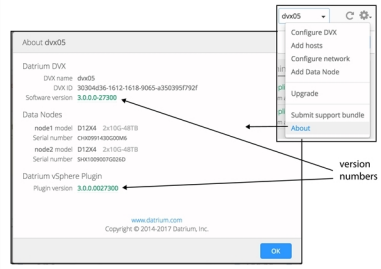

The vSphere plug-in version number must match the DVX System software version number. The vCenter Server might not load the plug-in. You must restart the vSphere Web Client service to force the Web Client to load the plug-in. Use the following command(s) to restart the Web Client on vCenter Server Virtual Appliance (vCSA) or on a Windows installation:

vCenter Server/VMware KB article	Command(s)
vCSA	service vsphere-client restart
Stopping, starting, or restarting vCenter Server Appliance services (2054085)	
pre-6.0 Windows vCenter installation	services.msc
How to stop, start, or restart vCenter Server services (1003895)	
6.0 and later Windows vCenter installation	service-control --start vsphere-client service-control --stop vsphere-client
Stopping, starting, or restarting VMware vCenter Server 6.0 services (2109881)	

Upgrade (CLI)

The DVX CLI provides the following commands to support DVX software upgrade:

`dvx software show`	Displays the versions of the currently installed software and of the software that is available for upgrade. It also displays the upgrade task history.
`dvx software download`	Downloads DVX software but does not install it.
`dvx software upgrade-check`	Checks the DVX system to determine if you can perform an upgrade. The upgrade-check process uses an image that you have downloaded from the Datrium Upgrade Center. If the upgrade-check command completes successfully, the Data Node and the hosts that use it are in a state that can be upgraded to the latest version.
`dvx software upgrade`	Downloads DVX software and installs it. During upgrade, your CLI session will be terminated. You must log in again.

VAAI VIB Upgrade

The DVX software that runs on a host is divided into two VIBs:

- Hyperdriver VIB – The DVX System performs the Hyperdriver VIB upgrade automatically.
- VAAI integration VIB – The VAAI VIB is not required for DVX operation. Install it to use the DVX support for cloning. The VAAI VIB is available from the Datrium Support portal. The VIB is available from the portal as a zip file.

If you have previously used VUM (vSphere Update Manager) to install or upgrade the VAAI VIB, then VUM will remember the settings for the DVX VAAI VIB. The following sections provide information about installing the DVX VAAI VIB.

Hyperdriver VIB for ESXi 7 upgrade considerations with DVX 5.1.4.2

- If you are running DVX 5.1.4.1, your DVX system will not be able to perform a Hyperdriver 7 VIB scan and install correctly. If you want to use ESXi 7 and install

Hyperdriver 7 VIBs, you must first upgrade to the DVX 5.1.4.2 patch.

- In a multi-DVX setup where vCenter is connected to both a DVX version 5.1.4.1 (vSphere 6.5 or 6.7) and DVX version 5.1.4.2 (vSphere 7), if you try to add an ESXi 7 host to a DVX 5.1.4.1 system using the vCenter GUI plugin, the GUI will show an invalid ESXi version error, but the error message will incorrectly state that the host needs to be ESXi 6.0 or later. For a workaround, upgrade your DVX system to the 5.1.4.2 patch.
- In vSphere 7, the vSphere Update Manager (VUM) has been renamed to vSphere Lifecycle Manager (vLCM).

VIB locations on the DVX controller

Both the currently signed VIBs and the new VIBs are packaged on the DVX controller in the DVX 5.1.4.2 release. When uploading VIB to vCenter, the admin can provide URL for VIB location, in this case controller, and then the admin should create a baseline specifying this VIB, and then finally apply the baseline to host(s).

Hyperdriver VIB

- Prior to DVX 5.1.4.2, the DVX Hyperdriver VIB was located in the `ESXVIBHyperDriver` directory on the DVX controller.
- With DVX 5.1.4.2, the old DVX Hyperdriver VIB has been moved into the `ESXVIBHyperDriver6` directory. The new Hyperdriver 7 VIB is located in the "`ESXVIBHyperDriver7`" directory on the controller.

VAAI integration VIB

- Prior to DVX 5.1.4.2, the DVX VAAI VIB was located in the "`esxVibVAAI`" directory on the DVX controller.
- Now with DVX 5.1.4.2, the old DVX VAAI VIB has been moved into the "`ESXVIBVAAI6`" directory. The new DVX VAAI VIB is located in the "`ESXVIBVAAI7`" directory on the controller.

VIB Offline Depot zip file locations on Data Node

ESX 6.x hyperdriver VIB:

```
http://mgmt-floating-IP-address/static/datrium-hyperdriver-
esx6.zip
```

ESX 7.x hyperdriver VIB:

```
http://mgmt-floating-IP-address/static/datrium-hyperdriver-
esx7.zip
```

```
ESXi 6 VAAI VIB:

http://mgmt-floating-IP-address/static/datrium-vaai-esx6.zip
```

ESXi 7 VAAI VIB

```
http://mgmt-floating-IP-address/static/datrium-vaai-esx7.zip
```

DVX VAAI VIB Installation

The VAAI VIB will require ESXi host reboot. The DVX system will not support VAAI storage operations until the next time you reboot the ESXi host(s).

Manual Installation of the DVX VAAI VIB

The following procedure provides a brief overview of how to use vSphere Update Manager (VUM) to install the DVX VAAI VIB on an ESXi host.

1. Add a download to the VUM repository. Use the Datrium VAAI VIB zip file:

   ```
   datrium-vaai_2.2.1.0-38226.zip
   ```
2. Validate and apply the download, then download the VIB.
3. Create a separate baseline for the VAAI VIB.
4. Install the Datrium VAAI NAS plugin (from the Hosts & Clusters page, Update Manager tab).
5. Attach the baseline.
6. Remediate.

Automatic Download and Installation of the DVX VAAI VIB

To use the VUM capability for automatic update, add the DVX VAAI VIB zip file to the patch repository:

```
datrium-vaai_2.2.1.0-38226.zip
```

If you use download sources for automatic update, use the DVX VAAI VIB that is located on the Data Node.

Install the DVX VAAI VIBs with esxcli

Use the following steps to copy the DVX VAAI VIB and install it on the ESXi host.

Install the VAAI VIB on the host for vSphere 6:

```
esxcli software vib install -d http://mgmt-floating-IP-
address/static/ESXVIBVAAI6
```

Install the VAAI VIB on the host for vSphere 7:

```
esxcli software vib install -d http://mgmt-floating-IP-
address/static/ESXVIBVAAI7
```

Install the Hyperdriver VIB on the host for vSphere 6:

```
esxcli software vib install -d http://mgmt-floating-IP-
address/static/ESXVIBHyperDriver6
```

Install the Hyperdriver VIB on the host for vSphere 7:

```
esxcli software vib install -d http://mgmt-floating-IP-
address/static/ESXVIBHyperDriver7
```

After installing the Hyperdriver VIB

After installing the Hyperdriver VIB on the ESXi host, sometimes the Datrium namespace is not added to the esxcli command.

In this case, you will see the following error:

```
esxcli datriumError: Unknown command or namespace datrium
```

This can be fixed by restarting hostd on the ESXi host. To restart hostd, SSH into the host as root, and then run this command:

```
/etc/init.d/hostd restart
```

Alternatively, you can run the localcli command to access the datrium namespace without the need to restart the hostd. For example:

```
localcli datrium
```