

NetWitness[®] Platform

Version 12.5

Storage Guide

Contact Information

NetWitness Community at <https://community.netwitness.com> contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA and other trademarks are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to <https://www.rsa.com/en-us/company/rsa-trademarks>. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability. This software is subject to change without notice and should not be construed as a commitment by RSA.

It is advised not to deploy third-party repos or perform any change to the underlying NetWitness Operating System that is not part of the supported NetWitness version. Any such change outside of the NetWitness approved image may result in a service or functionality conflict and require a reimage of the NetWitness system to bring NetWitness back to an optimized functional state. In the event a third-party repo is deployed, or other non-supported change is made by the customer without NetWitness approval, the customer takes full responsibility for any system malfunction until the issue can be remediated through troubleshooting efforts or a reimage of the service.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on NetWitness Community. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Miscellaneous

This product, this software, the associated documentations as well as the contents are subject to NetWitness' standard Terms and Conditions in effect as of the issuance date of this documentation and which can be found at <https://www.netwitness.com/standard-form-agreements/>.

© 2024 RSA Security LLC or its affiliates. All Rights Reserved.

September, 2024

Contents

Storage Overview	6
Storage Requirements	7
Drive Specifications	7
Required NetWitness Platform Storage Volumes	7
Performance Recommendations	9
Input/Output Operations Per Second	9
General Description of How NetWitness Platform Hosts Store Data	9
Enable Encryption on Series 6E or Series 7 Hosts	10
Configure Drive Pack(s)	11
Benefits of Series 6/6E and Series 7 Drive Pack	11
Decoder Meta Use Cases	12
Concentrator Index Use Cases	21
Enable Security on SED Capable Drives	24
Prepare Virtual or Cloud Storage	25
Decoder, Log Decoder, Concentrator, Archiver	25
NW Server, ESA Primary, ESA Secondary and Malware Analysis	25
Log Collector	26
Endpoint Log Hybrid	26
Additional Endpoint Log Hybrid Partitions	30
UEBA	31
Configure Storage Using the REST API	32
REST API Storage Configuration Commands	32
Storage Configuration Tasks	33
Task 1 - Attach Storage to the Host and Access the REST API Storage Commands	33
Task 2 - (Conditional) RAID Configuration for PowerVault and DACs	35
Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes	39
Task 4 - Allocate Volume Groups to NetWitness Services - srvAlloc	41
Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture	41
Prepare Unity Storage	44
Task 1 - Access Unisphere User Interface (UI)	45
Task 2 - Create Pools	46
Task 3 - Create LUNS	49
Task 4 - Register Hosts	51
Task 5 - Assign LUNS to Hosts	53
Task 6 - Install PowerPath	55

Migrate Data to Another Storage Type	57
Migrate Data Using the Warm and Hot Tier Option	57
Stop the Service	57
Set Up PowerVault	57
Configure The Mount Points	58
Set up Warm and Hot Tiers	59
Decommision the DAC	61
Move Data From DAC to PowerVault	62
Data on PowerVault After Move from DAC	65
SASE Node-x (Decoder/Concentrator) - GCP Persistent Disk (PD) Storage Configuration	66
Introduction	66
Identify Storage Requirements	67
Identify or Define Storage Model	68
Deploy SASE Node(s)	71
Configure SASE Node(s) Storage	72
Configure SASE Decoder Storage	72
Configure SASE Concentrator Storage	72
Extend Storage for SASE Node	72
Extend Decoder or Concentrator Storage	73
Appendix	73
Appendix A - Defining a Custom Host Model	73
Sample Custom Model Definition for Decoder	74
Sample Custom Model Definition for Concentrator	75
Appendix B - Sample Scenario for Configuring SASE Decoder Storage	77
Appendix C - Sample Scenario for Configuring SASE Concentrator Storage	81
Appendix D - Sample Scenario for Extending for SASE Decoder Storage	85
Appendix E - Sample Scenario for Extending SASE Concentrator Storage	86
Appendix A. How NetWitness Platform Hosts Store Data	87
Decoder Hosts	87
Concentrator Host	87
Archiver Host	88
Hybrid Hosts	88
Options for SAN Configurations	88
Performance Recommendations	88
Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives	88
Appendix B. Encrypt a Series 6E or Series 7 Core or Hybrid Host (encryptSedVd.py)	92
Enable SED on configured Drive Groups	94
Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage)	97

Enable SED on configured Virtual Drives/ Drive Groups on Power Vault (PERC 840)	99
Enable Virtual Drives / Drive Groups - PERC H840 Adaptors	99
Appendix C. Troubleshooting	107
Reconfigure Pre-Configured DAC Attached to Decoder Using REST API	107
Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs	108
Configure Storage for Archiver	108
Configure Storage for Network (Packet) Decoder	111
Configure Storage for Network Concentrator	123
Configure Storage for Log Decoder Hybrid	129
Appendix E. Sample Storage Configuration Scenarios for 8 or 12-Drive PowerVault	134
Configure Storage for Archiver using NW-PV-A/NW-PV-A-N	134
Configure Storage for Decoder using NW-PV-B/NW-PV-B-N	137
Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N	142
Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N	145
Configure Storage for Log Hybrid using NW-PV-A/NW-PV-A-N	148
Configure Storage for Network Hybrid using NW-PV-A/NW-PV-A-N	152
Configure Storage for Endpoint Log Hybrid using NW-PV-A/NW-PV-A-N	157
Appendix F. Sample Storage Configuration Scenarios for S7 Physical Hosts with 12-Drive PowerVault MD2412	163
Configure Storage for Archiver using MD2412	163
Configure Storage for Decoder using MD2412	170
Configure Storage for Concentrator using MD2412	178
Configure Storage for Log Hybrid using MD2412	211
Configure Storage for Network Hybrid using MD2412	219
Configure Storage for Endpoint Log Hybrid using MD2412	224

Storage Overview

This guide provides you with storage requirements and the instructions on how to allocate storage for physical (DACs, PowerVaults, Unity) and virtual storage devices for NetWitness Platform. It also includes the following topics.

- Detect Encryption on Existing PowerVault
- Migrate Data to Another Device

Refer to the following Hardware Setup Guides for information on how to connect these device to NetWitness Platform Core and Hybrid physical hosts:

- PowerVault (MD 2412) Setup Guide (see the "Enclosure Options" section of "Hardware Description")
- [NetWitness Community](#).
- PowerVault (MD 1400) Setup Guide (see the "Enclosure Options" section of "Hardware Description")
- [NetWitness Community](#).
- 60-Drive DAC Setup Guide - [NetWitness Community](#).
- 15-Drive DAC Setup Guide - [NetWitness Community](#).

Storage Requirements

This section contains all the storage requirements needed to successfully attach storage to your NetWitness Platform deployment host systems. It contains the required drive types, appropriate volumes, and performance IOPS that are needed.

IMPORTANT: For Series 7, Series 6E, MD2412 and MD1400, all drives are all self-encrypting drives. However, encryption should be turned on explicitly by the customer, so that they can own and manage the encryption key.

Drive Specifications

General specifications for core NetWitness Platform Hosts are:

- IO size 490/Dec
- Response/Latency < 20ms
- Decoder 10/90 read/write (low random I/O)
- Concentrator 50/50 read/write (high random I/O)

RAID Group	Suitable Volumes
NL-SAS or 10K SAS	All Packet Decoder volumes All Log Decoder volumes All Archiver volumes Concentrator meta volume
SSD	Concentrator index volume

Required NetWitness Platform Storage Volumes

Service Volume Names

Service	Volume Name	File Systems Created
Network Decoder	decoder	packetdb
Network Decoder	decodersmall	decoder root, index, sessiondb, metadb
Log Decoder	logdecoder	packetdb
Log Decoder	logdecodersmall	logdecoder root, index, sessiondb, metadb
Concentrator	concentrator	concentrator root, metadb, sessiondb

Service	Volume Name	File Systems Created
Concentrator	index	index
Archiver	archiver	database

Volume Sizing

The volume sizes below are automatically created when using the NetWitness Platform storage tool, described in [Configure Storage Using the REST API](#).

Volume	Filesystem	Mount Point	Size
decodersmall	decoroot	/var/netwitness/decoder	10 GB
decodersmall	index	/var/netwitness/decoder/index	30 GB
decodersmall	sessiondb	/var/netwitness/decoder/sessiondb	600 GB
decodersmall	metadb	/var/netwitness/decoder/metadb	100% of free space on decodersmall volume
decoder	packetdb	/var/netwitness/decoder/packetdb	100% of free space on decoder volume
logdecodersmall	decoroot	/var/netwitness/logdecoder	10 GB
logdecodersmall	index	/var/netwitness/logdecoder/index	30 GB
logdecodersmall	sessiondb	/var/netwitness/logdecoder/sessiondb	600 GB
logdecodersmall	metadb	/var/netwitness/logdecoder/metadb	100% of free space on logdecodersmall volume
logdecoder	packetdb	/var/netwitness/logdecoder/packetdb	100% of free space on logdecoder volume
concentrator	root	/var/netwitness/concentrator	30 GB
concentrator	sessiondb	/var/netwitness/concentrator/sessiondb	10% of free space on concentrator volume
concentrator	metadb	/var/netwitness/concentrator/metadb	100% of free space on concentrator volume
index	index	/var/netwitness/concentrator/index	100% of free space on index volume

Volume	Filesystem	Mount Point	Size
archiver	database	/var/netwitness/archiver/database	100% of free space on archiver volume

Performance Recommendations

NetWitness recommends that Packet and Log Decoders receive two LUNs or Block Devices, one for Packet data, the other for all other databases. This allows you to segregate the high-bandwidth Packet Database from the other databases so they do not compete for I/O bandwidth with other activity.

Concentrators require a separate SSD-based index volume for best performance. You must house this index volume on a different RAID group than the Concentrator Meta database volume, which you can stored on NL-SAS. Archivers can use a single large NL-SAS storage volume per appliance.

Input/Output Operations Per Second

The following table lists the IOPS requirements for the Decoder and Concentrator hosts.

Logs	Log Decoder	Concentrator
10K EPS	400	8,000
20K EPS	550	10,300
25K EPS	1,200	10,800

Packets	Network Decoder	Concentrator
1Gbps	600	6,050
2 Gbps	950	8,300
4 Gbps	1,650	12,800
6 Gbps	2,400	17,300
8 Gbps	3,200	21,800

General Description of How NetWitness Platform Hosts Store Data

For information about how NetWitness Platform hosts store data, see [Appendix A. How NetWitness Platform Hosts Store Data](#).

Enable Encryption on Series 6E or Series 7 Hosts

All drives in Series6E/Series7 and any configured external storage (MD1400 or MD2412) are self-encrypting (SED). However, encryption has to be turned on explicitly by the customer, so they can generate and manage their own encryption keys. NetWitness is not responsible for key generation and management and will not have access to customer's disk encryption keys. If the customer chooses to enable encryption, NetWitness recommends enabling it after completion of external storage configuration (if any).

Refer to [Appendix B. Encrypt a Series 6E or Series 7 Core or Hybrid Host \(encryptSedVd.py\)](#) for details on enabling encryption.

Configure Drive Pack(s)

Note: The terms 'Meta Disk Kit' and 'Meta Drive Pack' mean the same and are interchangeable.

Benefits of Series 6/6E and Series 7 Drive Pack

Series 6/6E Hosts

You can add additional drives to the Series 6 or 6E appliances to accommodate various use cases. These drives provide the capability for the decoder meta or concentrator index volumes to reside on the appliance. Each Meta Disk Kit has 3 drives. A maximum of 2 Meta Disk Kits can be installed on series 6/6E appliances. The Meta cache or index size determines the number of Meta Disk Kits. A standard Series 6/6E appliance has 4 drives in slots 0,1,2 and 3. The remaining slots, 4 through 9 are empty (highlighted in red in the Series 6/6E Disk Layout image below). These slots are used to install the Meta Disk Kit(s).


Series 7 Hosts

You can add additional drives to the Series 7 to accommodate various use cases. These drives provide the capability for the decoder meta to reside on the appliance. Each Meta Disk Kit has 3 drives. A maximum of 2 Meta Disk Kits can be installed on series 7 appliances. The Meta cache determines the number of Meta Disk Kits. A standard Series 7 appliance has 4 drives in slots 0,1,2 and 3. The remaining slots, 4 through 9 are empty (highlighted in red in the Series 7 Disk Layout image below). These slots are used to install the Meta Disk Kit(s).

NW Host - Meta Disk Kit support Matrix

NW Host	NW Service	External Storage	Meta Disk Kit Applicable (Yes/No)
Series 6/6E	Decoder	MD1400	Yes
	Concentrator	MD1400	Yes
Series 7	Decoder	MD2412	Yes
	Concentrator	MD2412	No

Series 6/6E and Series 7 Disk Layout



- **Maximize PowerVault Storage Capacity** - Traditionally, PowerVault storage allocates a volume for the Decoder metadata. This reduces the usable storage on the PowerVault. Drive Packs reduce this issue by providing 20TB of extra usable PV storage.
- **Reduces Cost for Meta Only Use Case** - In Meta Data Only deployments, the use of drive packs can help remove the requirement for the use of a single PowerVault on these Decoders..

- **Enable existing deployments to utilize compression options** - For existing deployments, an SSD index drive pack is required if you need to enable compression. When compressing the packetdb (Decoders) and metadb (Concentrators), additional indexing is needed to support compression of those databases.
- **Provides capability for expanding meta keys and associated indexing** - The index storage needs are scaled based on the NetWitness Platform deployment retention requirements. If additional meta keys are enabled and indexed, it may impact index retention.

Decoder Meta Use Cases

Note: Applicable to Series 6/6E and Series 7 Decoder.

- Meta-Only
- Maximize Power Vault Storage

Three or more 2.4TB 10K SAS SED drives can be added to a Decoder for the decodersmall or logdecodersmall volumes. These volumes are used to store the meta cache on the Decoders.

Both the Log Decoders and Network Decoders parse out meta data from the raw captured traffic. The meta data is then aggregated to a Concentrator for indexing.

The host requires storage to store a cache for the meta extracted during the data capture for Concentrator aggregation. The meta cache on a Decoder is generally fixed in size, but you can expand it to support additional cache to avoid possible connectivity loss between the Decoder and the corresponding Concentrator.

Typically, the decodersmall or logdecodersmall volumes are stored on the first three drives of the first PowerVault enclosure and, in 10G configurations, second PowerVault enclosure. By utilizing the drive pack option, these three drives can instead be used for the packetdb (maximizing Power Vault storage).



For meta-only scenarios, the decodersmall volume would be stored on the drive pack, therefore eliminating the need for a Power Vault.

Note: The steps noted in the below (Sample Storage Configuration Scenarios for Meta Disk Kit(s) on Series 6/6E Decoder) also applies to Series 7 Meta Disk Kit configuration for both Meta-Only and Maximize PowerVault Storage (MD2412).

Sample Storage Configuration Scenarios for Meta Disk Kit(s) on Series 6/6E Decoder

This section describes how to configure a Meta Disk kit on a Series 6/6E decoder as meta-only and maximized MD1400 PowerVault storage capacity.

Meta-Only (No Externally Attached Storage)

Install and configure one Meta Disk Kit (3 SED Drives configured as RAID 5) on a S6 Core Appliance that has been orchestrated as a Decoder with one attached and unconfigured MD1400 PowerVault.


Note:

1. The configuration for Log Decoder is similar to Network Decoder. Substitute the service and volume names for Log Decoder that would normally be associated with Network Decoder.
2. One Meta Disk pack is configured as RAID5 (3 drives) and two Meta Disk packs (6 drives) are configured as RAID6.
3. When configuring two Meta Disk Packs, the disks are installed in slots 4 through 9 and when adding second Meta Disk Pack, the disks are installed in slots 7 through 9. Refer [Series 6/6E and Series 7 Disk Layout](#) for slot details.

On the Series 6 (Dell R640) appliance, the Meta Drive Pack disks are installed in slots 4, 5 and 6. The virtual drive configuration requires identifying the controller ID and Enclosure ID (EID). On Series 6 appliance, the controller ID and Enclosure IDs are 0 and 64, respectively. The nwraidtool.py script, that is installed on every orchestrated server, can help to confirm these ID numbers.


1. Install the three drives in the Drive Pack in slots 4, 5 and 6. Refer [Series 6/6E and Series 7 Disk Layout](#) for slot details.
2. Identify the existing block devices using 'raidList' property. Login to **NW UI > Hosts > Select the Decoder Host > Actions > View > Explore > deviceAppliance > appliance (Right Click to access properties) > raidList** and click **Send**.

The existing devices on Controller 0, Enclosure 64 are highlighted in Yellow. The installed Drive Pack in slots 4,5 and 6 is highlighted in Red.




3. Access **NW UI->Hosts->Select the Decoder Host->Actions->View->Explore->deviceAppliance->appliance (Right Click to access properties)->raidNew**, specify the controller, enclosure, and scheme, click **Send**.

Note: For decoder the scheme is decoder-metakit, for logdecoder the scheme is logdecoder-metakit




- Identify the new device (highlighted in Yellow) using raidList command. In this case it is 'sdc'. Access NW UI->Hosts->Select the Decoder Host->Actions->View->Explore->deviceAppliance->Appliance (Right Click to access properties)->raidList->Click on Send.



- Execute the **partNew** command by selecting it from the Properties drop-down to create the decodersmall partition (decoder dir, index, metadb, sessiondb) with the following parameters.

name=sdc service=decoder volume=decodersmall commit=1

Note: For logdecoder, use the command name=sdc service=logdecoder volume=logdecodersmall commit=1




```
[root@s6Core ~]# df -hP
Filesystem                                Size  Used Avail Use% Mounted on
devtmpfs                                  63G   0    63G   0% /dev
tmpfs                                     63G  100K   63G   1% /dev/shm
tmpfs                                     63G   11M   63G   1% /run
tmpfs                                     63G   0    63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root           30G   4.1G   26G  14% /
/dev/mapper/netwitness_vg00-nwhome         2.7T  516M   2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog         10G   70M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome        10G   33M   10G   1% /home
/dev/sda1                                1014M   91M   924M   9% /boot
tmpfs                                      13G   0    13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot          10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index            30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb        600G   34M   600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb           3.8T   34M   3.8T   1% /var/netwitness/decoder/metadb
[root@s6Core ~]#
```

6. Allocate the decodersmall to Decoder service using 'srvAlloc'.

```
service=decoder volume=decodersmall commit=1
```

Note:

- For logdecoder, use the command `service=logdecoder volume=logdecodersmall commit=1`
- If a second Drive Pack is being configured, the volume for decoder would be 'decodersmall0'. For logdecoder, it is 'logdecodersmall0'.



Maximize PowerVault Storage Capacity

Installing and configuring a Drive pack on the Decoder appliance as decodersmall frees the first three drives on the attached Power Vault for storing additional packet data.

Best Practice Recommendation: When 1 to 4 PowerVaults are configured, one (1) Meta Disk kit is recommended. When 5 to 8 PowerVaults are configured, two (2) Meta Disk kits are recommended.


Note: This sample scenario assumes a S6/S6E appliance orchestrated as a Decoder and connected to a PowerVault (but not configured as storage).

1. Follow all the steps in [Meta Only: No Externally Attached Storage](#) section to configure one meta disk pack on decoder as the decodersmall partition.
2. Use the **Explorer** view (Login into UI, select **Hosts->Services->Decoder ->Actions->View->Explorer->deviceAppliance->appliance(right click)->properties**) -> **raidList** -> **Click on Send**, to identify and confirm the Controller Number, Enclosure Number, In Use, Drives, Devices, Drive Count, Size, and Vendor (highlighted in yellow).

You should see the following information.


In Use: FALSE

Devices: <EMPTY>




3. Create the raid using 'raidNew' from the properties drop-down (passing the below parameters) on the attached PowerVault. Identify the controller and enclosure from 'raidList'. Encryption can be turned on after configuring storage using the steps listed in [Appendix B. Encrypt a Series 6E Core or Hybrid Host](#).


```
controller=1 enclosure=250 scheme=packet-expansion commit=1
```



- Identify the block device created in the earlier step. Use 'raidList' to retrieve all the block devices. In this case, it is 'sdd' (highlighted in Yellow).



- Make partitions on the block device ('sdd') using 'partNew'.
 name=sdd service=decoder volume=decoder commit=1
 For logdecoder, use the following command.
 name=sdd service=logdecoder volume=logdecoder commit=1




```
[root@s6Core ~]# df -hP
Filesystem                                Size  Used Avail Use% Mounted on
devtmpfs                                  63G   0    63G   0% /dev
tmpfs                                     63G  80K   63G   1% /dev/shm
tmpfs                                     63G  11M   63G   1% /run
tmpfs                                     63G   0    63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root           30G  4.1G   26G  14% /
/dev/mapper/netwitness_vg00-nwhome        2.7T  515M  2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog        10G   67M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome       10G   33M   10G   1% /home
/dev/sda1                                1014M   91M  924M   9% /boot
tmpfs                                     13G   0    13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot         10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index           30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb       600G   34M  600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb          3.8T   34M  3.8T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb             107T   38M  107T   1% /var/netwitness/decoder/packetdb
[root@s6Core ~]#
```

6. Allocate the decoder volume to Decoder service using 'srvAlloc'.

```
service=decoder volume=decoder commit=1
```

For logdecoder, use the following command.

```
service=logdecoder volume=logdecoder commit=1
```



Concentrator Index Use Cases

Caution: The Drive Pack(s) for Concentrator Index are supported only on Series 6/6E hosts. These use cases are NOT supported on Series 7 Concentrator hosts. On Series 7 Concentrator hosts, the slots 4 through 10 on the host are reserved for Index database along with MD2412.

The following use cases are supported on Series 6/6E Concentrator Hosts configured with MD1400 Power Vaults.

- Support Additional Meta-Key Indexing
- Capability to Enable compression for Existing Deployments

Three or more 3.84 TB or 3.2 TB or 7.68 TB SSD SED drives can be added to a Concentrator to increase the index volume. The index storage needs are scaled based on the NetWitness Platform deployment retention requirements. If additional meta keys are enabled and indexed, it may impact index retention.

For existing deployments, an SSD index drive pack is required if you need to enable compression. When compressing the packetdb (Decoders) and metadb (Concentrators), additional indexing is needed to support compression of those databases.

Sample Storage Configuration for Concentrator Index with One Meta Disk Kit (Three SSD's)


This section describes the steps to configure a Meta Disk Kit (Three SSD's) installed on a Series 6/6E appliance orchestrated as Concentrator and **configured with One Power Vault**.

Note:

- Concentrator index usecase supports a maximum of Two Meta Disk Kits. A single Meta Disk pack is configured as RAID5 (3 drives) or two Meta Disk packs (6 drives) are configured as RAID6.
- Even though each Meta Disk Kit consist of 3 SSD's, In rare circumstances when there is no need for very large index expansion and to control costs, users may choose to install only two SSD's (in slots 4 and 5 and configured as RAID1) instead of 3 SSD's (One Meta Disk Kit).
- When configuring two Meta Disk Packs, the disks are installed in slots 4 through 9 and when adding second Meta Disk Pack, the disks are installed in slots 7 though 9. Refer to [Series 6/6E and Series 7 Disk Layout](#) for slot details.

On the Series 6 (Dell R640) appliance, the Meta Drive Pack disks are installed in slots 4, 5 and 6. The virtual drive configuration requires identifying the controller ID and Enclosure ID (EID). On Series 6 appliance, the controller ID and Enclosure IDs are 0 and 64. However, the `nwraidtool.py` script that is installed on every server can help to confirm these ID numbers.

1. Install the three SSDs in the Meta Disk kit in slots 4, 5 and 6 on the Concentrator Appliance. Refer to [Series 6/6E and Series 7 Disk Layout](#) for slot details.
2. Identify the existing block devices using **raidList** property. Login to NW UI > **Hosts** > **Select the Concentrator Host** > **Actions** > **View** > **Explore** > **deviceAppliance** > **appliance (Right Click to access properties)** > **raidList** and click **Send**.



- SSH to the Concentrator appliance or use iDRAC to connect to the console. Use **perccli** to create the virtual drive with installed Meta Pack kit. Create the Virtual Drive or Drive Group (DG) on the internal controller using the disks in slot 4 through 6 using the below command.


```
/opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6
strip=128
```

Note:

- To configure two Meta Disk Packs, use the following command.
/opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid6 drives=64:4-9
strip=128
- To add a second Meta Disk Pack (First Meta Disk Pack already configured in slots 4,5 and 6), use the following command.
/opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:7-9
strip=128
- To configure only 2 SSD's (instead of 3 SSD's) from a Meta Disk Pack, use the following command. /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid1
drives=64:4-5 strip=128


```
[root@S6CoreConc ~]# /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid5 drives=64:4-6 strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 0
Status = Success
Description = Add VD Succeeded.
```

- Identify the new device (highlighted in Yellow) using **raidList** command. In this case it is 'sde'. Login to NW UI->Hosts->Select the Concentrator Host->Actions->View->Explore->deviceAppliance->Appliance (Right Click to access properties)->raidList->Click on Send.



- Execute the **partNew** command with the below parameters to create the new index partition on the block device (sde) created in earlier step:


```
name=sde service=concentrator volume=index commit=1
```



- The concentrator appliance is already configured with a Power Vault (an 'index' volume is created and configured along with other volumes on the Power Vault) before adding the Meta Disk Kit. The new 'index' volume created in the earlier step is named as 'index0'.

Allocate the new index volume (index0) using `srvAlloc` property to concentrator service using the below parameters:

```
service=concentrator volume=index0 commit=1
```



```
[root@Concentrator ~]# df -hP
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  63G         0   63G   0% /dev
tmpfs                     63G        60K   63G   1% /dev/shm
tmpfs                     63G       11M   63G   1% /run
tmpfs                     63G         0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G       3.7G   27G  13% /
/dev/sda1                 1014M       91M   924M   9% /boot
/dev/mapper/netwitness_vg00-nwhome 2.7T     561M   2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome 10G       33M    10G   1% /home
/dev/mapper/netwitness_vg00-varlog 10G       48M    10G   1% /var/log
tmpfs                     13G         0   13G   0% /run/user/0
/dev/mapper/concentrator-root 30G       34M    30G   1% /var/netwitness/concentrator
/dev/mapper/concentrator-sessiondb 7.5T     34M   7.5T   1% /var/netwitness/concentrator/sessiondb
/dev/mapper/concentrator-metadb 68T       36M   68T   1% /var/netwitness/concentrator/metadb
/dev/mapper/index-index 7.0T       34M   7.0T   1% /var/netwitness/concentrator/index
/dev/mapper/index0-index 7.0T       34M   7.0T   1% /var/netwitness/concentrator/index0
[root@Concentrator ~]#
```

Enable Security on SED Capable Drives

To enable Security on the SED Capable Drive Group on Series 6E (PERC H740 Mini, PERC H750, PERC H840 Adaptors) or Series 7 Hosts (PERC H965i, PERC H965e Adaptors), see [Enable Encryption on Series 6E or Series 7 Hosts](#).

Prepare Virtual or Cloud Storage

This section describes how to set up virtual or cloud storage for the following types of component hosts:

- [Decoder, Log Decoder, Concentrator, Archiver](#)
- [NW Server, ESA Primary, ESA Secondary and Malware Analysis](#)
- [Log Collector](#)
- [Endpoint Log Hybrid](#)
- [Additional Endpoint Log Hybrid Partitions](#)
- [UEBA](#)

Decoder, Log Decoder, Concentrator, Archiver

Virtual or Cloud NetWitness hosts for Decoders, Log Decoders, Concentrators, and Archivers need block storage attached. Make sure that the allocated storage meets all of the storage requirements. Specifically, make sure that the required storage volumes are created (see "Required NetWitness Platform Storage Volumes" in [Storage Requirements](#)), and:

- At least two Block Devices are created for Decoders (meta /session and packet volumes)
- At least two block devices are created for Concentrators (index and meta volumes)
- Ensure that block devices can meet the minimum IOPS for expected ingestion rates

Attach the allocated storage to the NetWitness host by following the hosting platforms native procedure.

- VmWare – Vsphere Console (add disk to VM)
- Hyper-V – Manager Console (add disk to VM)
- Azure – Add Managed Disks to virtual instance
- AWS – Add EBS Storage to virtual instance
- Google Cloud Platform (GCP) - Add storage to virtual instance

After the storage is attached to the virtual host, proceed to "Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes" in [Configure Storage Using the REST API](#).

NW Server, ESA Primary, ESA Secondary and Malware Analysis

For an extension of /var/netwitness/ partition, attach an external volume.

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `/dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`

3. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
4. `xfs_growfs /dev/netwitness_vg00/nwhome`

NetWitness recommends the following partition definitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the Cloud Provider Block Storage setup (storage) tables.

Log Collector

For an extension of `/var/netwitness/` partition, attach an external volume

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate <pv_name>` (for example, `pv_name` is `dev/sdc`)
2. `vgextend netwitness_vg00 /dev/sdc`
3. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
4. `xfs_growfs /dev/netwitness_vg00/nwhome`

NetWitness recommends the following partition definitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the Cloud Provider Block Storage setup (storage) tables.

Endpoint Log Hybrid

The total disk size required depends on the data retention period. You can use the below per day disk usage indicative values to calculate the required disk size for your deployment. For example, to retain 30 days of data, multiply the below per day disk usage values with 30.

The following table provides disk usage for one full scan. The full scan disk usage values are based on the below event count:

- Files count -1100
- Processes count -100
- Dlls count - 500
- Drivers count -150
- Services count - 500
- Tasks count -100

Endpoint Log Hybrid(50K Advance Agents - Disk usage per full scan)

	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	220 GB	12 GB	5 GB	NA	237 GB
Concentrator	230 GB	NA	5 GB	6 GB	241 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 30 GB (Subsequent per scan increase)

The following tables provide per day disk usage for tracking data. The total tracking events per agent per day is 29000.

Endpoint Log Hybrid (50K Advance Agents - Tracking data without Expanded Network Visibility)

	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	1500 GB	140 GB	46 GB	NA	1,686 GB
Concentrator	1600 GB	NA	46 GB	30 GB	1,676 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 1.5 GB (Tracking data per day increase)

The following tables provide per day disk usage for tracking data. Total tracking events per agent per day is 33000

Endpoint Log Hybrid (50K Advance Agents - Tracking data with Expanded Network Visibility)

	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	1800 GB	152 GB	55 GB	NA	2007 GB
Concentrator	1900 GB	NA	55 GB	36 GB	1991 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 1.5 GB (Tracking data per day increase)

The following table provides per day disk usage for insight agents. The total tracking data per agent per day is 10800 plus 1 full scan daily.

Endpoint Log Hybrid (50K Insights Agents with Expanded Network Visibility)

	MetaDB	PacketDB	SessionDB	Index	Total
Log Decoder	500 GB	52 GB	18 GB	NA	570 GB

Endpoint Log Hybrid (50K Insights Agents with Expanded Network Visibility)

Concentrator	600 GB	NA	18 GB	13 GB	631 GB
MongoDB	NA	NA	NA	NA	35 GB (First full scan) 30 GB (Subsequent per scan increase)

The following table provides Endpoint Agents sizing based on the feature.

Feature	Description	Agent or Endpoint Server
Endpoint Only	Only scan and tracking data	Maximum 50K Endpoint Agents only
Windows Logs Only	Only Windows Logs from agents. Assuming 20K events per second supported by Hybrid.	Maximum 20K Agents: Generates 20K log events per second
File Collection Only	Only File Collection from agents. Assuming 20K events per second supported by Hybrid	Maximum 20K Agents : Generates 20K log events per second
Endpoint and Windows Logs	Event per second per agent (For Windows Logs) 1 event sent by 1 agent every second (For Tracking Events) 0.4 event sent by 1 agent every second 20K events per second supported by Hybrid Note: Total agents should be calculated as below: Hybrid events per second/ (Windows Logs Endpoint Server of 1 agent + Tracking Event Endpoint Server for 1 agent) For example, $20000 / (1.0 + 0.4)$	Maximum 15K (approximately) Agents: - Generates 15K (approximately) Windows log events Plus - Generates 15K (approximately) Agents EDR data

Feature	Description	Agent or Endpoint Server
Endpoint, Windows Logs and File Collection	Event per second per agent: • (For Windows Logs) 1 event sent by 1 agent every second • (For Tracking Events) 0.4 event sent by 1 agent every second • (For File Collection) 1 event sent by 1 agent every second • 20, 000 events per second supported by Hybrid	Maximum 10K (approximately) Agents: • Generates 10K (approximately) Windows log events Plus • Generates 10K (approximately) Endpoint Agents data Plus • Generates 10K (approximately) Agents File Collection data

Note: Total agents should be calculated as below:
Hybrid events per second/
(Windows Logs Endpoint Server of 1 agent + Tracking Event Endpoint Server for 1 agent + File Collection)
For example, $20000 / (1.0 + 1.0 + 0.4)$

Extending File Systems

For Endpoint Server, attach external disk for extension of `/var/netwitness/` partition, create an external disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see "Task 1. Add New Disk" in the *Virtual Hosts Installation Guide for NetWitness Platform*. Go to the [NetWitness All Versions Documents](#) page and find NetWitness Platform guides to troubleshoot issues.
2. Execute `lsblk` and get the physical volume name
3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc`
4. `vgextend netwitness_vg00 /dev/sdc`
5. `lvresize --resizefs --extents +100%FREE /dev/netwitness_vg00/nwhome`
6. `xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

NetWitness recommended partition for Endpoint Server (can be changed based on the retention days).

LVM	Folder	Size	Disk Type
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	6TB	HDD

For Mongo DB, attach external disk for extension of `/var/netwitness/mongo` partition, create an external disk with suffix as `nwhome`.

Follow these steps:

1. Ensure you have added a new disk. For more information, see "Task 1. Add New Disk" in the *Virtual Hosts Installation Guide*. Go to the [NetWitness All Versions Documents](#) page and find NetWitness Platform guides to troubleshoot issues.
 2. Execute `lsblk` and get the physical volume name
 3. `pvcreate <pv_name>` suppose the PV name is `/dev/sdc1`
 4. `vgextend hybrid /dev/sdc1`
 5. `lvresize --resizefs --extents +100%FREE /dev/hybrid-vmng`
 6. `xfs_growfs /dev/mapper/hybrid-vmng`
- NetWitness recommended partition for Mongo DB (Can be changed based on the retention days).
Minimum recommended size for `var/netwitness` is 500 GB.

LVM	Folder	Size	Disk Type
<code>/dev/hybrid-vmng</code>	<code>/var/netwitness/mongo</code>	6TB	HDD

Additional Endpoint Log Hybrid Partitions

The following partition should be on the volume group endpoint and should be in a single RAID 0 array.

Folder	LVM	Volume Group
<code>/var/netwitness/mongo</code>	hybrid-mongo	endpoint
<code>/var/netwitness/concentrator</code>	concentrator-concroot	endpoint
<code>/var/netwitness/concentrator/index</code>	hybrid-concindex	endpoint
<code>/var/netwitness/logdecoder</code>	hybrid-ldecroot	endpoint

Run `lsblk` to get the physical volume name and run the following commands:

1. `pvcreate /dev/md0`
2. `vgcreate -s 32 endpoint /dev/md0`
3. `lvcreate -L <disk_size> -n <lv_name> endpoint`
4. `mkfs.xfs /dev/ endpoint /<lv_name>`
5. Repeat the above steps for all the LVMs mentioned.

NetWitness recommends the following partitions. However, you can change these values based on the retention days.

LVM	Folder	Block Storage
<code>/dev/netwitness_vg00/nwhome</code>	<code>/var/netwitness/</code>	Refer to the Cloud Provider Block Storage setup (storage) tables.

LVM	Folder	Block Storage
/dev/endpoint/hybridmongo	/var/netwitness/mongo	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/concentratorconcroot	/var/netwitness/concentrator	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/hybridconcinde	/var/netwitness/concentrator/index	Refer to the Cloud Provider Block Storage setup (storage) tables.
/dev/endpoint/hybridldecroot	/var/netwitness/logdecoder	Refer to the Cloud Provider Block Storage setup (storage) tables.

UEBA

The following procedure attaches an external disk and extends the `/var/netwitness/` partition. You must use `nwhome` as the external disk suffix. This procedure illustrates how to add a 2TB disk.

Note: `/var/netwitness` is the only partition that can reside on this volume.

1. List the physical volume name.
`lsblk (for example, dev/mapper/sdc)`
2. Extend the `/var/netwitness/` partition.
`pvcreate <pv_name>where pv_name is dev/mapper/sdc`
`vgextend netwitness_vg00 /dev/mapper/sdc`
`lvresize --resizefs --extents +100%FREE /dev/mapper/netwitness_vg00/nwhome`
`xfs_growfs /dev/mapper/netwitness_vg00-nwhome`

This partition is the NetWitness recommended partition for UEBA. You can change it based on retention days.

Configure Storage Using the REST API

you can use the REST API for all storage configuration operations. For information about how to use the REST API, see the *RESTful API User Guide*. Go to the [NetWitness All Versions Documents](#) page and find NetWitness Platform guides to troubleshoot issues.

REST API Storage Configuration Commands

Each of the commands listed below has built-in help that describes their function and usage. If you are using the REST interface, select the command from the drop-down menu to see the help text. For examples of REST API storage configuration commands, see [Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs](#).

Commands for Direct-Attached RAID Volumes

- `raidList` : List the RAID controllers and direct-attach enclosures that are present on this host.
- `raidNew` : Allocate direct-attached enclosures to block devices.

Commands for Allocating Block Devices as Storage

- `devlist` : List available block devices on the host.
- `partNew` : Allocate partitions on a block device and create volume groups.
- `vgs` : Summarize how block devices are organized into volume groups.

Commands for Allocating Storage to Services

- `srvList` : List services on the host and their allocated storage paths.
- `srvAlloc` : Allocate a volume group to a service.
- `srvFree` : Remove a volume group from a service.
- `multipath-II` : To verify if SAN devices are attached.

Command to Reconfigure Services to Detect and Use All of the New Storage

- `reconfig` - After configuring new storage, detect and use new storage on the associated service and database.

Storage Configuration Tasks

Task 1 - Attach storage to the host and access the REST API storage configuration commands.

Task 2 - (Conditional) Configure RAID if necessary.

Task 3 - Allocate block devices to partitions, volume groups, and logical volumes.

Task 4 - Allocate volume groups to NetWitness services.

Task 5 - Reconfigure services and databases to detect and appropriately use new storage.

Task 1 - Attach Storage to the Host and Access the REST API Storage Commands

Complete the following steps to attach an external storage device to a host and access the storage configuration commands available through the REST API.

1. Attach the storage and make it available to this host.
 - To attach PV storage, refer to the [PowerVault \(Dell MD1400\) Setup Guide and PowerVault \(Dell MD2412\) Setup Guide](#).
 - For third-party storage, create the RAID groups to match the volumes listed in [Storage Requirements](#)
2. There are two ways that you can access the REST API storage commands: from a Browser, or from the **Services > Explore** view from the User Interface.


Note: Once you have accessed the REST API, the steps that you perform are the same, no matter which method you used to access it.

- From a Browser.
 - a. Open a Browser and specify the ip-address of the host with port **50106**.
The following example is the Decoder, but you need to use port 50106 for any host hardware for which you are configuring storage using the REST API.
`https://<decoder-ip-address>:50106`


- b. Log in with the `admin` account credentials.
The following REST API menu is displayed.



- c. Click on the **(*)** next to **appliance** to access the REST command set.
The **Properties for /appliance** dialog is displayed under the initial REST menu. The **Output (or command manual help)** section describes the commands that the REST API can send to the device, their usage, and their parameters.



- From the User Interface.
 - a. In the **NetWitness** menu, go to  (Admin) > **SERVICES**.
 - b. Select the service (for example, a Concentrator).
 - c. Under  (actions), select **View > Explore**.
 - d. Navigate to **deviceappliance/appliance**, right click, and click **Properties**.



You can now access the storage commands from the **Properties** dialog.

3. Proceed to:

- [Task 2](#) if you need to configure RAID for PowerVault or DACs.
- [Task 3](#) if you do not need to configure RAID and already have a block device available.

Task 2 - (Conditional) RAID Configuration for PowerVault and DACs

NetWitness Platform hardware uses direct-attached SAS drives for storage. These drives are housed in a SAS enclosure. SAS enclosures are shelves of drives attached to the NetWitness node by a cable connected to the SAS host bus adapter.

SAS enclosures are also known as other names, such as "DAC" (Direct-Attached Capacity), or "JBOD" (Jumbo Box of Disks), or "Dell PowerVault".

NetWitness Platform utilizes Dell PERC SAS host bus adapters. NetWitness Platform devices typically include two SAS host bus adapters. One is used for controller drives that are internal to the NetWitness Node, and another is used for controlling drives attached to the SAS enclosures. The internal controller and drives are configured when the node is built, but the external SAS enclosures are not. You execute the `raidList` and `raidNew` commands to identify and configure the external SAS enclosures.

These commands work with the following SAS enclosure types:

- EMC ESAS 15-drive enclosures
- EMC ESAS 60-drive enclosures
- Dell PowerVault 12-drive enclosures
- Dell PowerVault 8-drive enclosures

Note: EMC 60-drive enclosures are logically organized as four separate 15-drive sub-enclosures. They behave as if there are four 15-drive enclosures, each of which can be configured independently.

The `raidList` and `raidNew` commands operate on entire enclosures. Execute `raidList` to identify the enclosures. execute `raidNew` to configure an enclosure to perform one of the pre-determined roles within a NetWitness Platform node.

After you attach storage to the host and access the REST API storage commands, complete the following steps to create RAID if required.

1. Execute the `raidList` command to identify the controllers and enclosures that are attached to the system.

In the following example, Controller 1 does not display any block devices. This indicates the array is not configured.

Properties for /appliance
raidList Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0, Enclosure 32
  Vendor: DP
  Model: BP13G+EXP
  In Use: true
  Drives: 931.511 GB x 2
          1.818 TB x 2
  Devices: sda
           sdb

Controller 1, Enclosure 82
  Vendor: DELL
  Model: MD1400
  In Use: false
  Drives: 10.691 TB x 12
  Devices:

Controller 1, Enclosure 13
  Vendor: DELL
  Model: MD1400
  In Use: false
  Drives: 10.691 TB x 12
  Devices:
```

2. Select a RAID layout scheme for the Enclosure.

The following tables list the PV to Supported Hosts Mapping.

Type	SKU	Specification	Supported Hosts
High Density	NW-PV-A	Dell Storage MD1400 12 x 12 TB NL-SAS SED	Decoder, LogDecoder, Archiver, Log Hybrid, Network Hybrid, Endpoint Log Hybrid
High Density	NW-PV-B	Dell Storage MD 1400 8 x 12TB NL-SAS SED	Decoder, LogDecoder, Archiver, Log Hybrid, Network Hybrid, Endpoint Log Hybrid
High Performance	NW-PV-C	Dell Storage MD 1400 6 x 12TB NL-SAS SED, 2 x 3.8TB SSD SED	Concentrator

Type	SKU	Specification	Supported Hosts
High Performance	NW-PV-D	Dell Storage MD 1400 9 x 12TB NL-SAS SED, 3 x 3.8TB SSD SED	Concentrator

Type	SKU	Specification	Supported Hosts
High Density	192TB (NW-PV-A-N)	Dell Storage MD1400 12 x 16 TB NL-SAS SED	Decoder, LogDecoder, Archiver, Log Hybrid, Network Hybrid, Endpoint Log Hybrid
High Density	128TV (NW-PV-B-N)	Dell Storage MD 1400 8 x 16TB NL-SAS SED	Decoder, LogDecoder, Archiver, Log Hybrid, Network Hybrid, Endpoint Log Hybrid
High Performance	103TB (NW-PV-C-N)	Dell Storage MD 1400 6 x 16TB NL-SAS SED, 2 x 3.8TB SSD SED	Concentrator
High Performance	155TB (NW-PV-D-N)	Dell Storage MD 1400 9 x 16TB NL-SAS SED, 3 x 3.8TB SSD SED	Concentrator
High Density	S7-PV-HD-192	Dell Storage MD2412 12 x 16 TB SAS SED	Decoder, LogDecoder, Archiver, Log Hybrid, Network Hybrid, Endpoint Log Hybrid
High Density	S7-PV-HD-192 AND S7-PE-CON-SDF-3 OR S7-PV-HD-192 AND S7-PE-CON-SDF-7	Dell Storage MD2412 12 x 16 TB SAS SED One or more SSD Drive Pack (s)	Concentrator

The following tables show you the supported allocation schemes.

Note:

- On a Series 6 or Series 7 Network Decoder or newer with multiple PowerVault storage trays, use the decoder-hotspare RAID scheme for the first enclosure and the packet-expansion RAID scheme for subsequent enclosures.

Scheme	Drives Required	Allocation
decoder-hotspare	8 or 12 or 15 HDDs	2x Drives in RAID 1 for decoder small, 1 drive as hotspare, all remaining drives in RAID 5 for decoder

Scheme	Drives Required	Allocation
logdecoder-hotspare	8 or 12 or 15 HDDs	Same as decoder-hotspare configuration
archiver	8 or 12 or 15 HDDs	All drives in RAID 6 for archiver or decoder database volume
network-hybrid	8 or 12 or 15 HDDs	3x drives in RAID 5 for meta expansion, all remaining drives in RAID 5 for packet expansion
log-hybrid	8 or 12 or 15 HDDs	Half of the drives in RAID 5 for meta expansion, half the drives in RAID 5 for packet expansion
Note: log-hybrid scheme is also used to configure a PowerVault for Endpoint Log Hybrid host.		
concentrator	2 or more SSDs, 4 or more HDDs	All SSDs in RAID 1 or RAID 5 for index, all HDDs in RAID 6 for meta
packet-expansion	8 or 12 or 15 HDDs	All drives in RAID 6 for decoder volume, no drives allocated for decodersmall
decoder-metakit	1 metakit (3 HDDs) or 2 metakits (6 HDDs)	3x drives in RAID 5 or 6x drives in RAID 6 for meta
logdecoder-metakit	1 metakit (3 HDDs) or 2 metakits (6 HDDs)	3x drives in RAID 5 or 6x drives in RAID 6 for meta
concentrator-metakit	1 metakit (3 SDDs) or 2 metakits (6 SDDs)	3x drives in RAID 5 or 6x drives in RAID 6 for index. If two drive configuration, then 2x drives in RAID 1 for index.
decoder or logdecoder	8 or 12 or 15 HDDs	3x drives in RAID 5 for decodersmall or logdecodersmall, all remaining drives in RAID 5
Note: The decoder and logdecoder scheme has been deprecated in favour of decoder-hotspare and logdecoder-hotspare.		

- After the controller, enclosure, and scheme are identified, execute the `raidNew` command to create RAID Volumes. For example:

```
send /appliance raidNew controller=1 enclosure=82 scheme=decoder-hotspare preferSecure=false
```

Add the `commit=1` parameter to actually execute this operation. Execute the `raidList` command to

list the created block devices.

4. (Optional) Configure SEDs (Self-Encrypting Drives). If the `raidNew` command detects self-encrypting drives and a security key has been set on the controller, the `raidNew` command will attempt to create a secure array. To set a security key on the controller, execute the `raidKey` command. For example:

```
send /appliance raidKey controller=1 key=myPassssphrase keyId=1
```

 - To create a secured (that is, encrypted) array on physical devices attached to a controller with a security key set, specify `preferSecure=true` when using `raidNew`
 - To create an unsecured (that is, unencrypted) array on physical devices attached to a controller with a security key set, specify `preferSecure=false` when using `raidNew`.
5. Go to [Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes](#), after you create RAID volumes.

Task 3 - Allocate Block Devices to Partitions, Volume Groups, and Logical Volumes

The `partNew` command prepares a storage device to use in NetWitness Platform. It performs the following tasks.

- Creates the partition table on the block device.
- Creates the Linux Volume Manager physical device partition.
- Creates a volume group containing the physical device.
- Creates logical volumes in the volume group.
- Creates XFS filesystems on each logical volume.
- Creates `/etc/fstab` entries for each logical volume.
- Mounts each logical volume.

Complete the following steps to allocate block devices to partitions, volume groups, and logical volumes.

1. Run the `devlist` command to locate unused block devices. The following example shows the `devlist` command output.

Output (or command manual help)

```
sda: vendor=DELL model="PERC H730P Mini" size="931 GB" used=1
sdb: vendor=DELL model="PERC H730P Mini" size="1.81 TB" used=1
sdc: vendor=DELL model="PERC H830 Adp" size="21.38 TB" used=1
sdd: vendor=DELL model="PERC H830 Adp" size="85.53 TB" used=1
```

Also, you must provide a name for the service that will be used with the storage, for example, **decoder** for the Network Decoder service, or **concentrator** for the Concentrator service. You have the option of providing the volume type. The default volume type has the same name as the service.

Note: Run the `devlist` command to see if the multipath user-friendly names are listed correctly.

2. Run the `multipath -ll` command to make sure that SAN devices are attached. The following is an example when SAN devices are attached.

```
[root@116Decoder40GBDTrans block]# multipath -ll
mpathb (36006016001e04100babaab5acb9a24e0) dm-17 DGC ,VRAID
size=20T features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| '- 17:0:0:1 sdj 8:144 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
| '- 7:0:0:1 sdh 8:112 active ready running
mpatha (36006016001e04100e5baab5a5c2c6979) dm-2 DGC ,VRAID
size=10T features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| '- 17:0:0:0 sdi 8:128 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
| '- 7:0:0:0 sdg 8:96 active ready running
[root@116Decoder40GBDTrans block]#
```

Note: Block devices should be configured with a user-friendly name such as `mpatha`, `mpathb` etc.

3. Execute the `partNew` command to allocate block devices to partitions, volume groups, and logical volumes.

By default, the `partNew` command does not make changes. It displays the actions that will be taken if you commit the command string. To actually make the changes to the system, add the `commit=true` parameter to the command.

For example, to assign devices `sdd` and `sde` to Decoder:

```
send /appliance partNew name=sdc service=decoder volume=decodersmall
commit=true
```

```
send /appliance partNew name=sdd service=decoder volume=decoder commit=true
```

Caution: For the **decoder** and **concentrator** services, you must create storage volumes in a specific order.

- The **decoder** has the **decodersmall** and **decoder** volumes. Create the **decodersmall** volume before the **decoder** volume because **decodersmall** contains the small filesystem mounted at `/var/netwitness/decoder`.

- The **concentrator** has the **concentrator** and **index** volumes. Create the **concentrator** volume before **index** volume or it will fail and you receive the following message.

```
Failed to process message partNew for /appliance
com.rsa.netwitness.carlos.transport.TransportException: Volumes for index
require mount point /var/netwitness/concentrator to be created and
mounted first.
```

4. Execute the `vgs` command to validate that the `partNew` command created the correct Logical Volumes.

The output of this command:

- Enumerates all the volume groups on this host.
- Displays the physical volumes that the volume group consists of, and the logical volumes within the volume group.

5. Go to [Task 4 - Allocate Volume Groups to NetWitness Services- `srvAlloc`](#).

Task 4 - Allocate Volume Groups to NetWitness Services - `srvAlloc`

The `srvAlloc` command configures services on a host to use storage in a volume group. You must provide the name of the service to configure and the volume group to assign to the service (the service you provide must be installed on the host). For information about NetWitness Platform service volumes, see "NetWitness Platform Service Volume Reference" in [Storage Requirements](#).

Allocate services in the following order:

- For the Decoder, allocate `decodersmall` first then the `decoder`.
- For a Concentrator, allocate `concentrator` first then `index`.

Note: By default, the `srvAlloc` command does not make changes. You must append the `commit=1` parameter to the command string to actually make the changes to the system and restart the specified service after making changes.

1. Execute the `srvList` command to see a list of services installed on this host.
The `srvList` command communicates with the service through the SSL port. You install a Category on a host. A Category can be a single service, or multiple related services, located on the same host.
2. Execute the `srvAlloc` command to configure a service on a host to use storage in a volume group.
For example:


```
service=concentrator volume=concentrator commit=1  
service=concentrator volume=index commit=1
```
3. Go to Task 5 - Reconfigure Services and Databases to Detect and Appropriately Use New Storage.

Task 5 - (Optional) Reconfigure Storage Configuration for 10G Capture

You need to reconfigure the Decoder service and databases for 10G capture. Complete the following steps so that the Network Decoder service and its database detect and use new free space.


1. In the **NetWitness** menu, go to  (Admin) > **SERVICES**.
The **SERVICES** view is displayed.
2. Select the **decoder**.
3. Under  (actions), select **View > Explore**.
The **Explore** tree for the service is displayed.

4. Reconfigure space on the **decoder** service.
 - a. Navigate to the **decoder**, right click, and click **Properties**.




The **Properties** dialog is displayed.

- b. Execute the `reconfig` command by selecting it from the drop-down list, specify `update=1 op=10g` in **Parameters**, and click **Send**.




5. Reconfigure space on the database.
 - a. Navigate to **database** in the service **Explore** tree, right click, and click **Properties**.



The **Properties** dialog is displayed.

- b. Execute the `reconfig` command by selecting it from the drop-down list, specify `update=1 op=10g` in **Parameters**, and click **Send**.



Prepare Unity Storage

You must work with your Dell EMC Storage Engineer to allocate storage within your Unity environment for the NetWitness Platform and ensure the allocated storage meets all of the NetWitness Platform Storage Requirements. Specifically, make sure that:

- You have at least two LUNS created for Decoders (meta /session and packet volumes).
- You have at least two LUNS created for Concentrators (index and meta volumes).
- Ensure block devices can meet the minimum IOPS for expected ingestion rates.

You must add every NetWitness host that uses the Unity storage as a host within the Unity interface. After you create hosts and LUNs, you must assign the LUNs to the hosts. Assigning the LUNs to hosts makes the storage visible to the hosts so they can locate the storage through the host-based Dell EMC PowerPath software.

Note: A Dell EMC engineer will configure the following Unity Array.

You need to perform the following tasks to prepare Unity Storage.

[Task 1 - Access Unisphere User Interface \(UI\)](#)

[Task 2 - Create Pools](#)

[Task 3 - Create LUNS](#)

[Task 4 - Register Hosts](#)


[Task 5 - Assign LUNS to Hosts](#)

[Task 6 - Install PowerPath](#)


Task 1 - Access Unisphere User Interface (UI)

1. Connect your workstation on the same subnet as the UNITY.
2. Open a browser and go to **http://<unisphereIP>** to connect to the Unisphere UI.
3. Log in with the credentials provided by the Dell EMC CE. The default credentials are **admin/Password123#**.

Note: Unisphere will ask you to change password the first time log in. It also asks you to install the license before you can configure array (Dell EMC CE may do this for you. You must get the new admin password from them).




After you log in to Unisphere, the main dashboard is displayed.




Task 2 - Create Pools

The NetWitness configuration consists of two different pools. One pool is dedicated to the NL-SAS drives and the other pool is dedicated to the SSDs.

1. From **Storage Section**, click > **Pools** > **+** (Add) to launch the Create Pool Wizard.



2. Enter in a name for the pool (for example, **NLPool**) and click **Next**. Optionally, you can also enter a description for the pool.



3. Select **Capacity Tier** under **Tier** for the tier type (drive type) and click **Change**.

Create Pool

✓ Name and Description
• Tiers
Drives
Capability Profile Name
Summary
Results

Select Storage Tiers

Available Tiers

Tier	Drive Type	Unused Drives	Unused Capacity (GB)
☐ Extreme Performance Tier	SAS FLASH	1	165.1
☐ Performance Tier	SAS	6	3,301.8
☒ Capacity Tier	NL SAS	40	220,121.7

[Use FAST Guide](#)

Selected Tiers

Capacity Tier
RAID 6 (5+2), Maximum Usable Capacity 128.9 TB

[Change](#)

This pool will be created without any flash drives. For best performance, add flash drives to the pool according to the Best Practices Guide.

[Change](#)

Cancel [Back](#) [Next](#)

4. Choose the RAID type and from the drop down and select the RAID size.
The RAID type and size are a customer preference. The only requirement is to make sure you have enough IOPS within the pool to accommodate the log or packet capture and queries. In the following example, a **RAID 5 (8+1)** configuration is selected, however some customers may prefer a **RAID 6 (10+2 or 12 +2)**.
5. Make sure you have the correct Raid type and size selected.

Create Pool

✓ Name and Description
• Tiers
Drives
Capability Profile Name
Summary
Results

Select Storage Tiers

Available Tiers

Tier	Drive Type	Unused Drives	Unused Capacity (GB)
☐ Extreme Performance Tier	SAS FLASH	1	165.1
☐ Performance Tier	SAS	6	3,301.8
☒ Capacity Tier	NL SAS	40	220,121.7

[Use FAST Guide](#)

Selected Tiers

Capacity Tier
RAID 5 (8+1), Maximum Usable Capacity 171.9 TB

[Change](#)

This pool will be created without any flash drives. For best performance, add flash drives to the pool according to the Best Practices Guide.

Cancel [Back](#) [Next](#)


Proceed to the next step.

6. Choose the number of drives you want to add into the pool and click **Next**.

7. Skip the **VMware Capability** section and click **Next**.


8. Make sure that everything is correct on the Summary tab, and click **Finish**.
9. Create another pool for the SSDs using steps 2 – 8.
- Enter in a name for the other pool (for example, **SDDPool**) and click **Next**. Optionally, you can also enter a description for the pool.
 - Select **Extreme Performance Tier** under **Tier** for the tier type (drive type) and click **Change**.
 - Choose the RAID type and from the drop down, select the RAID size, and click **OK**.

Note: Raid 5 (4+1) RAID Configuration is different then Capacity Tier.



Task 3 - Create LUNS

1. From **Storage** section, click **Block** > **+** (Add) to launch the **Create LUN Wizard**.



The table below list all of the possible LUNS you may need to create. The ConIndex is the only LUN you need to assign to the SSD Pool. Make sure that the LUN sizes do not exceed what is listed below.

DecoderLarge01	75 TB orLess	NL Pool	No
DecoderSmall01	20 TB or Less	NL Pool	No
Concentrator01	15 TB or Less	NL Pool	No
Archiver01	75 TB or Less	NL Pool	No
ConIndex01	3 TB or Less	SSD Pool	No

2. Enter the LUN Name from the list. Optionally, you can enter a description of LUN.
3. Select the appropriate pool from the list on the drop-down menu.
4. Deselect the **Thin** checkbox (These will be fully provisioned LUNs).
5. Select **Next** to proceed to the next menu.

Create LUNs

Configure LUN(s)

Number of LUNs: 1

Name: DecoderLarge01

Description:

Pool: NLPool (Capacity Tier, 171.9 TB free)

Tiers Policy: Start High Then Auto-Tier

Size: 20 TB

☒ Thin

☐ Compression

Host I/O Limit: No Limit

Buttons: Cancel, Next

Proceed to the next step.

6. Click **Next** until you get to the summary section.
7. Verify that the **Name**, **Pool**, **Size** and **Thin** selections are all correct.
8. Click **Finish** to complete LUN creation.

Create LUNs

Summary

Name and Description

Name: DecoderLarge01

Description:

LUN Configuration

Pool: NLPool

Size: 20.0 TB

Thin: No

Compression: No

Tiers Policy: Start High Then Auto-Tier

Host I/O Limit: No Limit

Host Access

Access has not been configured for any hosts.

Snapshot

No snapshot schedule will be assigned to this LUN.

Replication

No replication is being created

Buttons: Cancel, Back, Finish


- Repeat steps 2- 8 for the remaining LUN creations.

Task 4 - Register Hosts


Before proceeding, record the hostname and IP address of the Head Unit and make sure that the HBAs in the head unit are properly cabled to the UNITY.

- From the Access section, click **Initiators**.
- Under the **Initiator Paths** tab, make sure that the correct HBAs are selected that you will use to register the Head Unit.

You should see two initiators per Head Unit. This represents the fiber connection from port 1 to SPA and port 1 to SPB. If you have multiple head units, the easiest method is to power each down and then power them up and register one by one.



- From the **Access** section, click **Hosts** >  (Add) > **Host** to add a host configuration.



- Enter the Hostname of the Head Unit.
- Under **Operating System**, select **Linux** from the-drop down menu.
- Enter the IP address of the Head Unit.

7. Click **Next** to proceed to the next section.

Add a Host

Specify a Name and Additional Information

Name: * 95Decoder

Description:

Operating System: Linux

Network Address: 10.25.66.32

Tenant: Select or enter a tenant.

While the host operating system information is not required, providing it will allow for more specific setup and troubleshooting instructions.

In order to customize access to NFS shares, the Network Address (name or IP address) is required. Port information is not allowed.

Network Address examples:
 IPv4 address: 192.168.1.2
 IPv6 address: FE80:3202:B3FF:FE1E:8329
 Host name: hostname

Tenant information is not required. To create a tenant, select the Tenants tab for a file system.

Cancel **Next**

Proceed to the next step.

8. In the Initiators section, select the two initiators that correspond to the correct port associated with the Head Unit and click **Next** to proceed.

Add a Host

Select Discovered Initiators or Manually Add Initiators

Automatically Discovered Initiators

Initiator IQN/WWN	Connected To
☒ 20:00:00:90:FA:A7:F5:E6	SP A iO Module 1 FC Port 1
☐ 20:00:00:90:FA:A7:BF:10:00:00:90:FA:A7:FA:BF	SP B iO Module 1 FC Port 2
☐ 20:00:00:90:FA:A7:F3:6A	SP B iO Module 1 FC Port 3
☒ 20:00:00:90:FA:A7:F3:6A	SP B iO Module 1 FC Port 3

Manually Added Initiators

No initiators have been manually added yet. Click the + button to manually add an initiator.

Cancel **Back** **Next**

Proceed to the next step.

- Make sure that the **Name**, **OS**, **IP** and **WWNs** are correct and click **Finish**.

Add a Host

Review the host configuration

Name: **SSDecoder**

Description: **Linux**

Operating System: **Linux**

Network Addresses: **10.25.66.32**

Tenant:

Initiators to be registered with this host


Protocol	Initiator IQN/WWN
FC	20 00 00 90 FA A7 F5 E6 10 00 00 90 FA A7 F5 E6
FC	20 00 00 90 FA A7 F5 E7 10 00 00 90 FA A7 F5 E7

Buttons: Cancel, Back, **Finish**


- Repeat steps 2-9 for all Head Units.
- In the Initiators section, select the two initiators that correspond to the correct port associated with the Head Unit. Then click “Next” to proceed.

Task 5 - Assign LUNS to Hosts


- From the **Access** section, click **Hosts**, select the head unit (for example, **Decoder**) and click  (edit) to view and edit details for the selected host.




2. In the **Properties** section, select the **LUNS** tab and click **+** (Add icon).




3. From the **Add LUNs to <Host>** popup, click **+** > **Add existing LUNs**.




4. Select the LUN to add to the Head Unit and **OK**.



5. Make sure that the correct LUN was added to the host and click **OK**.



6. (OPTIONAL) If you need to modify the HLU (Host LUNN Unique ID):
 - a. Select the LUN you want to change.
 - b. Click **Modify Host LUN IDs**.



7. Click  (edit), change the HLU to the number you want, and click **OK**.

Task 6 - Install PowerPath

1. Make sure that the Emulex ports on the Decoder host are attached to the Unity.
2. Log in to `root` on the Decoder attached to the Unity with the `admin` credentials.
3. Install PowerPath and register the Dell EMC PowerPath licenses for Unity hardware.

```
yum install DellEMCPower.LINUX-6.4.0.00.00-95.RHEL7.x86_64.rpm
```

Note: When you purchase an NetWitness Provided Unity, PowerPath licenses are sent to you. You can download PowerPath at support.dell.com.

Note: It is possible that the RPM downloaded from Dell EMC is not signed with a cert that the NetWitness device has available, which can cause the installation to fail with the package not signed error. Run the yum install with the `--ngpgcheck` option to enable the software to install.

4. Make sure that all the PowerPath connections are correct.

```
powermt display dev=all
```

The following output is an example of valid PowerPath connections.

```
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State   Q-I/Os Errors
=====
    15 lpfc           sde      SP A6    active  alive    0      0
    18 lpfc           sdg      SP B6    active  alive    0      0

Pseudo name=emcpowerb
Unity ID=APM00174407815 [Host_62]
Logical device ID=600601609D9046006996745A46B60AB6 [DecoderSmall01]
state=alive; policy=CLAROpt; queued-I/Os=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State   Q-I/Os Errors
=====
    15 lpfc           sdd      SP A6    active  alive    0      0
    18 lpfc           sdf      SP B6    active  alive    0      0
```

5. Verify that the PowerPath license is installed using the `emcpreg` command.

```
[root@NWAPPLIANCE24932 ~]# emcpreg -list
```

```
Key BQPO-DB4M-VFC2-Q24R-ML9Z-EQTU
```

```
Product: PowerPath
```

```
Capabilities: Al
```

6. Add the following string to the `/etc/lvm/lvm.conf` file to filter the LVM (Logical Volume Manager) so that it ignores duplicate volumes.

```
filter = [ "a|^/dev/sda2$|", "a|^/dev/sdb1$|",
"a|^/dev/emcpower.*|", "r|.*/|" ]
```

7. Run the following commands in this order:

```
a. systemctl enable PowerPath.service
```

```
b. systemctl start PowerPath.service
```

8. Reboot the Decoder.

9. Complete the instructions in [Configure Storage Using the REST API](#) to complete storage configuration.

Migrate Data to Another Storage Type

This section provides two options for moving data from DACs to PowerVaults:

[Migrate Data Using the Warm and Hot Tier Option](#)

[Move Data From DAC to PowerVault](#)

Refer to the Hardware Setup Guides on [NetWitness Community](#) for detailed instructions for setting up NetWitness Platform host and storage hardware.

Migrate Data Using the Warm and Hot Tier Option

In this procedure, you configure a warm tier for the DAC's, so that they do not write any new data. The warm tier continues to be available for analyst operations. You configure the PowerVaults as a hot tier, where new data can be written and available for analysts. When the required data retention is available on the hot tier, the warm tier can be decommissioned.

To set up the warm and hot tiers, perform the following tasks:

- [Stop the Service](#)
- [Set Up PowerVault](#)
- [Configure The Mount Points](#)
- [Set up Warm and Hot Tiers](#)
- [Decommission the DAC](#)

Stop the Service

1. Log in to the NetWitness Platform user interface.
2. Go to  (Admin) > **SERVICES** and select the service (for example, Log Decoder).
3. Click  > **View** > **Config**, and under Log Decoder Configuration, clear the **Capture Autostart** checkbox, and then click **Apply**.
4. In the menu bar, click the down arrow next to **Config**, select **System**, and at the top of the panel, click **Stop Capture**.
5. From the command line interface in NwConsole, stop the service by running the following command:
`systemctl stop nwlogdecoder`

Set Up PowerVault

1. Go to the REST API for the service by entering the IP address of the service, in this example, the Log Decoder. For example, `172.16.0.1:50106`.
2. Click the asterisk (*) next to the service. for example, **decoder (*)**.

3. Under **Properties for /decoder**, click the down arrow, select **RaidNew** and enter the following parameters, entering the name of the service for scheme. In this example, we use logdecoder.
`controller=1 enclosure=75 scheme=logdecoder-hotspare commit=1`
4. Click **Send**.
5. To configure the partitions, click the down arrow again, select **PartNew**, and enter the following parameters,
`name=sde service=logdecoder volume=logdecodersmall commit=1`
6. Click **Send**.
7. With **PartNew** still selected, enter the following parameters:
`name=sdf service=logdecoder volume=logdecoder commit=1`

Note: To validate the partition definitions before committing them, you can enter these parameters without `commit=1`, and click **Send**. After you validate the parameters, add `#commit=1` and then click **Send** to commit the parameter settings.

Configure The Mount Points

1. On the NwConsole at the root level of the service (for example, the Log Decoder), run `df -h`.
A list of mounted partitions is displayed.
2. Unmount all of the old storage points of the DAC and copy all the data to the Log Decoder. At the root level, run the `umount` command and the path name of each partition. You can concatenate the path names, for example:

```
umount /var/netwitness/logdecoder/index
/var/netwitness/logdecoder/sessiondb /var/netwitness/logdecoder/metadb
/var/netwitness/logdecoder/packetdb /var/netwitness/logdecoder/index0
/var/netwitness/logdecoder/sessiondb0 /var/netwitness/logdecoder/metadb0
/var/netwitness/logdecoder/packetdb0
```
3. Temporarily mount the petitions in the decoroot folder in the /mnt directory in order to access the files. For example:
`mount /dev/mapper/logdecodersmall-decoroot /mnt/decoroot/`
4. Copy the contents of decoroot from /mnt to /var/netwitness/logdecoder, answering Y (yes) to the prompts:
`cp -R statdb /var/netwitness/logdecoder/`
5. Unmount /mnt/decoroot.
`umount /mnt/decoroot`
6. Comment out decoroot from /etc/fstab, as this was on the DAC and the DAC will be decomissioned.

```
#/dev/logdecodersmall/decoroot
/var/netwitness/logdecoder/xfs/noatime,nosuid 1 2
```
7. Mount all the remaining file systems.
`mount -a`
8. Start the nwlogdecoder service (with capture still disabled).
`systemctl start nwlogdecoder`

Set up Warm and Hot Tiers

Caution: Before you set up warm and hot tiers, be sure that you know the right warm and hot tier entries for each collection so that you can set them up accurately.

1. Go to  (Admin) > **SERVICES** and select the service (for example, Log Decoder).
2. For the Log Decoder service, click  > **View** > **Explore**, and go to **database** > **config**.
 - a. Copy the contents of `meta.dir` and paste them to `meta.dir.warm` as shown in the following example:

logdecoder - Log Dec... <|

logdecoder - Log Decoder (LOG_DECODER)

connections

database

config

stats

decoder

deviceappliance

index

logs

rest

/database/config		logdecoder - Log Decoder
hash.algorithm		none
hash.databases		session,meta,packet
hash.dir		
manifest.dir		
meta.compression		none
meta.compression.level		0
meta.dir		/var/netwitness/logdecoder/metadb=4.58 TB
meta.dir.cold		
meta.dir.warm		
meta.file.size		auto
meta.files		auto

logdecoder - Log Dec... <|

logdecoder - Log Decoder (LOG_DECODER)

connections

database

config

stats

decoder

deviceappliance

index

/database/config		logdecoder - Log Decoder
hash.algorithm		none
hash.databases		session,meta,packet
hash.dir		
manifest.dir		
meta.compression		none
meta.compression.level		0
meta.dir		/var/netwitness/logdecoder/metadb=4.58 TB
meta.dir.cold		
meta.dir.warm		/var/netwitness/logdecoder/metadb=4.58 TB


- b. In the same way, copy the packet database in `packet.dir` to `packet.dir.warm`.
 - c. Copy the session database in `session.dir` to `session.dir.warm`.
3. Go to **index** > **config** and copy `index.dir` to `index.dir.warm`.

Note that the new volumes end in 0, so PowerVault will write to the directories ending in 0, for example:

```
[root@logdecoder ~]# df -h
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root          30G   3.3G   27G  11% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                     63G   12K   63G   1% /dev/shm
tmpfs                                     63G   34M   63G   1% /run
tmpfs                                     63G     0   63G   0% /sys/fs/cgroup
/dev/sdal                                1019M   96M   924M  10% /boot
/dev/mapper/netwitness_vg00-nwhome         3.3T   1.2G   3.3T   1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome        10G    33M    10G   1% /home
/dev/mapper/netwitness_vg00-varlog         10G   1.5G   8.6G  15% /var/log
tmpfs                                     13G     0    13G   0% /run/user/0
/dev/mapper/logdecodersmall-index          30G    54M    30G   1% /var/netwitness/logdecoder/index
/dev/mapper/logdecodersmall-sessiondb     600G   733M   599G   1% /var/netwitness/logdecoder/sessiondb
/dev/mapper/logdecodersmall-metadb         4.9T   11G   4.9T   1% /var/netwitness/logdecoder/metadb
/dev/mapper/logdecoder-packetdb           31T    12G   31T   1% /var/netwitness/logdecoder/packetdb
/dev/mapper/logdecodersmall0-index         30G    33M    30G   1% /var/netwitness/logdecoder/index0
/dev/mapper/logdecodersmall0-sessiondb    600G    34M   600G   1% /var/netwitness/logdecoder/sessiondb0
/dev/mapper/logdecodersmall0-metadb        21T    34M   21T   1% /var/netwitness/logdecoder/metadb0
/dev/mapper/logdecoder0-packetdb           86T    35M   86T   1% /var/netwitness/logdecoder/packetdb0
[root@logdecoder ~]#
```

Update the Decoder configuration with the path to the PowerVault mount by adding a 0 to the path.

1. In the /database/config column, right-click **meta.dir** and click **Properties**.
2. In **Properties for logdecoder**, select **set**, and in **Parameters**, enter `value="/var/netwitness/logdecoder/metadb0=4.58 TB" force=true`, as shown in this example, and then click **Send**.



3. Repeat step 2 for **session.dir**, **packet.dir**, and **index.dir**. Do not be concerned if the size is the same as the DAC in "`=xx GB`". This will be updated in the next step.

Note: We are only putting the PowerVault paths into the *.dir values.

4. Update the sizes for the live PowerVault volumes.
 - a. In the Log Decoder Explore view, in the left panel, right-click **database** and click **Properties**.
 - b. Select **reconfig** and in **Parameters**, enter `update=1` and click **Send**.
 - c. Repeat steps a and b for **index**.
5. Restart the service.


```
systemctl restart nwlogdecoder
```

6. Go to  (Admin) > **SERVICES**, select the Log Decoder service, and click  > **View** > **System**.
7. Click **Start Capture**.
8. Go to the **Config** view, select **Capture Autostart**, and click **Apply**.
9. Reboot the host.

Decommission the DAC

When the DAC data has aged, you should go back into the Explore view and remove all of the *.dir.warm configurations for session, meta, packet and index. You can determine when the DAC data

has aged by going to the Log Decoder  > **View** Explore view. Since we have a hot and warm tier, there are two sets of configuration stats that you need to be aware of. For example, for a packet Decoder, when you look at the packet oldest time in `packet.oldest.file.time`, look at the `packet.oldest.file.time.hot` value and if you see that your DAC had storage up until 30 days ago you can take your DAC offline and decommission it.

These are the basic steps for decommissioning a DAC. NetWitness recommends that you work with your Customer Support representative when you decommission your DACs.

1. Go to  (Admin) > **SERVICES** and select the service (for example, Log Decoder).
2. Click  > **View** > **Config**, and under Log Decoder Configuration, clear the **Capture Autostart** checkbox, and then click **Apply**.
3. In the menu bar, click the down arrow next to **Config**, select **System**, and at the top of the panel, click **Stop Capture**.
4. From the commandline interface in NwConsole, stop the service by running the following command:

```
systemctl stop nwlogdecoder
```
5. Unmount the warm tier. At the root level, run the `umount` command and the path name of each partition. You can concatenate the path names, for example:

```
umount /var/netwitness/logdecoder/index  
/var/netwitness/logdecoder/sessiondb /var/netwitness/logdecoder/metadb  
/var/netwitness/logdecoder/packetdb /var/netwitness/logdecoder/index0  
/var/netwitness/logdecoder/sessiondb0 /var/netwitness/logdecoder/metadb0  
/var/netwitness/logdecoder/packetdb0
```
6. Comment out all the old DAC dbs from `/etc/fstab`, so that only the PowerVault dbs remain.
7. Start the service.

```
systemctl start nwlogdecoder
```
8. In the user interface, go to  (Admin) > **SERVICES** and select the Log Decoder service.
9. Click  > **View** > **Explore** and remove the warm tier configurations:
 - a. In **database** > **config**, delete the content for `meta.dir.warm`, `packet.dir.warm`, `session.dir.warm`.

- b. In **index > config**, delete the content for `index.dir.warm`.
 - c. Go to the **Config** view, select **Capture Autostart**, and click **Apply**
 - d. Go to the **System** view and click **Start Capture**.
10. Restart the service.


```
systemctl restart nwlogdecoder
```

The DAC is now unmounted, and is no longer configured in the Decoder for warm storage and is ready to be wiped clean.

 1. Remove the logical volume. Run `lvscan` to get a list of the logical volumes.
 2. Run `lvremove` on the old logical volumes, for example:


```
/dev/logdecodersmall/decroot /dev/lvremove /dev/logdecodersmall/index
/dev/logdecodersmall/sessiondb /dev/logdecodersmall/metadb
/dev/logdecodersmall/packetdb
```
 3. Remove the volume groups. Run `vgscan` to get a list of volume groups.
 4. Run `vgremove` on the old volume groups (be careful not to remove any volume groups that end in 0, as they are PowerVault).
 5. Run `pvs` to view block devices that are freed up.
 6. When the DAC has been successfully removed, reboot the host.

Move Data From DAC to PowerVault

The following procedure describes how to move data from DAC to PowerVault. Before you move data from 2 DACs to 2 PowerVaults, a table, similar to the following table, is displayed if you run the `pvs` (Physical Volume Size) command from the Decoder Linux console (or SSH to the Decoder) with 2 DACs attached and configured to the Decoder. The column headings are Physical Volume (PV), Volume Group(VG), Linux Format (Fmt), Linux Attribute (Attr), Physical Volume Size (PSize), and Physical Volume Free Space(PFree).

PV	VG	Fmt	Attr	PSize	PFree
/dev/sda2	netwitness_vg00	lvm2	a--	<930.00g	0
/dev/sdb1	netwitness_vg00	lvm2	a--	<1.82t	0
/dev/sdc	decodersmall	lvm2	a--	<5.46t	0
/dev/sdd	decoder	lvm2	a--	<27.29t	0
/dev/sde	decodersmall0	lvm2	a--	<5.46t	0
/dev/sdf	decoder0	lvm2	a--	<27.29t	0

Complete the following steps to move data from a DAC to a PowerVault.

1. Attach two PowerVaults to a separate PERC controller on the Decoder.
2. Create the devices.
 - a. Open a Browser and specify the ip-address of the Network Decoder and port **50106** to access the REST tool.
 - b. Log in with the `admin` account credentials.



- c. Click on the `(*)` next to **appliance** to access the REST command set.
- d. Run `raidList` to display the Controller/Enclosure combination with the new PowerVault enclosures.

In the following example, the output shows **dev/sdg** and **/dev/sdh** on **Controller 2, Enclosure 246**.

```
Controller 2, Enclosure 246
Vendor:  DELL
Model:   MD1400
In Use:  true
Drives: 10.691 TB x 12
Devices: sdg
         sdh
```

- e. Under **Properties for /appliance**, select `raidNew`, specify `controller=<PowerVault-controller-id>` `enclosure=<PowerVault-enclosure-id>` `scheme=decoder-hotspare` `preferSecure=false`, and click **Send**.

Note: You specify `preferSecure=false` if the PowerVault drives are not SED drives. If PowerVault drives are SED drives and you do not want to encrypt them you specify `preferSecure=false`. You must specify `preferSecure=true` if PowerVault drives are SED drives and you want to encrypt them.

3. Go to the Decoder Linux console or SSH to the Decoder and run the following commands.

```
parted -s /dev/sdg mklabel gpt
parted -s -a optimal /dev/sdg mkpart LVM 0% 100%
pvcreate -f /dev/sdg
parted -s /dev/sdh mklabel gpt
parted -s -a optimal /dev/sdh mkpart LVM 0% 100%
```

```
pvccreate -f /dev/sdh
```

If the volume is created successfully, the following message is displayed.

```
Physical volume "/dev/sdg" successfully created
```

Note: Repeat this step for every block device. The block device names may be different depending on how many enclosures per perc card slot.

4. Run the following command strings to extend the DAC volume group (**decoder**, **decodersmall**) to the PowerVault Physical volume.

```
vgextend decoder /dev/sdg
```

```
vgextend decodersmall /dev/sdh
```

5. Run the following command strings to move the data from the DAC to the PowerVault. In this following command string, the DAC is **/dev/sdc** and the PowerVault is **/dev/sdg**.

```
pvmove /dev/sdc /dev/sdg
```

```
pvmove /dev/sdd /dev/sdh
```

Note: 1.) The `pvmove` command synchronizes data across volumes so that NetWitness can continue ingesting or aggregating data while the migration is executing. You can run the `pvmove` command multiple times if it fails. 2.) Depending on the amount of data on the drives, the move can take a long time complete depending on the amount of data. For example, in a test, it took four hours to move one TB of data.

6. After the move is complete, run the following commands to reduce and remove the DAC drive.

```
vgreduce decoder /dev/sdc
```

```
pvremove /dev/sdc
```

```
vgreduce decodersmall /dev/sdd
```

```
pvremove /dev/sdd
```

7. Detach the physical connections from the DACs to the host.
8. Verify that the Physical volumes are moved from the DACs to the PowerVaults.

- a. Reboot the host.

```
reboot
```

- b. Verify that the **/etc/fstab** file is correct.

- c. Run the `pvs` command and make sure that the **PSize** and **PFree** values are correct on the PowerVault.

```
root@nitritiercdecoder:~# pvs
```

PU	VG	Fmt	Attr	PSize	PFree
/dev/sda2	netwitness_vg00	lvm2	a--	<930.00g	0
/dev/sdb1	netwitness_vg00	lvm2	a--	<1.82t	0
/dev/sdc1	decodersmall	lvm2	a--	21.38t	<15.93t
/dev/sdd1	decoder	lvm2	a--	<85.54t	58.25t

Data on PowerVault After Move from DAC

After you move data from 2 DACs to 2 PowerVaults, a table, similar to the following table, is displayed if you run the `pvs` (Physical Volume Size) command from the Decoder Linux console (or SSH to the Decoder) with 2 PowerVaults attached and configured to the Decoder. The column headings are Physical Volume (PV), Volume Group(VG), Linux Format (Fmt), Linux Attribute (Attr), Physical Volume Size (PSize), and Physical Volume Free Space(PFree).

PV	VG	Fmt	Attr	PSize	PFree
/dev/sda2	netwitness_vg00	lvm2	a--	<930.00g	0
/dev/sdb1	netwitness_vg00	lvm2	a--	<1.82t	0
/dev/sdc1	decodersmall	lvm2	a--	21.38t	<15.93t
/dev/sdd1	decoder	lvm2	a--	<85.54t	58.25t

SASE Node-x (Decoder/Concentrator) - GCP Persistent Disk (PD) Storage Configuration

This section contains:

- [Introduction](#)
- [Identify Storage Requirements](#)
- [Identify or Define Storage Model](#)
- [Deploy SASE Node\(s\)](#)
- [Configure SASE Node\(s\) Storage](#)
- [Extend Storage for SASE Node](#)
- [Appendix](#)

Introduction

The SASE node-x (decoder/concentrator) storage configuration is dependent on a few configuration attributes defined in `/opt/rsa/saTools/cloud/host-models.yml` (referred to as `host-models` file) and `/opt/rsa/saTools/cloud/sase-deployment-models.yml` (referred to as `sase-deployment-models` file).

The configuration files are installed on *Admin Server/Node-z* during orchestration (*Refer to SASE Installation Guide for details*). When the SASE deployment script (`nw-create-cloud-hybrid`) is run for the first time, it copies both the configuration files (`/opt/rsa/saTools/cloud/host-models.yml` and `/opt/rsa/saTools/cloud/sase-deployment-models.yml`) to `/root/.sase/` directory on the Admin Server/Node-z. Any subsequent SASE deployments refer to `/root/.sase/host-models.yml` and `/root/.sase/sase-deployment-models.yml`. All subsequent updates such as changing a model-name must be made to `/root/.sase/sase-deployment-models.yml`.

The SASE node's disk specifications such as `disk_name`, `disk_type`, `disk_size` for each node along with gcp virtual machine type (`machine_type`) is collectively referred to as a model. Three models: `c1r6m30`, `c1r12m60` and `c1r23m120` are defined in the `host-models.yml`. One of these models is assigned as `model-name` attribute value in `sase-deployment-models.yml`.

The SASE node(s) deployment and storage configuration are tightly coupled and storage needs should be considered before installation or deployment. The SASE Installation guide should be cross referenced for complete attribute definition in `host-models` and `sase-deployment-models` files. As a best practice, the storage requirements should be identified and `host-models.yml` and `sase-deployment-models` files updated before deploying or installation of the nodes.

The below outlines the steps for successful installation and storage configuration of SASE nodes in GCP.

1. Identify the storage requirements for each of the SASE services - decoder and concentrator.
2. Identify or define the storage model: Identify a predefined model or define a new custom model. If defining a custom model, copy the `/opt/rsa/saTools/cloud/host-models.yml` and `/opt/rsa/saTools/cloud/sase-deployment-models.yml` to `/root/.sase/` and update the `model-name` attribute for every node in `/root/.sase/sase-deployment-models.yml`.

3. Deploy the SASE GCP nodes: Execute the `nw-create-cloud-hybrid` script on admin node or node-z to complete the installation. *Refer to SASE Installation Guide for details.*
4. Complete storage configuration using **Explore** view of the Decode or Concentrator by logging into *Admin Server* or *UI*.

Identify Storage Requirements

This section covers step 1 defined in Introduction.

Based on storage needs, an appropriate model-name (`c1r6m30`, `c1r12m60` or `c1r23m120`) is assigned to `model_name` attribute in `/root/.sase/sase-deployment-models.yml` file. During SASE Node-x deployment the setup script(`/usr/bin/nw-create-cloud-hybrid`) refers to `additional_disks` attribute value in `sase-deployment-models.yml` to determine whether storage disks must be created or not. When the `additional_disks` is set to `true` for a node, the corresponding `model_name` attribute value is used to identify the model specification (such as `disk_name`, `disk_type` and `disk_size`) in `/root/.sase/host-models.yml` and the corresponding disks are created and assigned (but not configured i.e these disks are ready for storing NW databases) to the SASE node.

When the predefined models don't satisfy the customer requirements, custom models (*Refer to Define a custom storage model*) can be defined. Please contact Professional Services/ NW Support for details on defining custom models.

Example: Sample content of `sase-deployment-models.yml` identifying `model_name` and `additional_storage` attribute values:

```
# Define model name to configure host and storage
model_name: c1r12m60
additional_storage: true
```

Note:

- The `storage_class` and `warm_retention` attributes are not supported in NW 12.4. These are developmental features and should be ignored.
- The `disk_size` value is always defined in GB (Gigabytes).
- `c1default` is for testing purposes only.

Model	Description
c1r6m30	- Creates and attaches persistent disks to decoder service that are capable of 6 days of packet retention at 1gbps line rate (capture) with 100% utilization. - Creates and attaches persistent disks to concentrator service (concentrator volume) and pd-ssd disks for index that are capable of 30 days of concentrator meta retention at 1gbps line rate (capture) with 100% utilization.

Model	Description
c1r12m60	- Creates and attaches persistent disks to decoder service that are capable of 12 days of packet retention at 1gbps line rate (capture) with 100% utilization. - Creates and attaches persistent disks to concentrator service (concentrator volume) and pd-ssd disks for index that are capable of 60 days of concentrator meta retention at 1gbps line rate (capture) with 100% utilization.
c1r23m120	- Creates and attaches persistent disks to decoder service that are capable of 23 days of packet retention at 1gbps line rate (capture) with 100% utilization. - Creates and attaches persistent disks to concentrator service (concentrator volume) and pd-ssd disks for index that are capable of 120 days of concentrator meta retention at 1gbps line rate (capture) with 100% utilization.

Identify or Define Storage Model

This section covers step 2 defined in the [Introduction](#) section.

Every predefined model defines *machine_type* attribute for boot disk and a minimum of two disks (one set) for storage (referred to as initial disks or default storage disks or default disks). The default storage disks for decoder service are named (*disk_name* attribute in *host-models.yml*) as *decoder* and *decodersmall*. The default disks for concentrator are *index* and *concentrator*. The disk naming convention is also tied to subsequent configuration steps such a creation of partitions, volume allocation to service. It also helps to associate the corresponding partition scheme that is created on that particular disk when using NW REST API to complete storage configuration

Each node (*Decoder* or *Concentrator*) has its own storage attributes defined under disks (*Refer to screen shot below*). The attributes are further classified into two groups for each service with each group corresponding to a specific logical volume name. For decoder service, the two groupings refer to the logical volume names: *decoder* and *decodersmall* and for concentrator service these are *concentrator* and *index*.

Note: The grouping names (*decoder* and *decodersmall* for decoder and *concentrator* and *index* for concentrator) under *disks* match the *disk_name* attribute values that they define. This relationship (underlined in red below) **MUST** be maintained when defining a Custom storage model since the disk creation process requires this 1-1 mapping of grouping name to *disk_name*.

host-models.yml snippet highlighting the disk groupings and disk_name relationship (underlined in red):

```
c1r12m60:
  Decoder:
    machine_type: n2-standard-32
    storage_class: STANDARD
    # retention size in TB
    warm_retention: 1
    disks:
      # allocate to decoder root, index, sessiondb, metadb
      decodersmall:
        disk_name: decodersmall
        disk_type: pd-standard
        disk_size: 3000
      # allocate to packetdb
      decoder:
        disk_name: decoder
        disk_type: pd-standard
        disk_size: 65000
      # allocate to decoder root, index, sessiondb, metadb
      decodersmall0:
        disk_name: decodersmall0
        disk_type: pd-standard
        disk_size: 3000
      # allocate to packetdb
      decoder0:
        disk_name: decoder0
        disk_type: pd-standard
        disk_size: 65000
```

Concentrator:

```

machine_type: n2-standard-32
storage_class: STANDARD
# retention size in TB
warm_retention: 1
disks:
  # allocate to concentrator root, metadb, sessiondb
  concentrator:
    disk_name: concentrator
    disk_type: pd-standard
    disk_size: 40000
  # allocate to index
  index:
    disk_name: index
    disk_type: pd-ssd
    disk_size: 2000
  # allocate to concentrator root, metadb, sessiondb
  concentrator0:
    disk_name: concentrator0
    disk_type: pd-standard
    disk_size: 40000
  # allocate to index
  index0:
    disk_name: index0
    disk_type: pd-ssd
    disk_size: 2000

```

Each disk attached to a SASE Decoder node is configured as a separate logical volume to host the service's database(s) such as *packetdb*, *metadb*, *sessiondb* and *index*. The disk named (*disk_name* attribute in *host-models* file) as decoder corresponds to decoder volume and hosts *packetdb* and the disk named (*disk_name* attribute in *host-models* file) as *decodersmall* corresponds to *decodersmall* volume and *host sessiondb*, *metadb* and *index*. If decoder storage requirements mandate larger disks (i.e 12 days of Packet retention versus 6 days of packet retention or 23 days of packet retention versus 12 days or 6 days) then multiple sets of disks are defined and created.

If a model defines two sets of disks, the first is the default disks named (*disk_name*) as *decoder* and *decodersmall* and the second set is named as *decoder0* and *decodersmall0*, likewise when a model defines four sets of disks ex: *c1r23m120*, first set is default disks (named *decoder* and *decodersmall*), second set of disk named as *decoder0* and *decodersmall0* and the third set of disks are named as *decoder1* and *decodersmall1* and fourth set is *decoder2* and *decodersmall2*. Any additional disk set follow the similar naming convention as described above. i.e *decoder<COUNTER>* and *decodersmall<COUNTER>* where COUNTER is incremented by 1 for any new set of disks starting with 0. The first part of the name (*decoder*, *decodersmall*) helps in associating the disk with volume name/type that is created later. The volume name is used to identify the appropriate Decoder service's database during service allocation.

Similar pattern is followed with *concentrator disks*. The default disk set is *index* and *concentrator*. The concentrator is *pd-standard* (*disk_type*) and index disks are *pd-balanced*. The *index disk* is used to host *index database* and *concentrator disk* is used to host *root*, *sessiondb* and *metadb*. When more than one disk set is created, these are named as *index<COUNTER>* and *concentrator<COUNTER>* where COUNTER starts with 0 and increments by 1 (similar pattern as decoder disks described above). Refer to above screen shots for pattern identification for *c1r12m60*.

The *additional_storage* attribute value in *sase-deployment-models.yml* determines whether the storage disks are created or not. A value of *true* creates the disks and *false* skips the disk creation. The default value is *false*. Updating the *additional_storage* to *true* and re-running the *nw-create-cloud-hybrid* creates the storage disks. These disks are not deleted by re-running *nw-create-cloud-hybrid* after setting *additional_storage* to *false*.

The storage requirements must be identified to select the appropriate storage model. The retention days along with capture rate is used to identify the appropriate model. After the model is identified, the *model_name* value for every SASE node must be updated in the *sase-deployment-models.yml* file. If the multiple nodes are deployed, then each node's *model_name* must be assigned with appropriate value.

For both predefined and custom models, after identifying the model (Refer to *Identify a Pre-defined Storage Model or Define a Custom Storage Model* section(s) below) the **model_name** attribute value (in *sase-deployment-models.yml*) is set to the identified model name.

Refer Appendix B for SASE Decoder Storage Configuration and Appendix C for SASE Concentrator Storage Configuration.

Deploy SASE Node(s)

This section covers step 3 defined in the [Introduction](#) section.

Assumptions:

- The *model_name* is updated with correct host-model and *additional_storage* attributes set to *true* in *sase-deployment-models.yml*
- When opting for custom storage model, the */root/.sase/host-model.yml* is updated with the custom model definition.

Execute the *nw-create-cloud-hybrid* script on *node-z* or *Admin Server* to complete the installation of SASE node-x. The storage disks are created and attached to the SASE nodes but not configured to host the NW service's database.

The *gcp SASE instance name* (SASE Decoder deployed) is a combination of *nw-<name>-<region_name>-<zone_suffix>*. The *name*, *region_name* and *zone_suffix* attributes are defined in *sase-deployment-model* file for the decoder node. Every SASE node follows the similar naming convention.

The deployed SASE node can be accessed by logging into *GCP account* → *Compute Engines* → *Virtual Machines* → *Search* for the SASE instance using the name.

Refer to SASE Installation Guide for details on installation.

Configure SASE Node(s) Storage

This section covers step 4 defined in the [Introduction](#) section.

Assumptions:

- The SASE decoder node is successfully bootstrap and orchestrated in gcp.
- All the storage disks associated with the node's model are created. As noted earlier, the disks are not created when an incorrect *model_name* is used or the *additional_disks* attribute is set to false. Incorrect *model_name* requires uninstalling the node and recreating it. If incorrect *additional_disks* value (false) is used, set the value to true and re-run the *nw-create-cloud-hybrid* script.

```
nw-create-cloud-hybrid --enable-cloud-sase
```

Configure SASE Decoder Storage

Refer to Appendix B: Sample scenario for Configuring SASE Decoder Storage to complete the configuration of the storage disks created during installation.

Configure SASE Concentrator Storage

Refer to Appendix C: Sample scenario for Configuring SASE Concentrator Storage to complete the configuration of the storage disks created during installation.

Extend Storage for SASE Node

Note: Only pre-defined host-models can be extended. Custom host models can not be extended.

Customer storage requirements such as higher retention for packet and meta data may change over time after initial storage configuration. SASE storage models support additional storage allocations. The first column in *Storage Extension Matrix* table below lists the current host model deployed and the second column shows the available host-models that the initial model can be extended.

Storage Extension Matrix:

Current Model (model_name)	Supported Storage Extensions (model_name)
c1r6m30	c1r12m60 c1r23m120
c1r12m60	c1r23m120
c1r23m120	N/A (cannot be extended)

Extend Decoder or Concentrator Storage

Extending decoder storage involves the following steps on the Admin Server or Node-z:

1. Identify the new model in *host-models.yml* and update the node's *model_name* value with the new model, *additional_storage* to *true* and click **Save**.
2. Execute the below command to extend the storage.

```
nw-create-cloud-hybrid --enable-cloud-sase
```
3. Complete the storage configuration using NW REST API by logging into Admin Server or UI and Use REST API utility to create partitions and service allocations (Navigate to **decoder's explorer view ->deviceappliance->properties->Right Click** and select **Properties** drop-down).
4. Complete the storage configuration of new disks added in the above step using REST API. *Refer to [Configure Storage Using the REST API](#) section for details. Refer to [Appendix B - Sample Scenario for Configuring SASE Decoder Storage](#) and [Appendix C - Sample Scenario for Configuring SASE Concentrator Storage](#) for sample scenarios.*

Appendix

This section contains:

- [Appendix A - Defining a Custom Host Model](#)
- [Appendix B - Sample Scenario for Configuring SASE Decoder Storage](#)
- [Appendix C - Sample Scenario for Configuring SASE Concentrator Storage](#)
- [Appendix D - Sample Scenario for Extending for SASE Decoder Storage](#)
- [Appendix E - Sample Scenario for Extending SASE Concentrator Storage](#)

Appendix A - Defining a Custom Host Model

Custom storage models can be created subject to Google Cloud Persistent Disk restrictions.

The below steps must be executed in sequence for using custom model before starting the SASE installation.

1. Copy */opt/rsa/saTools/cloud/sase-deployment-models.yml* and */opt/rsa/saTools/cloud/host-models.yml* to */root/.sase/*. Update */root/.sase/host-models.yml* with the custom model specification and **Save**.
2. Update the *model_name* attribute value with the above custom model name in */root/.sase/sase-deployment-models.yml* and **Save**.
3. Complete the SASE node deployment. *Refer to SASE Installation Guide.*

All the attributes must be assigned a valid value. It is recommended to append custom model definition to existing content in */root/.sase/host-models.yml*. *Refer to the SASE Installation Guide for details on the attribute description.*

Sample Custom Model Definition for Decoder

The sample custom configuration assumes defining a model capable of 1gbps capture with 10 days of packet retention and 40 days of concentrator meta retention : *c1r10m40* , the *machine_type* as *n2-standard-32* (Refer to Google Cloud documentation for details on all the available VM types) and disk_type as *pd-standard* and the storage requirements of **98000 GB** for packet db and **8000 GB** for other databases (metadb, sessiondb and index).

Due to gcp restrictions on the disk size, the model requires two sets of disks. Following the naming convention described earlier, the two disk sets are named (disk_name) as *decoder*, *decodersmall*, *decoder0* and *decodersmall0*.

Custom model definition for model_name: *c1r10m40*.

host-models.yml:

```

1  c1r10m40:
2      machine_type: n2-standard-32
3      disks:
4          Decoder:
5              decodersmall:
6                  disk_name: decodersmall
7                  disk_type: pd-standard
8                  disk_size: 4000
9              decoder:
10                 disk_name: decoder
11                 disk_type: pd-standard
12                 disk_size: 47000
13             decodersmall0:
14                 disk_name: decodersmall0
15                 disk_type: pd-standard
16                 disk_size: 4000
17             decoder0:
18                 disk_name: decoder0
19                 disk_type: pd-standard
20                 disk_size: 47000

```

Partial Content of sase-deployment-models.yml updated with custom model_name Value:

```
#Container element defining all nodes that will be created within this region.
nw_nodes:

  # First Node to be created. This element is just an arbitrary name for the type of
  # node to be created and provisioned in this region's subnet.
  decoder:

    # Name of the instance known to the nw-ppn network.
    name: decoder_w

    # This node's ip in C.I.D.R format. This address MUST be within the ppn_cidr range.
    ppn_cidr_ip: 172.30.30.6/24

    # This value is concatenated with the region to define the zone
    # that the nw node will be installed into on GCP.
    zone_suffix: '-b'

    # The size and type of boot disk attached to nw node when it is created.
    boot_disk_size: 196
    boot_disk_type: pd-standard

    # Define model name to configure host and storage
    model_name: c1r10m40
    additional_storage: true

    # Used to determine if the calling script will automatically bootstrap and accept
    # the node keys in the Admin Server. This allows for either automated or
    # manual orchestration of a NetWitness Category to the node.
    bootstrap: true

    # Used to determine if the calling script will automatically orchestrate
    # a NetWitness Category to the node.
    orchestrate: true

    # The NetWitness Category to be orchestrated. Must be an exact value (Case Sensitive)
    category: Decoder
```

Sample Custom Model Definition for Concentrator

The sample custom configuration model (*c1r10m40*) assumes the *machine_type* as *n2-standard-32* (Refer to Google Cloud documentation for details on all the available VM types) and concentrator meta retention of 40 days. The concentrator *disk_type* is *pd-standard*, index *disk_type* is *pd-balanced* and the storage requirements for metadb, sessiondb, root is **24000 GB** (for 40 days) and index database is **6000 GB** (for 40 days).

Since these sizes fall within the GCP limitations on per disk disk maximum size, only one set (default name) is needed: *concentrator*, *index*.

Custom model definition for model_name: c1r10m40

host-models.yml:

```
c1r10m40:
  machine_type: n2-standard-32
  disks:
    Concentrator:
      concentrator:
        disk_name: concentrator
        disk_type: pd-standard
        disk_size: 24000
    index:
      disk_name: index
      disk_type: pd-ssd
      disk_size: 6000
```

Partial Content of sase-deployment-models.yml updated with custom model_name Value:

```
#Container element defining all nodes that will be created within this region
nw_nodes:

  # First Node to be created. This element is just an arbitrary name for the
  # node to be created and provisioned in this region's subnet.
  concentrator:

    # Name of the instance known to the nw-ppn network.
    name: concentrator

    # This node's ip in C.I.D.R format. This address MUST be within the ppn_c
    ppn_cidr_ip: 172.30.30.6/24

    # This value is concatenated with the region to define the zone
    # that the nw node will be installed into on GCP.
    zone_suffix: '-b'

    # The size and type of boot disk attached to nw node when it is created.
    boot_disk_size: 196
    boot_disk_type: pd-standard

    # Define model name to configure host and storage
    model_name: c1r10m40
    additional_storage: true

    # Used to determine if the calling script will automatically bootstrap an
    # the node keys in the Admin Server. This allows for either automated or
    # manual orchestration of a NetWitness Category to the node.
    bootstrap: true

    # Used to determine if the calling script will automatically orchestrate
    # a NetWitness Category to the node.
    orchestrate: true

    # The NetWitness Category to be orchestrated. Must be an exact value (Case
    category: Concentrator
```

Appendix B - Sample Scenario for Configuring SASE Decoder Storage

The sample storage configuration in this section uses REST API utility to complete the configuration of the storage disks created during installation.

Refer Configure Storage Using the REST API section for details on REST API usage.

Assumptions:

- The SASE Decoder node is deployed using *c1default* (*model_name*) and *additional_storage: true* in *sase-deployment-models.yml* during SASE deployment. *c1default* is a test model for illustration purposes.
- The storage disks corresponding to *c1default* model are created and attached to the decoder node (by the installation scripts). These disks are not yet configured.

```

c1default:
  Decoder:
    machine_type: n2-standard-4
    storage_class: STANDARD
    # retention size in TB
    warm_retention: 1
    disks:
      # allocate to decoder root, index, sessiondb, metadb
    decodersmall:
      disk_name: decodersmall
      disk_type: pd-standard
      disk_size: 641
      # allocate to packetdb
    decoder:
      disk_name: decoder
      disk_type: pd-standard
      disk_size: 69

```

Follow these steps to configure SASE Decoder storage:


1. Login to SA UI. To list all the devices, Navigate to **Hosts -> Select the Decoder node -> Services -> Actions -> View -> Explore -> deviceappliance -> appliance-> Right Click -> Properties** drop-down -> Select **devlist** -> **Send**.

All the block devices (both configured and unconfigured) are returned. The configured devices has the '*used=1*' and unconfigured has '*used=0*'. In this case, the block devices are sdb and sdc. sda is the boot disk and no changes are allowed.

Note: decodersmall must be partitioned before decoder volume.


2. Use **partNew** command to partition unused block devices. For decoder service, the larger devices (decoder) are always allocated partitioned for packetdb and smaller devices (decodersmall) are allocated for *root/session/meta* and *index* databases. To partition sdc, select **partNew** from the **Properties** drop-down with the following parameters and click **Send**.

```
name=sdc volume=decodersmall service=decoder commit=1
```




- To partition *sdb* device, select 'partNew' from the **Properties** drop-down with the following parameters and click **Send**.

name=sdb volume=decoder service=decoder commit=1




- Allocate the above configured partitions to decoder service by selecting **srvAlloc** in the **Properties** drop-down and click **Send**.

service=decoder volume=decoder commit=1




- Allocate decodersmall to decoder service using *srvAlloc* property.

```
service=decoder volume=decodersmall commit=1
```



- View the storage allocation by issuing the below command in SSH-in-browser. The configured storage is highlighted in yellow.

```
df -hP
```



```
@NWAPPLIANCE26051 ~]$ df -hP
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  7.9G         0   7.9G   0% /dev
tmpfs                     7.9G   8.0K   7.9G   1% /dev/shm
tmpfs                     7.9G   8.5M   7.9G   1% /run
tmpfs                     7.9G         0   7.9G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G   4.2G   26G  14% /
/dev/mapper/netwitness_vg00-usrhome 10G   104M   9.9G   2% /home
/dev/mapper/netwitness_vg00-varlog 10G   113M   9.9G   2% /var/log
/dev/mapper/netwitness_vg00-nwhome 141G   1.2G  140G   1% /var/netwitness
/dev/sdal                 1014M   185M   830M  19% /boot
tmpfs                     1.6G         0   1.6G   0% /run/user/1269
/dev/mapper/decodersmall-decoroot 10G   105M   9.9G   2% /var/netwitness/decoder
/dev/mapper/decodersmall-index 30G   247M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb 600G   4.3G  596G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb 1015M   40M   976M   4% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb 69G   525M   69G   1% /var/netwitness/decoder/packetdb
```

Appendix C - Sample Scenario for Configuring SASE Concentrator

Storage

The sample storage configuration in this section uses REST API utility to complete the configuration of the storage disks created during installation. Refer to [Configure Storage Using the REST API](#) section for details on REST API usage.

Assumptions:

- The SASE Concentrator node is deployed using *c2default* (*model_name*) and *additional_storage: true* values in *sase-deployment-models.yml* during SASE deployment. *c2default* is a test or custom model for illustration purposes.
- The storage disks corresponding to *c2default* model are successfully created and attached to the

concentrator node. These disks are not yet configured.

```
Concentrator:
machine_type: n2-standard-4
storage_class: STANDARD
# retention size in TB
warm_retention: 1
disks:
  # allocate to concentrator root, metadb, sessiondb
  concentrator:
    disk_name: concentrator
    disk_type: pd-standard
    disk_size: 38
  # allocate to index
  index:
    disk_name: index
    disk_type: pd-ssd
    disk_size: 11
  # allocate to concentrator root, metadb, sessiondb
  concentrator0:
    disk_name: concentrator0
    disk_type: pd-standard
    disk_size: 38
  # allocate to index
  index0:
    disk_name: index0
    disk_type: pd-ssd
    disk_size: 11
```

Follow these steps to configure Concentrator storage:


1. Log into SA UI. To list all the devices, Navigate to **Hosts- > Select the Concentrator node -> Services -> Actions -> View -> Explore -> deviceappliance -> appliance -> Right Click -> Properties drop-down -> Select devlist -> Send.**

Response Output:

In this case, the block devices are *sdb*, *sdc*, *sdd* and *sde*. The '**used=0**' indicates that this device is not configured yet. When configured, the 'used' attribute value is set to '**1**'. As noted earlier, *sda* is the boot disk and no changes are allowed.

2. Use **partNew** command to partition unused block devices. For concentrator service, the smaller devices (disk_type: pd-ssd) are allocated to index database and larger devices are allocated to root,session and metadb. To partition *sdb*, select 'partNew' from the **Properties** drop-down with the following parameters and click **Send**.

```
name=sdb volume=concentrator service=concentrator commit=1
```




3. Repeat the above step for **concentrator0** volume. Use **partNew** command to partition unused block devices. To partition **sdc**, select 'partNew' from the **Properties** drop-down with the following parameters and click **Send**.

name=sdc volume=concentrator0 service=concentrator commit=1

4. To partition **sdd** device, select 'partNew' from the **Properties** drop-down with the following parameters and click **Send**.

name=sdd volume=index service=concentrator commit=1




Repeat the above for **sde**.

name=sde volume=index service=concentrator commit=1

5. Allocate the index volume to concentrator service using **srvAlloc** from the **Properties** drop-down with the following parameters and clicking **Send**.

service=concentrator volume=index commit=1



- Complete the service allocation of the three remaining volumes (index, index0, concentrator0) using srvAlloc property using the above step.

```
service=concentrator volume=index0 commit=1
```

```
service=concentrator volume=concentrator commit=1
```

```
service=concentrator volume=concentrator0 commit=1
```

- The configured storage can be viewed by issuing the below command in SSH-in-browser. The configured storage is highlighted in yellow.

```
df -hP
```

```
[root@NWAPPLIANCE18597 ~]# df -hP
```

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	7.9G	0	7.9G	0%	/dev
tmpfs	7.9G	8.0K	7.9G	1%	/dev/shm
tmpfs	7.9G	8.6M	7.9G	1%	/run
tmpfs	7.9G	0	7.9G	0%	/sys/fs/cgroup
/dev/mapper/netwitness_vg00-root	30G	4.1G	26G	14%	/
/dev/mapper/netwitness_vg00-usrhome	10G	104M	9.9G	2%	/home
/dev/mapper/netwitness_vg00-varlog	10G	115M	9.9G	2%	/var/log
/dev/mapper/netwitness_vg00-nwhome	141G	1.2G	140G	1%	/var/netwitness
/dev/sdal	1014M	185M	830M	19%	/boot
tmpfs	1.6G	0	1.6G	0%	/run/user/1269
/dev/mapper/concentrator-root	30G	248M	30G	1%	/var/netwitness/concentrator
/dev/mapper/concentrator-sessiondb	811M	38M	773M	5%	/var/netwitness/concentrator/sessiondb
/dev/mapper/concentrator-metadb	7.2G	84M	7.2G	2%	/var/netwitness/concentrator/metadb
/dev/mapper/index-index	11G	111M	11G	1%	/var/netwitness/concentrator/index
/dev/mapper/index0-index	11G	111M	11G	1%	/var/netwitness/concentrator/index0
/dev/mapper/concentrator0-sessiondb	3.8G	60M	3.8G	2%	/var/netwitness/concentrator/sessiondb0
/dev/mapper/concentrator0-metadb	35G	277M	34G	1%	/var/netwitness/concentrator/metadb0

Appendix D - Sample Scenario for Extending for SASE Decoder

Storage

1. SSH into node-z (Admin Server). Identify the new model and update *model_name*, *additional_storage* attribute values in */root/.sase/sase-deployment-models.yml* and save the changes.

The below screen shot shows the updated storage configuration in */root/.sase/sase-deployment-models.yml*: *model_name: c1r12m60*, *additional_storage: true*. Additional storage disks are created for the updated *model_name* only when *additional_storage* is set to true.

```
# Define model name to configure host and storage
model_name: c1r12m60
additional_storage: true
```

2. Run the *nw-create-cloud-hybrid* script on node-z.

```
nw-create-cloud-hybrid --enable-cloud-sase
```

```
[root@124Admin026 12.4.0.0]# nw-create-cloud-hybrid --help
Usage:
nw-create-cloud-hybrid command [options]

Commands:
--enable-cloud-sase           Deploys NetWitness SASE focused assets to the cloud
--disable-cloud-sase          Undeploys NetWitness SASE focused assets from the cloud
--upgrade-overlay-network      Upgrades overlay network resources
--reissue-all-certs            reissues all overlay network certificates
--reissue-node-certs           reissues overlay network certificates for a specific node
    required parameter:
    --uuid                     UUID of NetWitness Node (see nw-manage -l)
--backup-cloud-nodes, -b       Backup configuration of all cloud nodes
--restore-cloud-node, -r       Restores configuration of specified Cloud node
    required parameter:
    --uuid                     UUID of NetWitness Cloud Node (see nw-manage -l)
--check-overlay-status, -c     Checks inter-connectivity of all nw-ppn overlay network hosts
--check-cert-status, -s       Checks overlay network certificate expiration status

Command Options:
--cloud-provider               Required: Destination cloud provider (gcp|aws)
--deployment-model             Optional Name of deployment model in template
                                defaults to pre-defined '(gcp|aws) default'
--cloud-key-path               Optional Cloud based Service Account key data path
                                defaults to .(gcp|aws) specific file
```

3. Login to Admin Server or UI. To list the all the block devices, Navigate to **Hosts -> Select the Decoder node -> Services -> Actions -> View -> Explore -> deviceappliance -> appliance -> Right Click -> Properties** drop-down -> **Select devlist -> Click Send.**
 - In this case, the model *c1r12m60* defines four disks for decoder service: *decoder*, *decodersmall*, *decoder0* and *decodersmall0*.
 - The response output for devlist returns both configured and unconfigured block devices. '*used=1*' indicate that the devices are configured. '*used=0*' indicates that the device(s) are unconfigured. No changes are required for configured blockdevice(s).
 - All existing configuration is preserved. When extending storage, existing disks as defined in *host-models.yml* (in this case *decoder* and *decodersmall*) are not recreated. Only new disks (*decoder0* and *decodersmall0*) are created.

- Repeat steps 2 through 6 from [Appendix B - Sample Scenario for Configuring SASE Decoder Storage](#) section to complete configuration of the unused (used=0) block devices.

Appendix E - Sample Scenario for Extending SASE Concentrator Storage

This sample scenario describes the Storage extension for concentrator configured with *c1r12m60* model. The supported extension models are *c1r6m30*, *c1r12m60* and *c1r23m120*.

Follow these steps to extend storage for Concentrator:

- SSH into node-z (Admin Server). Identify the new model and update *model_name* and *additional_storage* values in */root/.sase/sase-deployment-models.yml* and save the changes.

The below screen shot shows an example configuration updated with extended storage *model_name: c1r12m60*, *additional_storage: true*. If the *additional_storage* is not set to true, additional storage is not created even though the *model_name* is updated.

The below screen shot shows the updated storage configuration in */root/.sase/sase-deployment-models.yml*: *model_name: c1r12m60*, *additional_storage: true*. Additional storage disks are created for the updated *model_name* only when *additional_storage* is set to true.

```
# Define model name to configure host and storage
model_name: c1r12m60
additional_storage: true
```

- Run the *nw-create-cloud-hybrid* script on node-z.

```
nw-create-cloud-hybrid --enable-cloud-sase
```

- Login to SA UI. To list the all the block devices, Navigate to **Hosts** -> **Select the current Host (concentrator node)** -> **Services** -> **Actions** -> **View** -> **Explore** -> **deviceappliance** -> **appliance** -> **Right Click** -> **Properties** drop-down -> **Select devlist** -> Click **Send**.
 - The new disks corresponding to 'used=0' indicate that these disks are not configured.
 - During storage extension, all existing configuration is preserved.
- Repeat steps 2 through 7 from [Appendix C - Sample Scenario for Configuring SASE Concentrator Storage](#) section to complete configuration of the unused (used=0) block devices.

Appendix A. How NetWitness Platform Hosts Store Data

In most deployments, NetWitness Platform Decoders, Log Decoders, Concentrators, Archivers, and Hybrid hosts require external storage to house their data. Each host uses the external storage in different ways and with different expectations on throughput and performance of the external storage. Some hosts have a higher occurrence of sequential writes and some hosts have a higher occurrence of random reads and writes.

Decoder Hosts

Log Decoders and Network Decoders capture data and parse meta. The difference between these two hosts is in the type of data they capture:

- Log Decoder captures logs.
- Network Decoder captures packets.

Both Log Decoders and Network Decoders parse out meta data from the raw captured traffic. The meta data is then aggregated to a Concentrator for indexing. The host requires storage to house the raw payload data (raw packets or raw logs) and a cache for the meta extracted during data capture for Concentrator aggregation.

Your retention requirements is a key factor in determining the amount of storage you need for the raw packets or raw logs. In most deployments, you add storage over time based on increased retention requirements and increased capture rates. The storage for the raw data must support a high amount of sequential writes with random reads. Especially in the case of higher speed Network Decoder environments, it is recommended to have a minimum of two partitions exposed to the host to support the throttling between partitions for reads and writes.

The meta cache on a Decoder is generally fixed in size but you can expand it to support additional cache the possible loss of connectivity between the Decoder and a corresponding Concentrator. The meta cache must support a random IOPS rate for sustained writes from the Decoder of meta extracted and the corresponding reads from the Concentrator as meta is aggregated to a Concentrator.

Concentrator Host

A Concentrator aggregates and indexes the meta data from a Decoder. Both the meta and index storage needs are scaled based on your NetWitness Platform deployment retention requirements. Similar to raw data stored on the Decoders, you may need to increase the storage for both meta data and index data over time to meet your retention requirements.

The meta storage houses all meta data extracted from either a Network Decoder or Log Decoder. Although the ratio of how much meta is extracted may change, the expectations for performance against meta storage is the same for both packet capture and log capture environments. The meta storage must support a sustained amount of sequential writes with random reads of meta data.

The index storage houses the live index generated from the meta data aggregated to a Concentrator. The size of the index is directly related to the size of the meta store. In addition to supporting IOPS for sustained writes, the index also needs to support a much higher rate IOPS for reads than meta based on interactive queries run through analyst interaction and reports and alerts.

Archiver Host

The Archiver host requires a single partition for both meta and raw log storage. The storage pool deals primarily with sequential writes for long term data written from a Log Decoder or Network Decoder and random reads for reports and analysis.

Hybrid Hosts

A Hybrid hosts two or more services on a single host. For example:

- A Network Hybrid hosts both the Decoder and Concentrator services handling packets exclusively. It captures packet data and indexes this data to the Concentrator service. Expectations for storage performance match what is outlined for a dedicated Network Decoder host and dedicated Concentrator host.
- A Log Hybrid hosts both the Log Decoder and Concentrator services handling logs exclusively. It captures log data and indexes the data to a Concentrator service. Expectations for performance match what is outlined for a dedicated Log Decoder and dedicated Concentrator.
- An Endpoint Log Hybrid hosts the Endpoint Server, Log Decoder, Concentrator, Log Collector, and Endpoint Broker services. It collects and manages endpoint (host) data from Windows, Mac, and Linux hosts, collects log files and Windows logs from Windows hosts, and generates metadata to correlate endpoint data with sessions from other events sources, such as logs and packets.

Options for SAN Configurations

If you want to use a Storage Area Network (SAN) , use the same basic drive groups and partition organization that you use for the other NetWitness storage devices. Depending on the SAN configuration and overhead, SAN configurations may require more enclosures and drives to operate with the same performance as on PowerVault or DAC. When deciding whether to use SAN, PowerVault, or DAC, any additional overhead on the SAN will be important to determine the minimum storage required.

Performance Recommendations

NetWitness recommends that Packet and Log Decoders receive two LUNs or Block Devices, one for Packet data, the other for all other databases. This allows you to segregate the high-bandwidth Packet Database from the other databases so they do not compete for I/O bandwidth with other activity.

Concentrators require a separate SSD-based index volume for best performance. You must house this index volume on a different RAID group than the Concentrator Meta database volume, which you can stored on NL-SAS. Archivers can use a single large NL-SAS storage volume per appliance.

Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives

The encryptSedVd.py may fail to identify the SED Capable Virtual Drives when there is mix of both SED and NON-SED drives on the appliance. The below steps are applicable when both SED and NON-SED capable virtual drives exist on the host.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.

To list all the controllers on the appliance:

```
/opt/MegaRAID/percccli/percccli64 show | egrep -A3 'Model'
```

The first column (**Ctl**) lists out the controller index on the appliance. In this case, the controller '**0**' corresponds to '**PERC H740 Mini**' and controller '**1**' corresponds to '**PERC H840 Adaptor**'. The columns '**DGs**' and '**VDs**' displays the virtual drives and drive groups on the controller.

```
[root@116Decoder percccli]# /opt/MegaRAID/percccli/percccli64 show | egrep -A3 'Model'
Ctl Model          Ports PDs DGs DN0pt VDs VN0pt BBU sPR DS EHS AS0s Hlth
-----
0 PERCH740PMini    8 10 3 0 3 0 Opt On - N 0 Opt
1 PERCH840Adapter  8 12 1 0 1 0 Opt On - N 0 Opt
[root@116Decoder percccli]#
```

2. To enable the security on the 'PERC H740 (mini) Adaptor', for example, Controller '**0**', execute the following command:

```
/opt/MegaRAID/percccli/percccli64 /c0 set securitykey='<SOME_STRING_VALUE>'!'
keyid='< SOME_STRING_VALUE >'
```

Example:

```
/opt/MegaRAID/percccli/percccli64 /c0 set securitykey='Netwitness1!' keyid=1
'Netwitness1' is the securityKey and '1' is ID. Preserve both the Key and
keyID securely.
```

```
[root@116Decoder percccli]# /opt/MegaRAID/percccli/percccli64 /c0 set securitykey='Netwitness1!' keyid='netwitness'
Controller = 0
Status = Success
Description = None

Controller Properties :
=====

Ctrl Method Result
-----
0 set Key Success
```

3. Identify the correct Drive group (DG) / Virtual Drive (VD) corresponding to the SED Capable drives that you are trying to enable security.

```
/opt/MegaRAID/percccli/percccli64 /c0 /vall show | egrep -A5 'DG/VD'
```

Refer to first two and last column to identify the correct Drive Group (DG) / Virtual Drive (VD) correspond to the 6 SED enabled drives. On Series 6 appliances, there is only one DG/VD with **RAID6**. '**NAME**' column can be used to identify the VD/DG. In this case, the DG/VD is '**2**'. Using a combination of '**Type**', '**Name**' and '**Size**' columns (these are defined by the user when the VDs are created above).

```
[root@116Decoder percccli]# /opt/MegaRAID/percccli/percccli64 /c0 /vall show | egrep -A5 'DG/VD'
DG/VD TYPE State Access Consist Cache Cac sCC Size Name
-----
0/0 RAID1 Optl RW Yes RWBD - OFF 931.0 GB
1/1 RAID1 Optl RW Yes RWBD - OFF 1.818 TB
2/2 RAID6 Optl RW Yes RWBD - OFF 8.730 TB Virtual Disk 2
[root@116Decoder percccli]#
```

4. To turn on Security on the disk group (created out of the 6 SED Capable drives), execute the below command:

```
/opt/MegaRAID/percccli/percccli64 /c0 /d2 set security=on
[root@116Decoder percccli]# /opt/MegaRAID/percccli/percccli64 /c0 /d2 set security=on
Controller = 0
Status = Success
Description = Success
```

5. Get the Enclosure ID (**EID**) using on the controller '0'. In this case, it is '64'

```
/opt/MegaRAID/percccli/percccli64 /c0 /eall show
```

```
[root@116Decoder percccli]# /opt/MegaRAID/percccli/percccli64 /c0 /eall show
Controller = 0
Status = Success
Description = None

Properties :
=====

EID State Slots PD PS Fans TSs Alms SIM Port#      ProdID      VendorSpecific
-----
64 OK      10 10 0    0  0  0  1 00 & 00 x8 BP14G+EXP +
-----

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count

[root@116Decoder percccli]#
```

6. To confirm that the drives / Drive Groups (DG) are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as 'Yes' for drives in slots 4 (**s4**) through 9 (**s9**).

```
/opt/MegaRAID/percccli/percccli64 /c0 /e64/sall show all | egrep -i '
(Policies/Settings |SED Capable|Secured|SED Enabled) '
```

```
Drive /c0/e64/s0 Policies/Settings :
```

```
SED Capable = No
```

```
SED Enabled = No
```

```
Secured = No
```

```
Drive /c0/e64/s1 Policies/Settings :
```

```
SED Capable = No
```

```
SED Enabled = No
```

```
Secured = No
```

```
Drive /c0/e64/s2 Policies/Settings :
```

```
SED Capable = No
```

```
SED Enabled = No
```

```
Secured = No
```

```
Drive /c0/e64/s3 Policies/Settings :
```

```
SED Capable = No
```

```
SED Enabled = No
```

Secured = No
Drive /c0/e64/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c0/e64/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes

Appendix B. Encrypt a Series 6E or Series 7 Core or Hybrid Host (encryptSedVd.py)

NetWitness Series 6E Core and Hybrid hosts have Self-Encrypting Drives (SED). The `encryptSedVd.py` script:

- Validates that the Series 6E host has the correct setup for encryption.
- Encrypts unencrypted drives.

Note: For external storage devices such as PowerVault, refer to "[Configure Storage Using the REST API](#)" under "Using the REST API to Configure Storage" for instructions on how to encrypt their SED drives.

The following scenarios are examples of why you would use the `encryptSedVd.py` script.

- You want to know if a physical host has encryption. In this case, if the script determines that the device does not have encryption, it gives you the opportunity to encrypt it.
- You set up a device without encryption and you want to encrypt it.

You will find this script in the `rsa-sa-tools` directory for releases 12.5 and later. The following directory is for 12.5.

```
rsa-sa-tools-12.5.0.0-xxxx.noarch.rpm
```

The following procedure illustrates how to use the script.

1. Log in as `root`.
2. Change the directory to the `rsa-sa-tools` RPM base directory:

```
cd /opt/rsa/saTools/supportScript/
```

3. Execute the following command:

```
OWB_ALLOW_NON_FIPS=1 ./encryptSedVd.py
```

The script tells you if the disks are encrypted or not encrypted.

- If the drives are encrypted, the script displays the following message.
No unencrypted RAID virtual drives with SED physical drives found.
- If the drives are not encrypted, the script identifies the unencrypted drives as shown in the following example.

```
Detected unencrypted RAID Virtual Drives with SED Physical Disks
Please select the drives to encrypt
Navigation: <Tab><Up/Down Arrow> move vertical
<Esc> Quit, <Enter> Save, <Space> Select/Deselect, <A> Select All, <D> Deselect All

  ID VD  DG  RAID  SIZE  HBA
( ) 0  0  0   RAID1 1.1TB PERC H740P Mini
( ) 0  1  1   RAID1 2.2TB PERC H740P Mini
```

4. If the drives are not encrypted and you want to encrypt them:
 - a. Select the drives you want to encrypt with the space bar and press **Enter**.

The following prompt is displayed.

```
Please enter a passphrase for the PERC H740P Mini security key, minimum length 8 characters, maximum 32
The passphrase must contain a mix of lowercase, uppercase, numeric and non-alphanumeric characters
Optionally enter a key identifier, a default id will be created if not specified

Editing: <Backspace> clear cursor left, <Delete> clear cursor right
Navigation: <Tab><Up/Down Arrow> move vertical, <Left/Right Arrow> move horizontal
<Esc> quit without saving, <Enter> save, trailing spaces are ignored

Enter Passphrase:
[ ]

Verifv Passphrase:
[ ]

Key ID (optional):
[ ]
```

- b. In the **Enter Passphrase** text box, type the <passphrase>, for example nFreDaW\$792, and press **Tab**.
 - c. In the **Verify Passphrase** text box, re-enter passphrase again for validation.
 - d. In the **Key ID (optional)** text box, enter an optional ID string for the security key less than 256 characters or press Enter for none.

The following prompt is displayed.

```
The Passphrase for the security key *Must* be securely backed up in case of PERC adapter hardware
failure and/or replacement, without it the data on all encrypted disks will be unrecoverable.

Entered Passphrase('Quoted'): 'Testing$123'
Entered KeyId('Quoted'): '1'

( ) I understand the risks and have added the passphrase to my organization's permanent record
<Esc> Cancel, <Y> Acknowledge Backup, <D> Decline Backup, <Enter> Save
```

- e. Select <Y> and press **Enter** to confirm that you added the Passphrase.
 - f. Submit the following command string to verify that the SED drives are encrypted.

```
/opt/MegaRAID/percccli/percccli64 /c0 show more
```

The following information is displayed. You can see that all four SED drives are encrypted (that is, Y is displayed for each drive in the SED column).

```
Physical Drives = 4
```

```
PD LIST :
```

```
=====
```

EID:SlT	DID	State	DG	Size	Intf	Med	SED	PI	SeSz	Model	Sp
64:0	0	Onln	0	1.090 TB	SAS	HDD	Y	N	512B	ST1200MM0069	U
64:1	1	Onln	0	1.090 TB	SAS	HDD	Y	N	512B	ST1200MM0069	U
64:2	2	Onln	1	2.182 TB	SAS	HDD	Y	N	512B	ST2400MM0149	U
64:3	3	Onln	1	2.182 TB	SAS	HDD	Y	N	512B	ST2400MM0149	U

```
=====
```

Note: The SED Enabled and Secured label values are set to Yes, if the drives are SED enabled and secured.

To check the drives on controller 0 and enclosure 247 use the below command:

```
/opt/MegaRAID/perccli/perccli64 /c1 /e247/sall show all | egrep -i '
(Policies/Settings|SED Capable|Secured|SED Enabled) '
```

You will find detailed information on perccli commands in the Dell PowerEdge RAID Controller CLI Reference Guide (http://l4u-00.jinr.ru/pub/misc/h-w/LSI/dell-sas-hba-12gbps_reference-guide_en-us.pdf).

Enable SED on configured Drive Groups

Virtual Drives configured are SED Capable but are NOT SED Enabled.

To enable virtual drives or drive groups using PERC H840 Adaptors (External storage):

1. SSH to the appliance and run the below script to encrypt the virtual drive (on external storage).

Note: The encryptSedVd.py script turn on the SED feature only on Virtual Drives or Drive Groups on the PERC H840 Adaptors (external storage) and NOT on PERC H740 mini. Refer to [Enable Virtual Drives / Drive Groups - PERC H740 \(Mini\) Adaptors \(Internal storage\)](#) to enable SED on PERC H740 Mini .

```
OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py
```

2. Select the Virtual Drive and press **Enter**.
Passphrase screen is displayed.

3. Enter the Passphrase and press **Enter**.
For Example,

Passphrase : **Netwitness1!**

keyID: **netwitness**

Please enter a passphrase for the PERC H840 Adapter security key; minimum length 8 characters, maximum 32
The passphrase must contain a mix of lowercase, uppercase, numeric and non-alphanumeric characters
Optionally enter a key identifier, a default id will be created if not specified

Editing: <Backspace> clear cursor left, <Delete> clear cursor right
Navigation: <Tab><Up/Down Arrow> move vertical, <Left/Right Arrow> move horizontal
<Esc> quit without saving, <Enter> save, trailing spaces are ignored

Enter Passphrase:

Verify Passphrase:

Key ID (optional):

4. Acknowledge the message and Press **Enter** to Save.

```
The Passphrase for the security key *Must* be securely backed up in case of PERC adapter hardware failure and/or replacement, without it the data on all encrypted disks will be unrecoverable.
Entered Passphrase('Quoted'): 'Netwitness!!'
Entered KeyID('Quoted'): 'netwitness'

( ) I understand the risks and have added the passphrase to my organization's permanent record
<Esc> Cancel, <Y> Acknowledge Backup, <D> Decline Backup, <Enter> Save
```

5. Press any Key to Exit.

```
Successfully Encrypted All Selected RAID Virtual Drives
If you set a PERC controller security key passphrase or key ID,
Please be sure to add them to your organization's permanent record.
Press any key to exit
```

6. To confirm that the drives are SED Enabled and secured, run the following command and verify the SED Enabled and Secured returns Yes.

```
/opt/MegaRAID/percccli/percccli64 /c1 /e247/sall show all | egrep -i '
(Policies/Settings|SED Capable|Secured|SED Enabled) '
```

```
Drive /c1/e247/s0 Policies/Settings :
```

```
SED Capable = Yes
```

```
SED Enabled = Yes
```

```
Secured = Yes
```

```
Drive /c1/e247/s1 Policies/Settings :
```

```
SED Capable = Yes
```

SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s2 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s3 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s4 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s10 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s11 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes

Secured = Yes

Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage)

You can enable the SED capability on the Virtual Drive or Drive Groups created out of on-board SED capable drives (in slots 4 through 9 – total of 6 drives) using the percli64 utility. You cannot use `/opt/rsa/saTools/supportScript/encryptSedVd.py` to turn on Security on the Virtual drives on the PERC H740 (mini) Adaptor.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.

To list all the controllers on the appliance run the following command:

```
/opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
```

The first column (**Ctl**) lists out the controller index on the appliance. In this case, the controller **0** corresponds to **PERC H740 Mini** and controller **1** corresponds to **PERC H840 Adaptor**. The columns **DGs** and **VDs** displays the virtual drives and drive groups on the controller.

2. To enable the security on the **PERC H740 (mini) Adaptor**, for example, Controller **0**, run the following command:

```
/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='<String>'!  
keyid='<String>'
```

Example:

```
/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!'!  
keyid='netwitness'
```

'Netwitness1' is the securityKey and 'netwintess' is ID.

Make a note of both the Key and keyID securely.

3. Identify the correct Drive group (DG) or Virtual Drive (VD) corresponding to the SED Capable drives that you want to enable security.

```
/opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
```

Check the first two and last column to identify the correct Drive Group / Virtual Drive correspond to the 6 SED enabled drives that are SED Capable. On Series 6 appliances, there is only one DG or VD with **RAID6** type. Name column can be used to identify the VD or DG. In this case, the DG or VD is **2**. Using a combination of **Type**, **Name** and **Size** columns (these are defined when you created VDs above).

4. To turn on Security on the disk group (created out of the 6 SED Capable drives) for **decodersmall** volume group, run the below command:

```
/opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
```

5. Get the Enclosure ID (**EID**) using on the controller **0**. In this case, it is **64**

```
/opt/MegaRAID/perccli/perccli64 /c0 /eall show
```

6. To confirm that the drives or drive groups are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as **Yes** for drives in slots 4 (s4) through 9 (s9).

```
/opt/MegaRAID/percccli/percccli64 /c0 /e64/sall show all | egrep -i '  
(Policies/Settings |SED Capable|Secured|SED Enabled) '
```

Drive /c0/e64/s0 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s1 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s2 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s3 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s4 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s5 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s6 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s7 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s8 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s9 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Enable SED on configured Virtual Drives/ Drive Groups on Power Vault (PERC 840)

Enable Virtual Drives / Drive Groups - PERC H840 Adaptors

Note: The virtual disk created in *Configure Block Devices for PowerVaults* section in [Configure Drive Pack\(s\)](#) is SED capable but NOT SED Enabled.

1. To enable, SSH into the appliance and run the below script to encrypt the VD (on external storage).
`OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py`

Note: The encryptSedVd.py script turn on the SED feature only on Virtual Drives or Drive Groups on the PERC H840 Adaptors (external storage) and NOT on PERC H740 mini. Refer to **Enable Virtual Drives / Drive Groups - PERC H740 (Mini) Adaptors (Internal storage)** to enable SED on PERC H740 Mini

`OWB_ALLOW_NON_FIPS=true /opt/rsa/saTools/supportScript/encryptSedVd.py`

```
Detected unencrypted RAID Virtual Drives with SED Physical Disks
Please select the drives to encrypt
Navigation: <Tab><Up/Down Arrow> move vertical
<Esc> Quit, <Enter> Save, <Space> Select/Deselect, <A> Select All, <D> Deselect All

ID VD DG RAID SIZE HBA
(X) 1 0 0 RAID6 106.9TB PERC H840 Adapter
```

2. Select both the **Virtula Disks** and press **Enter**.
The Passphrase screen is displayed.

```
Please enter a passphrase for the PERC H840 Adapter security key, minimum length 8 characters, maximum 32
The passphrase must contain a mix of lowercase, uppercase, numeric and non-alphanumeric characters
Optionally enter a key identifier, a default id will be created if not specified

Editing: <Backspace> clear cursor left, <Delete> clear cursor right
Navigation: <Tab><Up/Down Arrow> move vertical, <Left/Right Arrow> move horizontal
<Esc> quit without saving, <Enter> save, trailing spaces are ignored

Enter Passphrase:


Verify Passphrase:


Key ID (optional):

```

3. Enter the Passphrase and press **Enter**.

For example,

Passphrase : **Netwitness1!**

keyID: **netwitness**

Please enter a passphrase for the PERC H840 Adapter security key, minimum length 8 characters, maximum 32
The passphrase must contain a mix of lowercase, uppercase, numeric and non-alphanumeric characters.
Optionally enter a key identifier, a default id will be created if not specified

Editing: <Backspace> clear cursor left, <Delete> clear cursor right
Navigation: <Tab><Up/Down Arrow> move vertical, <Left/Right Arrow> move horizontal
<Esc> quit without saving, <Enter> save, trailing spaces are ignored

Enter Passphrase:

Netwitness1!

Verify Passphrase:

Netwitness1!

Key ID (optional):

netwitness

4. Acknowledge the message and Press **Enter** to Save.


The Passphrase for the security key *Must* be securely backed up in case of PERC adapter hardware failure and/or replacement, without it the data on all encrypted disks will be unrecoverable.

Entered Passphrase('Quoted'): 'Netwitness1!'

Entered Keyid('Quoted'): 'netwitness'

() I understand the risks and have added the passphrase to my organization's permanent record
<Esc> Cancel, <Y> Acknowledge Backup, <D> Decline Backup, <Enter> Save

5. Press any Key to Exit.



```
Successfully Encrypted All Selected RAID Virtual Drives
If you set a PERC controller security key passphrase or key ID,
Please be sure to add them to your organization's permanent record.
Press any key to exit
```

6. To confirm that the drives are SED Enabled and secured, run the below command and verify the SED Enabled and Secured returns Yes.

```
/opt/MegaRAID/perccli/perccli64 /c1 /e247/sall show all | egrep -i '
(Policies/Settings|SED Capable|Secured|SED Enabled) '
```

Drive /c1/e247/s0 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s1 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s2 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s3 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c1/e247/s4 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

```
Secured = Yes
Drive /c1/e247/s5 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s6 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s7 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s8 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s9 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s10 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
Drive /c1/e247/s11 Policies/Settings :
SED Capable = Yes
SED Enabled = Yes
Secured = Yes
```

Enable Security on SED Capable Drive groups on Host with a mix of SED and NON SED Drives

The `encryptSedVd.py` may fail to identify the SED Capable Virtual Drives when there is mix of both SED and NON-SED drives on the appliance. The below steps are applicable when both SED and NON-SED capable virtual drives exist on the host.

1. SSH to the appliance and enable security on the PERC H740 (mini) Adaptor. The controller number for this adaptor is **0**. The PERC H840 Adaptor is shown as **1**.

To list all the controllers on the appliance:

```
/opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
```

The first column (**Ctl**) lists out the controller index on the appliance. In this case, the controller '**0**' corresponds to '**PERC H740 Mini**' and controller '**1**' corresponds to '**PERC H840 Adaptor**'. The columns '**DGs**' and '**VDs**' displays the drive groups and virtual drives on the controller.

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 show | egrep -A3 'Model'
Ctl Model
-----
0 PERCH740PMini 8 10 3 0 3 0 Opt On - N 0 Opt
1 PERCH840Adaptor 8 12 1 0 1 0 Opt On - N 0 Opt
[root@116Decoder perccli]#
```

2. To enable the security on the 'PERC H740 (mini) Adaptor' i.e Controller '**0**', execute the following command:

```
/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='<SOME_STRING_VALUE>' !'
keyid='< SOME_STRING_VALUE >'
```

Example:

```
/opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid=1
```

'Netwitness1' is the securityKey and '1' is ID. Preserve both the Key and keyID securely.

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 set securitykey='Netwitness1!' keyid='netwitness'
Controller = 0
Status = Success
Description = None

Controller Properties :
=====

-----
Ctrl Method Result
-----
0 set Key Success
-----
```

3. Identify the correct Drive group (DG) / Virtual Drive (VD) corresponding to the SED Capable drives that we are trying to enable security.

```
/opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
```

Refer to first two and last column to identify the correct Drive Group (DG) / Virtual Drive (VD) correspond to the 6 SED enabled drives. On Series 6 appliances, there is only one DG/VD with **RAID6**. '**Name**' column can be used to identify the VD/DG. In this case, the DG/VD is '**2**'. Using a combination of '**Type**', '**Name**' and '**Size**' columns (these were defined by the user when the VDs are created above).

```
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /vall show | egrep -A5 'DG/VD'
DG/VD TYPE State Access Consist Cache Cac sCC Size Name
-----
0/0 RAID1 Optl RW Yes RWBD - OFF 931.0 GB
1/1 RAID1 Optl RW Yes RWBD - OFF 1.818 TB
2/2 RAID6 Optl RW Yes RWBD - OFF 8.730 TB Virtual Disk 2
-----
[root@116Decoder perccli]#
```

- To turn on Security on the disk group (created out of the 6 SED Capable drives), execute the below command:

```
/opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /d2 set security=on
Controller = 0
Status = Success
Description = Success
```

- Get the Enclosure ID (EID) using on the controller '0'. In this case, it is '64'

```
/opt/MegaRAID/perccli/perccli64 /c0 /eall show
[root@116Decoder perccli]# /opt/MegaRAID/perccli/perccli64 /c0 /eall show
Controller = 0
Status = Success
Description = None

Properties :
=====

EID State Slots PD PS Fans TSs Alms SIM Port#      ProdID      VendorSpecific
-----
64 OK      10 10  0   0   0   0   1 00 & 00 x8 BP14G+EXP +
-----

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSs-Temperature sensor count |Alms-Alarm count |SIM-SIM Count
```

- To confirm that the drives / Drive Groups (DG) are **SED Enabled** and **Secured**, run the below command and verify the **SED Capable**, **Secured**, **SED Enabled** flags are set as 'Yes' for drives in slots 4 (s4) through 9 (s9).

```
/opt/MegaRAID/perccli/perccli64 /c0 /e64/sall show all | egrep -i '
(Policies/Settings |SED Capable|Secured|SED Enabled) '
```

Drive /c0/e64/s0 Policies/Settings :

```
SED Capable = No
SED Enabled = No
Secured = No
```

Drive /c0/e64/s1 Policies/Settings :

```
SED Capable = No
SED Enabled = No
Secured = No
```

Drive /c0/e64/s2 Policies/Settings :

```
SED Capable = No
SED Enabled = No
Secured = No
```

Drive /c0/e64/s3 Policies/Settings :

SED Capable = No

SED Enabled = No

Secured = No

Drive /c0/e64/s4 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s5 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s6 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s7 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s8 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Drive /c0/e64/s9 Policies/Settings :

SED Capable = Yes

SED Enabled = Yes

Secured = Yes

Note: The steps for enabling encryption on Series 7 hardware (with or without external storage configuration i.e MD2412) are like Series 6 steps described in this section. On Series 7 hardware:

1. The internal HBA is PERC H965i and external HBA is PERC H965e.
2. Perccli64 is not supported on Series 7 hardware. Use */opt/MegaRAID/perccli/perccli2*. Substitute perccli64 with perccli2 when executing the perccli commands on Series 7 hardware.

Appendix C. Troubleshooting

This section contains instructions on how to resolve various storage tasks using the REST API.

Reconfigure Pre-Configured DAC Attached to Decoder Using REST API

This scenario covers how to reconfigure a DAC using the REST API that was configured using another tool and clear any pre-existing data (if no longer need or backed up to another storage device).

The following information describes the state of the host and storage hardware prior to the attempt to reconfigure the storage devices using the REST API.

When the DAC was added, it had old data and was configured (but not using the REST API). This prevented the REST API from executing the `raidNew` command and returned the `Physical disk does not have appropriate attributes` error message.

The following steps describe the scenario and with its resolution.

1. From the Decoder Linux console (or SSH to Decoder), submitted the following command string.
`/opt/MegaRAID/perccli/perccli64 /c2/fall del`
You will find detailed information on perccli commands in the **Dell PowerEdge RAID Controller CLI Reference Guide** (https://topics-cdn.dell.com/pdf/dell-sas-hba-12gbps_reference-guide_en-us.pdf).
This deleted all foreign configuration from controller 2 and cleared all data from the DAC.
2. Tried to partition the DAC, but the `partNew` command failed because that information was already defined on the DAC. `partNew` displayed that you must use one an available device, but `devList` showed it in use.
3. Assuming that the partitions were defined, tried to allocate the storage devices, but this did not work because the DAC was not mounted.
4. Tried to mount the DAC from the command line, but received `mount failed: structure needs to be cleaned` error message.
5. There was no data that needed to be preserved on the DAC, so submitted the following command strings to clean the structure.
`mkfs.xfs -f /dev/decoder0/packetdb`
`mkfs.xfs -f /dev/decoder1/packetdb`
6. Mounted devices to their appropriate locations in `/var/netwitness/decoder`.
7. Completed the remainder of the applicable steps as described in [Configure Storage Using the REST API](#) to reconfigure the DAC

Appendix D. Sample Storage Configuration Scenarios for 15-Drive DACs

This appendix illustrates the following example of how to configure storage on two non-encrypted 15-drive DAC external storage devices.

- [Configure Storage for Archiver](#)
- [Configure Storage for Network \(Packet\) Decoder](#)
- [Configure Storage for Network Concentrator](#)
- [Configure Storage for Log Decoder Hybrid](#)

Configure Storage for Archiver

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for an Archiver physical host.

1. Execute the `raidList` command.
 - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
You should see the following information.
In Use: FALSE
Devices: <empty>
 - b. Verify the Drive Count, Size, and Vendor.
The following example illustrates what you should see before you create a RAID array.

The screenshot displays the NetScout Systems management interface. On the left is a navigation tree with categories like archiver, connections, deviceappliance, and appliance. The 'appliance' category is selected, showing a list of sub-items including connections, logs, rest, services, storedproc, sys, and users. The main panel on the right shows the 'Properties for NWHOST2100 – Archiver (ARCHIVER)/deviceappliance/appliance.' The 'raidList' command is entered in the 'Parameters' field. Below the command, the 'Response Output' is displayed, showing details for two controllers. The first controller (Controller 0, Enclosure 32) is in use and has 2 drives. The second controller (Controller 1, Enclosure 0) is not in use and has 15 drives. The 'In Use' status and 'Devices' for the second controller are highlighted in yellow.

Controller	Enclosure	Vendor	Model	In Use	Drives	Devices
Controller 0	Enclosure 32	DP	BP13G+EXP	true	931.511 GB x 2	sda, sdb
Controller 1	Enclosure 0	EMC	ESES Enclosure	false	3.637 TB x 15	

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.
`controller=1 enclosure=0 scheme=archiver commit=1`

The following example illustrates what you should see after you create a RAID array.

Properties for NWHOST2100 – Archiver (ARCHIVER)/deviceappliance/appliance.

raidNew Parameters controller=1 enclosure=0 scheme=archiver commit=1

Message Help

enclosure - <uint32, (enum-one:32.0)> Enclosure number of the shelf to clear
 scheme - <string, (enum-one:decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid)> Type of RAID volumes to allocate
 preferSecure - <bool, optional, (bool:0,1,yes,no,true,false,on,off)> Prefer creation of a secure array given compatible physical drives and a controller with a security key set
 commit - <bool, optional> commit changes

Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r6 drives=0:0,0:1,0:2,0:3,0:4,0:5,0:6,0:7,0:8,0:9,0:10,0:11,0:12,0:13,0:14 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
```

- Execute the `raidList` command to verify the new RAID array.

You should now see the following information.

In Use: TRUE

Devices: <device> (for example, `sdc`)

Properties for NWHOST2100 – Archiver (ARCHIVER)/deviceappliance/appliance.

raidList Parameters

Message Help

list drive shelves attached to this appliance
 security.roles: appliance.manage

Response Output

```
Controller 0, Enclosure 32
Vendor: DP
Model: BP13G+EXP
In Use: true
Drives: 931.511 GB x 2
1.818 TB x 2
Devices: sda
sdb

Controller 1, Enclosure 0
Vendor: EMC
Model: ESES Enclosure
In Use: true
Drives: 3.637 TB x 15
Devices: sdc
```

- Execute the `partNew` command with the following parameters to create partitions and mount points in the `etc/fstab` file.
`name=<device>` (for example, `sdc`) `service=archiver` `volume=archiver` `commit=1`
- Execute the `srvAlloc` command with the following parameters to allocate the space to the archiver service. This adds storage to the archiver service configuration and restarts the service every time it is executed.
`service=archiver` `volume=archiver0` `commit=1`

Properties for NWHOST2100 – Archiver (ARCHIVER)/device/appliance/appliance.

srvAlloc ▼ Parameters `service=archiver volume=archiver0 commit=1`

Message Help

service - <string, (enum-one:archiver|concentrator|decoder|logdecoder)> service that will use storage
 volume - <string, (enum-one:archiver0|netwitness_vg00)> volume group name
 commit - <bool, optional> commit changes

Change Service | NWHOST2100 - Archiver | System ⊙

Start Aggregation | Stop Aggregation | Host Tasks | Shutdown Service

Archiver Service Information

Name NWHOST2100 (Archiver)
 Version 11.3.0.0 (Rev null)
 Memory Usage 30016 KB (0.02% of 126 GB)
 CPU 0%
 Running Since 2019-Jun-12 13:12:17
 Uptime 1 minute 10 seconds
 Current Time 2019-Jun-12 13:13:27

6. Confirm the “Hot Storage” in “Data Retention”.

Change Service | NWHOST2100 - Archiver | Config ⊙

General | **Data Retention** | Files | Appliance Service Configuration

Configure the rollover criteria for removing database records from primary storage using an age-based threshold, and schedule the timing for checking if t

1. Configure hot, warm and cold storage
2. Configure collections
3. Define retention rules

Total Hot Storage **47.29 TB** ⚙️
 Total Warm Storage **Not Configured** ⚙️
 Cold Storage **Not Configured** ⚙️

1 Mount Point

Collections

+ - ✎ | ↻

☐	Collection	Usage / Hot Storage	Usage / Warm Storage	Cold Storage	Retention
☐	default	0 B / 44.93 TB (95%)	Disabled	☐	No Limit
Total Storage		0 B / 44.93 TB	0 B / 0 B		

Retention Rules

+ - ✎ | ↻ | ⬆ Move Up ⬇ Move Down | Apply ↻ Revert

☐	Order ^	Rule Name	Condition
☐		default	*

Configure Storage for Network (Packet) Decoder

The following scenario configures storage on two, non-encrypted, 15-Drive DACs for a Network Decoder for 10G Capture physical host.

1. Execute the `raidList` command.

- a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

In Use: FALSE

Devices: <empty>

- b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

The screenshot shows a web interface for the 'raidList' command. At the top, there is a dropdown menu set to 'raidList' and a 'Parameters' input field. Below this is a blue 'Send' button. Under the button is a 'Message Help' section with the text: 'list drive shelves attached to this appliance' and 'security.roles: appliance.manage'. The main part of the screenshot is the 'Response Output' section, which displays the following information for two controllers:

```
Controller 1 at PCI Address 3b:00.0, Enclosure 231, SCSI Channel 2
Vendor: EMC
Model: ESES Enclosure
In Use: false
Drives: 3.638 TB HDD x 15
Devices:

Controller 1 at PCI Address 3b:00.0, Enclosure 239, SCSI Channel 2
Vendor: EMC
Model: ESES Enclosure
In Use: false
Drives: 3.638 TB HDD x 15
Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

- Parameters for the first enclosure:

```
controller=1 enclosure=231 scheme=decoder-hotspare commit=1
```

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

raidNew Parameters controller=1 enclosure=231 scheme=decoder-hotspare commit=1

Message Help

allocate RAID devices in a drive shelf

security.roles: appliance.manage

parameters:

controller - <uint32, (enum-one:The value must be one of the following: 0,1)> Controller the shelf is attached to

Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r1 drives=231:0,231:1 ra Strip=128
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86_64

Controller = 1

Status = Success

Description = Add VD Succeeded.

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=231:2,231:3,231:4,231:5,231:6,231:7,231:8,231:9,231:10,231:11,231:12,231:13 ra Strip=128
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86_64

Controller = 1

Status = Success

Description = Add VD Succeeded.

```
/opt/MegaRAID/perccli/perccli64 /c1 /e231 /s14 add hotsparedrive
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86_64

Controller = 1

Status = Success

Description = Add Hot Spare Succeeded.

- Parameters for the second enclosure:

```
controller=1 enclosure=239 scheme=decoder-hotspare commit=1
```

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

raidNew Parameters controller=1 enclosure=239 scheme=decoder-hotspare commit=1

Message Help

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 64|231|239}> Enclosure number of the shelf to clear. Required if the controller is attached.

scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-decoder|expansion|decoder-hotspare|logdecoder-hotspare}> Type of RAID volumes to allocate

Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r1 drives=239:0,239:1 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.
```

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=239:2,239:3,239:4,239:5,239:6,239:7,239:8,239:9,239:10,239:11,239:12,239:13 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.
```

```
/opt/MegaRAID/perccli/perccli64 /c1 /e239 /s14 add hotsparedrive
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add Hot Spare Succeeded.
```

- Use the `raidList` command to display block devices for enclosures so you can verify In Use: TRUE.

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

raidList Parameters

Message Help

list drive shelves attached to this appliance
security.roles: appliance.manage

Response Output

```

1.819 TB HDD x 2
Devices: sda /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:0:0
        sdb /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 231, SCSI Channel 2
Vendor: EMC
Model: ESES Enclosure
In Use: true
Drives: 3.638 TB HDD x 15
Devices: sdc /dev/disk/by-path/pci-0000:3b:00.0-scsi-0:2:0:0
        sdd /dev/disk/by-path/pci-0000:3b:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 239, SCSI Channel 2
Vendor: EMC
Model: ESES Enclosure
In Use: true
Drives: 3.638 TB HDD x 15
Devices: sde /dev/disk/by-path/pci-0000:3b:00.0-scsi-0:2:2:0
        sdf /dev/disk/by-path/pci-0000:3b:00.0-scsi-0:2:3:0

```

4. Use `devlist` to view the new block devices and their sizes. The new devices are highlighted in red.

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

devlist Parameters

Message Help

list storage devices
security.roles: appliance.manage

Response Output

```

sda:vendor=DELL model="PERC H740P Mini" size="931 GB" used=1
sdb:vendor=DELL model="PERC H740P Mini" size="1.82 TB" used=1
sdc:vendor=DELL model="PERC H840 Adp" size="3.64 TB" used=0
sdd:vendor=DELL model="PERC H840 Adp" size="40.02 TB" used=0
sde:vendor=DELL model="PERC H840 Adp" size="3.64 TB" used=0
sdf:vendor=DELL model="PERC H840 Adp" size="40.02 TB" used=0

```

5. SSH to the Network Decoder and use the `lsblk` command to confirm the block device sizes. The smaller block devices are always allocated to decodersmall and larger devices are allocated to

decoder volume.

```
[root@~:Decoder ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda                                  8:0      0  931G  0 disk
├─sda1                              8:1      0    1G  0 part  /boot
├─sda2                              8:2      0  930G  0 part
│   ├─netwitness_vg00-root          253:0      0   30G  0 lvm    /
│   ├─netwitness_vg00-swap          253:1      0    4G  0 lvm    [SWAP]
│   ├─netwitness_vg00-nwhome        253:4      0   2.7T  0 lvm    /var/netwitness
│   └─netwitness_vg00-varlog        253:5      0   10G  0 lvm    /var/log
└─netwitness_vg00-usrhome          253:6      0   10G  0 lvm    /home
sdb                                  8:16      0  1.8T  0 disk
├─sdb1                              8:17      0  1.8T  0 part
└─netwitness_vg00-nwhome          253:4      0   2.7T  0 lvm    /var/netwitness
sdc                                  8:32      0   3.7T  0 disk
sdd                                  8:48      0   40T  0 disk
sde                                  8:64      0   3.7T  0 disk
sdf                                  8:80      0   40T  0 disk
```

Note: When configuring for 10g capture, use decoder-hotspare for both the enclosures for performance reasons. For non 10g captures, use decoder-hotspare for the first enclosure and packet-expansion for the second enclosure.

- Execute the `partNew` command to create the **decodersmall** partition first (decoder dir, index, metadb, sessiondb) (First Enclosure, SDC, SDD) with the following parameters.

name=sdc service=decoder volume=decodersmall commit=1

partNew Parameters name=sdc service=decoder volume=decodersmall commit=1

Message Help

name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
 service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
 commit - <bool, optional> commit changes

Response Output

```
Logical volume "decoroot" created.
/sbin/mkfs.xfs /dev/decodersmall/decoroot
meta-data=/dev/decodersmall/decoroot isize=512  agcount=4, agsize=655360 blks
=          sectsz=512  attr=2, projid32bit=1
=          crc=1      finobt=0, sparse=0
data =          bsize=4096  blocks=2621440, imaxpct=25
=          sunit=0   swidth=0 blks
naming   =version 2      bsize=4096  ascii-ci=0 ftype=1
log      =internal log  bsize=4096  blocks=2560, version=2
=          sectsz=512  sunit=0 blks, lazy-count=1
realtime =none          extsz=4096  blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder
/bin/mount /var/netwitness/decoder
/sbin/lvcreate -y -n index -L 30G decodersmall
Logical volume "index" created.
/sbin/mkfs.xfs /dev/decodersmall/index
meta-data=/dev/decodersmall/index isize=512  agcount=4, agsize=1966080 blks
=          sectsz=512  attr=2, projid32bit=1
=          crc=1      finobt=0, sparse=0
data =          bsize=4096  blocks=7864320, imaxpct=25
```

```
[root@NWHOST2000 ~]# df -h
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/netwitness_vg00-root          30G   2.5G   28G   9% /
devtmpfs                                 63G     0   63G   0% /dev
tmpfs                                     63G   12K   63G   1% /dev/shm
tmpfs                                     63G   26M   63G   1% /run
tmpfs                                     63G     0   63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-nwhome        2.7T   98M   2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-varlog         10G   49M   10G   1% /var/log
/dev/mapper/netwitness_vg00-usrhome        10G   33M   10G   1% /home
/dev/sda1                                1014M   88M   927M   9% /boot
tmpfs                                     13G     0   13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot         10G   33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index            30G   33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb        600G   33M   600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb           6.7T   33M   6.7T   1% /var/netwitness/decoder/metadb
[root@NWHOST2000 ~]#
```

- Execute the `partNew` command to create the decoder volume (packetdb) (First Enclosure, SDC, SDD) with the following parameters.

name==sdd service=decoder volume=decoder commit=1

partNew name=sdd service=decoder volume=decoder commit=1

Message Help

name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
 service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
 commit - <bool, optional> commit changes

Response Output

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f decoder /dev/sdd1
Volume group "decoder" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder/packetdb
meta-data=/dev/decoder/packetdb isize=512  agcount=41, agsize=268435455 blks
        =           sectsz=512  attr=2, projid32bit=1
        =           crc=1      finobt=0, sparse=0
data      =           bsize=4096  blocks=10742791168, imaxpct=5
        =           sunit=0    swidth=0 blks
naming    =version 2           bsize=4096  ascii-ci=0 ftype=1
log       =internal log      bsize=4096  blocks=521728, version=2
        =           sectsz=512  sunit=0 blks, lazy-count=1
realtime  =none              extsz=4096  blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb
/bin/mount /var/netwitness/decoder/packetdb
```

```
[root@netwitness-decoder ~]# df -hP
```

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	63G	0	63G	0%	/dev
tmpfs	63G	60K	63G	1%	/dev/shm
tmpfs	63G	11M	63G	1%	/run
tmpfs	63G	0	63G	0%	/sys/fs/cgroup
/dev/mapper/netwitness_vg00-root	30G	3.9G	27G	13%	/
/dev/sda1	1014M	89M	926M	9%	/boot
/dev/mapper/netwitness_vg00-varlog	10G	48M	10G	1%	/var/log
/dev/mapper/netwitness_vg00-nwhome	2.7T	609M	2.7T	1%	/var/netwitness
/dev/mapper/netwitness_vg00-usrhome	10G	33M	10G	1%	/home
tmpfs	13G	0	13G	0%	/run/user/0
/dev/mapper/decodersmall-decoroot	10G	33M	10G	1%	/var/netwitness/decoder
/dev/mapper/decodersmall-index	30G	33M	30G	1%	/var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb	600G	34M	600G	1%	/var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb	3.1T	34M	3.1T	1%	/var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb	41T	34M	41T	1%	/var/netwitness/decoder/packetdb

In this example, the below partitions (highlighted in Yellow) are created on sdc and sdd (Enclosure 231).

```
[root@netwitness-decoder ~]# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	931G	0	disk	
└─sda1	8:1	0	1G	0	part	/boot
└─sda2	8:2	0	930G	0	part	
└─netwitness_vg00-root	253:0	0	30G	0	lvm	/
└─netwitness_vg00-swap	253:1	0	4G	0	lvm	[SWAP]
└─netwitness_vg00-nwhome	253:4	0	2.7T	0	lvm	/var/netwitness
└─netwitness_vg00-varlog	253:5	0	10G	0	lvm	/var/log
└─netwitness_vg00-usrhome	253:6	0	10G	0	lvm	/home
sdb	8:16	0	1.8T	0	disk	
└─sdb1	8:17	0	1.8T	0	part	
└─netwitness_vg00-nwhome	253:4	0	2.7T	0	lvm	/var/netwitness
sdc	8:32	0	3.7T	0	disk	
└─sdc1	8:33	0	3.7T	0	part	
└─decodersmall-decoroot	253:7	0	10G	0	lvm	/var/netwitness/decoder
└─decodersmall-index	253:8	0	30G	0	lvm	/var/netwitness/decoder/index
└─decodersmall-sessiondb	253:9	0	600G	0	lvm	/var/netwitness/decoder/sessiondb
└─decodersmall-metadb	253:10	0	3T	0	lvm	/var/netwitness/decoder/metadb
sdd	8:48	0	40T	0	disk	
└─sdd1	8:49	0	40T	0	part	
└─decoder-packetdb	253:11	0	40T	0	lvm	/var/netwitness/decoder/packetdb

At this point, you add the second DAC enclosure.

- Execute the `partNew` command to create the decodersmall partition first (Second Enclosure, SDE, SDF) with the following parameters.

```
name=sde service=decoder volume=decodersmall commit=1
```

Properties for 11mtlnxnwpacket01 - Decoder (DECODER) /deviceappliance/appliance.

partNew name=sde service=decoder volume=decodersmall commit=1

Message Help

name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
 service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
 commit - <bool, optional> commit changes

Response Output

```
/sbin/parted -s /dev/sde mklabel gpt
/sbin/parted -s -a optimal /dev/sde mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sde1
Physical volume "/dev/sde1" successfully created.
/sbin/vgcreate -f decodersmall0 /dev/sde1
Volume group "decodersmall0" successfully created
/sbin/lvcreate -y -n index -L 30G decodersmall0
Logical volume "index" created.
/sbin/mkfs.xfs /dev/decodersmall0/index
meta-data=/dev/decodersmall0/index isize=512  agcount=4, agsize=1966080 blks
       =      sectsz=512   attr=2, projid32bit=1
       =      crc=1      finobt=0, sparse=0
data    =      bsize=4096   blocks=7864320, imaxpct=25
       =      sunit=0      swidth=0 blks
naming  =version 2        bsize=4096   ascii-ci=0 ftype=1
log     =internal log     bsize=4096   blocks=3840, version=2
       =      sectsz=512   sunit=0 blks, lazy-count=1
realtime =none           extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/index0
/bin/mount /var/netwitness/decoder/index0
```

9. Execute the partNew command to create the packetdb decoder volume (Second Enclosure SDE, SDF) with the following parameters.

```
name=sdf service=decoder volume=decoder commit=1
```

partNew

Parameters

name=sdf service=decoder volume=decoder commit=1

Message Help

name - <string, {enum-one:sdc,sdd,sde,sdf}> block device name
service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create
commit - <bool, optional> commit changes

Response Output

```

/sbin/parted -s /dev/sdf mklabel gpt
/sbin/parted -s -a optimal /dev/sdf mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdf1
Physical volume "/dev/sdf1" successfully created.
/sbin/vgcreate -f decoder0 /dev/sdf1
Volume group "decoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder0
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder0/packetdb
meta-data=/dev/decoder0/packetdb isize=512  agcount=41, agsize=268435455 blks
=               sectsz=512   attr=2, projid32bit=1
=               crc=1        finobt=0, sparse=0
data  =               bsize=4096   blocks=10742791168, imaxpct=5
=               sunit=0       swidth=0 blks
naming  =version 2           bsize=4096   ascii-ci=0 ftype=1
log    =internal log        bsize=4096   blocks=521728, version=2
=               sectsz=512   sunit=0 blks, lazy-count=1
realtime =none             extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb0
/bin/mount /var/netwitness/decoder/packetdb0

```

```

[root@netwitness-Decoder ~]# df -hP
Filesystem                                Size  Used Avail Use% Mounted on
devtmpfs                                  63G   0    63G   0% /dev
tmpfs                                     63G   60K   63G   1% /dev/shm
tmpfs                                     63G   11M   63G   1% /run
tmpfs                                     63G   0    63G   0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root           30G   3.9G   27G   13% /
/dev/sda1                                 1014M   89M   926M   9% /boot
/dev/mapper/netwitness_vg00-varlog         10G    48M   10G   1% /var/log
/dev/mapper/netwitness_vg00-nwhome        2.7T   609M   2.7T   1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome       10G    33M   10G   1% /home
tmpfs                                      13G   0    13G   0% /run/user/0
/dev/mapper/decodersmall-decoroot         10G    33M   10G   1% /var/netwitness/decoder
/dev/mapper/decodersmall-index            30G    33M   30G   1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb       600G    34M   600G   1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb          3.1T    34M   3.1T   1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb              41T    34M   41T   1% /var/netwitness/decoder/packetdb
/dev/mapper/decodersmall0-index           30G    33M   30G   1% /var/netwitness/decoder/index0
/dev/mapper/decodersmall0-sessiondb      600G    34M   600G   1% /var/netwitness/decoder/sessiondb0
/dev/mapper/decodersmall0-metadb          3.1T    34M   3.1T   1% /var/netwitness/decoder/metadb0
/dev/mapper/decoder0-packetdb             41T    34M   41T   1% /var/netwitness/decoder/packetdb0

```

```

[root@netwitness Decoder ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda                                  8:0      0  931G  0 disk
├─sda1                              8:1      0    1G  0 part  /boot
├─sda2                              8:2      0  930G  0 part
│   ├─netwitness_vg00-root          253:0      0   30G  0 lvm    /
│   ├─netwitness_vg00-swap          253:1      0    4G  0 lvm    [SWAP]
│   ├─netwitness_vg00-nwhome        253:4      0   2.7T  0 lvm    /var/netwitness
│   ├─netwitness_vg00-varlog        253:5      0   10G  0 lvm    /var/log
│   └─netwitness_vg00-usrhome        253:6      0   10G  0 lvm    /home
└─sdb                                8:16      0   1.8T  0 disk
    ├─sdb1                          8:17      0   1.8T  0 part
    │   └─netwitness_vg00-nwhome      253:4      0   2.7T  0 lvm    /var/netwitness
└─sdc                                8:32      0   3.7T  0 disk
    ├─sdc1                          8:33      0   3.7T  0 part
    │   ├─decodersmall-decoroot      253:7      0   10G  0 lvm    /var/netwitness/decoder
    │   ├─decodersmall-index         253:8      0   30G  0 lvm    /var/netwitness/decoder/index
    │   ├─decodersmall-sessiondb     253:9      0  600G  0 lvm    /var/netwitness/decoder/sessiondb
    │   └─decodersmall-metadb        253:10     0    3T  0 lvm    /var/netwitness/decoder/metadb
    └─sdd                            8:48      0   40T  0 disk
        ├─sdd1                      8:49      0   40T  0 part
        │   └─decoder-packetdb       253:11     0   40T  0 lvm    /var/netwitness/decoder/packetdb
└─sde                                8:64      0   3.7T  0 disk
    ├─sde1                          8:65      0   3.7T  0 part
    │   ├─decodersmall0-index        253:12     0   30G  0 lvm    /var/netwitness/decoder/index0
    │   ├─decodersmall0-sessiondb    253:13     0  600G  0 lvm    /var/netwitness/decoder/sessiondb0
    │   └─decodersmall0-metadb       253:14     0    3T  0 lvm    /var/netwitness/decoder/metadb0
    └─sdf                            8:80      0   40T  0 disk
        ├─sdf1                      8:81      0   40T  0 part
        │   └─decoder0-packetdb      253:15     0   40T  0 lvm    /var/netwitness/decoder/packetdb0

```

10. Execute the `srvAlloc` command with the following parameters to add the storage information into the Service Configuration settings.
 - `service=decoder volume=decodersmall commit=1`
 - `service=decoder volume=decodersmall0 commit=1`
 - `service=decoder volume=decoder commit=1`
 - `service=decoder volume=decoder0 commit=1`
11. Allocate `decodersmall` and `decodersmall0` volumes to `decoder` service as shown below. Similarly, `decoder` and `decoder0` volumes are allocated to `decoder` service.

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

srvAlloc Parameters service=decoder volume=decodersmall commit=1

Message Help

parameters:

service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage

volume - <string, (enum-one:The value must be one of the following: decoder|decoder0|decodersmall|decodersmall0|netwitness_vg00)> volume group name

commit - <bool, optional> commit changes

Response Output

Set /database/config/meta.dir to /var/netwitness/decoder/metadb==2.86 TB

Set /database/config/session.dir to /var/netwitness/decoder/sessiondb==569.72 GB

Set /index/config/index.dir to /var/netwitness/decoder/index==28.49 GB

Properties for Decoder - Decoder (DECODER) /deviceappliance/appliance.

srvAlloc Parameters service=decoder volume=decodersmall0 commit=1

Message Help

parameters:

service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage

volume - <string, (enum-one:The value must be one of the following: decoder|decoder0|decodersmall|decodersmall0|netwitness_vg00)> volume group name

commit - <bool, optional> commit changes

Response Output

Set /database/config/meta.dir to /var/netwitness/decoder/metadb==2.86 TB;/var/netwitness/decoder/metadb0==2.87 TB

Set /database/config/session.dir to /var/netwitness/decoder/sessiondb==569.72 GB;/var/netwitness/decoder/sessiondb0==569.72 GB

Set /index/config/index.dir to /var/netwitness/decoder/index==28.49 GB;/var/netwitness/decoder/index0==28.49 GB

12. Use **Explore->database->config->Meta Database Directory / Session Database Directory / Packet/Log Database Directory** parameter to confirm the values set using srvAlloc.

Change Service | Decoder - Decoder | Explore

Decoder - Decoder

- Decoder - Decoder (DECODER)
 - connections
 - database
 - config
 - stats
 - decoder
 - deviceappliance
 - appliance
 - connections
 - logs
 - rest
 - services
 - storedproc
 - sys
 - users
 - index

Meta Database Directory (meta.dir) /var/netwitness/decoder/metadb==2.86 TB;/var/netwitness/decoder/metadb0==2.87 TB

Packet/Log Database Directory (packet.dir) /var/netwitness/decoder/packetdb==38.02 TB;/var/netwitness/decoder/packetdb0==38.02 TB

Packet Open Files (packet.files)	auto
Packet Minimum Free Space (packet.free.space.min)	23 GB
Packet Index Fidelity (packet.index.fidelity)	1
Packet Integrity Flush (packet.integrity.flush)	sync
Packet Write Block Size (packet.write.block.size)	64 KB
Session Database Directory (session.dir)	/var/netwitness/decoder/sessiondb==569.72 GB;/var/netwitness/decoder/sessiondb0==569.72 GB
Cold Session Database Directory (session.dir.cold)	
Warm Session Database Directory (session.dir.warm)	
Session File Size (session.file.size)	auto
Session Open Files (session.files)	auto
Session Minimum Free Space (session.free.space.min)	23 GB
Session Integrity Flush (session.integrity.flush)	sync
Session Write Block Size (session.write.block.size)	32 KB

13. Reconfigure the following Network Decoder service and its database to detect and take advantage of all of the free space as described in [Task 5 - \(Optional\) Reconfigure Storage Configuration for 10G Capture](#).

Configure Storage for Network Concentrator

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for a Network Concentrator physical host.

1. Execute the `raidList` command.

raidList

▼

Parameters

Message Help

list drive shelves attached to this appliance
security.roles: appliance.manage

Response Output

Controller 0, Enclosure 32

Vendor: DP
Model: BP13G+EXP
In Use: true
Drives: 931.511 GB x 2
1.818 TB x 2
Devices: sda
sdb

Controller 1, Enclosure 6

Vendor: EMC
Model: ESES Enclosure
In Use: false
Drives: 186.309 GB x 6
3.637 TB x 9
Devices:

2. Execute the `raidNew` command with the following parameters.

controller=1 enclosure=6 scheme=concentrator

raidNew	Parameters	controller=1 enclosure=6 scheme=concentrator commit=1
Message Help		
parameters: controller - <int32, {enum-one:0,1}> Controller the shelf is attached to enclosure - <uint32, {enum-one:32,6}> Enclosure number of the shelf to clear scheme - <string, {enum-one:decoder logdecoder concentrator archiver network-hybrid log-hybrid}> Type of RAID volumes to allocate		

Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=6:0,6:1,6:2,6:3,6:4,6:5 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded

/opt/MegaRAID/perccli/perccli64 /c1 add vd r6 drives=6:6,6:7,6:8,6:9,6:10,6:11,6:12,6:13,6:14 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
```

```
[root@NWHOST1500 ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      0  931G  0 disk
├─sda1                              8:1      0    1G  0 part /boot
└─sda2                              8:2      0  930G  0 part
   ├─netwitness_vg00-root            253:0    0   30G  0 lvm  /
   ├─netwitness_vg00-swap            253:1    0    4G  0 lvm  [SWAP]
   ├─netwitness_vg00-nwhome          253:2    0  2.7T  0 lvm  /var/netwitness
   ├─netwitness_vg00-varlog          253:3    0   10G  0 lvm  /var/log
   └─netwitness_vg00-usrhome          253:4    0   10G  0 lvm  /home
sdb                                  8:16     0   1.8T  0 disk
├─sdb1                              8:17     0   1.8T  0 part
└─netwitness_vg00-nwhome            253:2    0  2.7T  0 lvm  /var/netwitness
sdc                                  8:32     0  928.8G  0 disk
sdd                                  8:48     0  25.5T  0 disk
[root@NWHOST1500 ~]#
```

3. Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the **concentrator** volume before **index** volume or it will fail.

```
name=sdd service=concentrator volume=concentrator commit=1
```

partNew Parameters name=sdd service=concentrator volume=concentrator commit=1

Message Help

parameters:

- name - <string, {enum-one:sdc,sdd}> block device name
- service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
- volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create

Response Output

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f concentrator /dev/sdd1
Volume group "concentrator" successfully created
/sbin/lvcreate -y -n root -L 30G concentrator
Logical volume "root" created.
/sbin/mkfs.xfs /dev/concentrator/root
meta-data=/dev/concentrator/root isize=512 agcount=4, agsize=1966080 blks
=          sectsz=512 attr=2, projid32bit=1
=          crc=1 finobt=0, sparse=0
data      =          bsize=4096 blocks=7864320, imaxpct=25
=          sunit=0 swidth=0 blks
naming    =version 2          bsize=4096 ascii-ci=0 ftype=1
log       =internal log      bsize=4096 blocks=3840, version=2
=          sectsz=512 sunit=0 blks, lazy-count=1
realtime  =none             extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator
/bin/mount /var/netwitness/concentrator
```

```
[root@NWHOST1500 ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                 8:0    0  931G  0 disk
├─sda1                             8:1    0    1G  0 part /boot
└─sda2                             8:2    0  930G  0 part
   ├─netwitness_vg00-root           253:0    0   30G  0 lvm /
   ├─netwitness_vg00-swap           253:1    0    4G  0 lvm [SWAP]
   ├─netwitness_vg00-nwhome         253:2    0  2.7T  0 lvm /var/netwitness
   ├─netwitness_vg00-varlog         253:3    0   10G  0 lvm /var/log
   └─netwitness_vg00-usrhome         253:4    0   10G  0 lvm /home
sdb                                 8:16   0   1.8T  0 disk
├─sdb1                             8:17   0   1.8T  0 part
└─netwitness_vg00-nwhome           253:2    0  2.7T  0 lvm /var/netwitness
sdc                                 8:32   0  928.8G  0 disk
sdd                                 8:48   0  25.5T  0 disk
├─sdd1                             8:49   0  25.5T  0 part
│   ├─concentrator-root            253:5    0   30G  0 lvm /var/netwitness/concentrator
│   ├─concentrator-sessiondb       253:6    0   600G  0 lvm /var/netwitness/concentrator/sessiondb
│   └─concentrator-metadb          253:7    0  24.9T  0 lvm /var/netwitness/concentrator/metadb
```

4. Execute the partNew command with the following parameters with the following parameters to create an index on SSDs.

```
name=sdc service=concentrator volume=index commit=1
```

partNew
Parameters
name=sdc service=concentrator volume=index commit=1

Message Help

parameters:
name - <string, {enum-one:sdc,sdd}> block device name
service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create

Response Output

```

/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f index /dev/sdc1
Volume group "index" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index
Wiping xfs signature on /dev/index/index.
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index  isize=512  agcount=4, agsize=60866304 blks
=          sectsz=4096  attr=2, projid32bit=1
=          crc=1      finobt=0, sparse=0
data      =          bsize=4096  blocks=243465216, imaxpct=25
=          sunit=0    swidth=0 blks
naming    =version 2      bsize=4096  ascii-ci=0 ftype=1
log        =internal log  bsize=4096  blocks=118879, version=2
=          sectsz=4096  sunit=1 blks, lazy-count=1
realtime   =none         extsz=4096  blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index

[root@NWHOST1500 ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      0  931G  0 disk
├─sda1                              8:1      0    1G  0 part /boot
└─sda2                              8:2      0   930G  0 part
   ├─netwitness_vg00-root            253:0    0    30G  0 lvm  /
   ├─netwitness_vg00-swap            253:1    0     4G  0 lvm  [SWAP]
   ├─netwitness_vg00-nwhome          253:2    0   2.7T  0 lvm  /var/netwitness
   ├─netwitness_vg00-varlog          253:3    0    10G  0 lvm  /var/log
   └─netwitness_vg00-usrhome          253:4    0    10G  0 lvm  /home
bdb                                  8:16     0   1.8T  0 disk
├─sdb1                              8:17     0   1.8T  0 part
└─netwitness_vg00-nwhome            253:2    0   2.7T  0 lvm  /var/netwitness
sdc                                  8:32     0  928.8G  0 disk
├─sdc1                              8:33     0  928.8G  0 part
└─index-index                       253:8    0  928.8G  0 lvm  /var/netwitness/concentrator/index
sdd                                  8:48     0   25.5T  0 disk
├─sdd1                              8:49     0   25.5T  0 part
├─concentrator-root                 253:5    0    30G  0 lvm  /var/netwitness/concentrator
├─concentrator-sessiondb            253:6    0   600G  0 lvm  /var/netwitness/concentrator/sessiondb
└─concentrator-metadb               253:7    0   24.9T  0 lvm  /var/netwitness/concentrator/metadb

```

```
[root@NWHOST1500 ~]# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/netwitness_vg00-root	30G	2.1G	28G	7%	/
devtmpfs	63G	0	63G	0%	/dev
tmpfs	63G	12K	63G	1%	/dev/shm
tmpfs	63G	10M	63G	1%	/run
tmpfs	63G	0	63G	0%	/sys/fs/cgroup
/dev/sda1	1014M	91M	924M	9%	/boot
/dev/mapper/netwitness_vg00-varlog	10G	52M	10G	1%	/var/log
/dev/mapper/netwitness_vg00-usrhome	10G	33M	10G	1%	/home
/dev/mapper/netwitness_vg00-nwhome	2.7T	98M	2.7T	1%	/var/netwitness
tmpfs	13G	0	13G	0%	/run/user/0
/dev/mapper/concentrator-root	30G	33M	30G	1%	/var/netwitness/concentrator
/dev/mapper/concentrator-sessiondb	600G	33M	600G	1%	/var/netwitness/concentrator/sessiondb
/dev/mapper/concentrator-metadb	25T	33M	25T	1%	/var/netwitness/concentrator/metadb
/dev/mapper/index-index	929G	33M	929G	1%	/var/netwitness/concentrator/index

5. Execute the `srvAlloc` command with the following parameters.

```
service=concentrator volume=index commit=1
```

srvAlloc Parameters service=concentrator volume=index commit=1

Message Help

parameters:

service - <string, {enum-one:archiver | concentrator | decoder | logdecoder}> service that will use storage

volume - <string, {enum-one:concentrator,index,netwitness_vg00}> volume group name

commit - <bool, optional> commit changes

Response Output

```
Set /index/config/index.dir to /var/netwitness/concentrator/index==881.87 GB
```

Change Service	NWHOST1500 - Concentrator	Explore
NWHOST1500 - Concentrator	/index/config	NWHOST1500 - Concentrator (CONC)
NWHOST1500 - Concentrator (CONC)	index.dir	/var/netwitness/concentrator/index==881.87 GB
concentrator	index.dir.cold	
connections	index.dir.warm	
database	index.slices.open	42
deviceappliance	page.compression	huffhybrid
index	reindex.enable	true
config	save.session.count	auto

6. Execute the `srvAlloc` command with the following parameters.
- ```
service=concentrator volume=concentrator commit=1
```

srvAlloc

Parameters

service=concentrator volume=concentrator commit=1

Message Help

parameters:

service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage

volume - <string, {enum-one:concentrator,index,netwitness\_vg00}> volume group name

commit - <bool, optional> commit changes

## Response Output

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==23.6 TB

Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==569.72 GB

| Change Service   NWHOST1500 - Concentrator (CONCENTRAT   Explore |                        |                                                   |
|------------------------------------------------------------------|------------------------|---------------------------------------------------|
| NWHOST1500 - Concentrator                                        | /database/config       | NWHOST1500 - Concentrator (CONC                   |
| NWHOST1500 - Concentrator (CONC                                  | hash.algorithm         | none                                              |
| concentrator                                                     | hash.databases         | session,meta                                      |
| connections                                                      | hash.dir               |                                                   |
| database                                                         | manifest.dir           |                                                   |
| config                                                           | meta.compression       | none                                              |
| stats                                                            | meta.compression.level | 0                                                 |
| deviceappliance                                                  | meta.dir               | /var/netwitness/concentrator/metadb==23.6 TB      |
| index                                                            | meta.dir.cold          |                                                   |
| logs                                                             | meta.dir.warm          |                                                   |
| rest                                                             | meta.file.size         | auto                                              |
| sdk                                                              | meta.files             | auto                                              |
| services                                                         | meta.free.space.min    | 23 GB                                             |
| storedproc                                                       | meta.index.fidelity    | 4                                                 |
| sys                                                              | meta.integrity.flush   | sync                                              |
| users                                                            | meta.write.block.size  | 64 KB                                             |
|                                                                  | session.dir            | /var/netwitness/concentrator/sessiondb==569.72 GB |

## Configure Storage for Log Decoder Hybrid

The following scenario configures storage on one, non-encrypted, 15-Drive DAC for a Log Decoder Hybrid physical host.

1. Execute the `raidList` command.

raidList

▼

Parameters

Message Help

list drive shelves attached to this appliance  
security.roles: appliance.manage

### Response Output

#### Controller 0, Enclosure 32

Vendor: DP  
Model: BP13G+EXP  
In Use: true  
Drives: 745.21 GB x 2  
931.511 GB x 4  
5.457 TB x 8

#### Devices: sda

sdb  
sdc  
sdd  
sde

#### Controller 1, Enclosure 31

Vendor: EMC  
Model: ESES Enclosure  
In Use: false  
Drives: 3.637 TB x 15  
Devices:

2. Execute the `raidNew` command with the following parameters.  
`controller=1 enclosure=31 scheme=log-hybrid commit=1`

raidNew  Parameters controller=1 enclosure=31 scheme=log-hybrid commit=1

Message Help

controller - <int32, {enum-one:0,1}> Controller the shelf is attached to  
 enclosure - <uint32, {enum-one:32,31}> Enclosure number of the shelf to clear  
 scheme - <string, {enum-one:decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid}> Type of RAID volumes to allocate  
 preferSecure - <bool, optional, {bool:0,1,yes,no,true,false,on,off}> Prefer creation of a secure array given compatible physical drives and a controller with a security key set

## Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=31:0,31:1,31:2,31:3,31:4,31:5,31:6 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded

/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=31:7,31:8,31:9,31:10,31:11,31:12,31:13,31:14 wb ra cached Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
```

```
[root@NWHOST1700 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:11 0 876G 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:12 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:13 0 10G 0 lvm /home
sdb 8:16 0 931G 0 disk
├─sdb1 8:17 0 931G 0 part
├─decodermeta-vlnwdm 253:9 0 931G 0 lvm /var/netwitness/decoder/metadb
sdc 8:32 0 16.4T 0 disk
├─sdc1 8:33 0 16.4T 0 part
│ ├─decoderpacket-vlnwdp 253:2 0 16.2T 0 lvm /var/netwitness/decoder/packetdb
│ ├─decoderpacket-vlnwds 253:3 0 100G 0 lvm /var/netwitness/decoder/sessiondb
│ ├─decoderpacket-vlnwdi 253:4 0 50G 0 lvm /var/netwitness/decoder/index
│ └─decoderpacket-vlnwd 253:5 0 30G 0 lvm /var/netwitness/decoder
sdd 8:48 0 16.4T 0 disk
├─sdd1 8:49 0 16.4T 0 part
│ ├─concentrator-vlnwcm 253:6 0 14.9T 0 lvm /var/netwitness/concentrator/metadb
│ ├─concentrator-vlnwcs 253:7 0 1.5T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-vlnwc 253:8 0 30G 0 lvm /var/netwitness/concentrator
sde 8:64 0 744.6G 0 disk
├─sde1 8:65 0 744.6G 0 part
├─index-vlnwci 253:10 0 744.6G 0 lvm /var/netwitness/concentrator/index
sdf 8:80 0 21.8T 0 disk
sdg 8:96 0 25.5T 0 disk
```

3. Execute the `partNew` command with the following parameters with the following parameters.

- `name=sdf service=concentrator volume=concentrator commit=1`

partNew  `name=sdf service=concentrator volume=concentrator commit=1`

Message Help

name - <string, {enum-one:sdf,sdg}> block device name  
 service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage  
 volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create  
 commit - <bool, optional> commit changes

## Response Output

```
/sbin/parted -s /dev/sdf mklabel gpt
/sbin/parted -s -a optimal /dev/sdf mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdf1
Physical volume "/dev/sdf1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdf1
Volume group "concentrator0" successfully created
```

```
[root@NWHOST1700 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:11 0 876G 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:12 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:13 0 10G 0 lvm /home
sdb 8:16 0 931G 0 disk
├─sdb1 8:17 0 931G 0 part
│ └─decodermeta-vlnwdm 253:9 0 931G 0 lvm /var/netwitness/decoder/metadb
sdc 8:32 0 16.4T 0 disk
├─sdc1 8:33 0 16.4T 0 part
│ ├─decoderpacket-vlnwdp 253:2 0 16.2T 0 lvm /var/netwitness/decoder/packetdb
│ ├─decoderpacket-vlnwds 253:3 0 100G 0 lvm /var/netwitness/decoder/sessiondb
│ ├─decoderpacket-vlnwdi 253:4 0 50G 0 lvm /var/netwitness/decoder/index
│ └─decoderpacket-vlnwd 253:5 0 30G 0 lvm /var/netwitness/decoder
sdd 8:48 0 16.4T 0 disk
├─sdd1 8:49 0 16.4T 0 part
│ ├─concentrator-vlnwcm 253:6 0 14.9T 0 lvm /var/netwitness/concentrator/metadb
│ ├─concentrator-vlnwcs 253:7 0 1.5T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-vlnwc 253:8 0 30G 0 lvm /var/netwitness/concentrator
sde 8:64 0 744.6G 0 disk
├─sde1 8:65 0 744.6G 0 part
│ └─index-vlnwci 253:10 0 744.6G 0 lvm /var/netwitness/concentrator/index
sdf 8:80 0 21.8T 0 disk
├─sdf1 8:81 0 21.8T 0 part
│ ├─concentrator0-sessiondb 253:14 0 600G 0 lvm /var/netwitness/concentrator/sessiondb0
│ └─concentrator0-metadb 253:15 0 21.2T 0 lvm /var/netwitness/concentrator/metadb0
sdg 8:96 0 25.5T 0 disk
```

- name=sdg service=logdecoder volume=logdecoder commit=1

partNew  name=sdg service=logdecoder volume=logdecoder commit=1

Message Help

name - <string, {enum-one:sdf,sdg}> block device name  
 service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage  
 volume - <string, optional, {enum-one:index|concentrator|decodersmall|decoder|logdecodersmall|logdecoder|archiver}> volume to create  
 commit - <bool, optional> commit changes

Response Output

```
/sbin/parted -s /dev/sdg mklabel gpt
/sbin/parted -s -a optimal /dev/sdg mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdg1
Physical volume "/dev/sdg1" successfully created.
/sbin/vgcreate -f logdecoder0 /dev/sdg1
Volume group "logdecoder0" successfully created
```

```
[root@NWHOST1700 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
└─sda2 8:2 0 930G 0 part
 ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
 ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
 ├─netwitness_vg00-nwhome 253:11 0 876G 0 lvm /var/netwitness
 ├─netwitness_vg00-varlog 253:12 0 10G 0 lvm /var/log
 └─netwitness_vg00-usrhome 253:13 0 10G 0 lvm /home
sdb 8:16 0 931G 0 disk
├─sdb1 8:17 0 931G 0 part
└─decodermeta-vlnwdm 253:9 0 931G 0 lvm /var/netwitness/decoder/metadb
sdc 8:32 0 16.4T 0 disk
├─sdc1 8:33 0 16.4T 0 part
│ ├─decoderpacket-vlnwdp 253:2 0 16.2T 0 lvm /var/netwitness/decoder/packetdb
│ ├─decoderpacket-vlnwds 253:3 0 100G 0 lvm /var/netwitness/decoder/sessiondb
│ ├─decoderpacket-vlnwdi 253:4 0 50G 0 lvm /var/netwitness/decoder/index
│ └─decoderpacket-vlnwd 253:5 0 30G 0 lvm /var/netwitness/decoder
sdd 8:48 0 16.4T 0 disk
├─sdd1 8:49 0 16.4T 0 part
│ ├─concentrator-vlnwcm 253:6 0 14.9T 0 lvm /var/netwitness/concentrator/metadb
│ ├─concentrator-vlnwcs 253:7 0 1.5T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-vlnwc 253:8 0 30G 0 lvm /var/netwitness/concentrator
sde 8:64 0 744.6G 0 disk
├─sde1 8:65 0 744.6G 0 part
└─index-vlnwc1 253:10 0 744.6G 0 lvm /var/netwitness/concentrator/index
sdf 8:80 0 21.8T 0 disk
├─sdf1 8:81 0 21.8T 0 part
│ ├─concentrator0-sessiondb 253:14 0 600G 0 lvm /var/netwitness/concentrator/sessiondb0
│ └─concentrator0-metadb 253:15 0 21.2T 0 lvm /var/netwitness/concentrator/metadb0
sdg 8:96 0 25.5T 0 disk
├─sdg1 8:97 0 25.5T 0 part
└─logdecoder0-packetdb 253:16 0 25.5T 0 lvm /var/netwitness/decoder/packetdb0
```

4. Execute the `srvAlloc` command with the following parameters.

- `service=concentrator volume=concentrator0 commit=1`

Message Help

service - <string, {enum-one:archiver|concentrator|decoder|logdecoder}> service that will use storage  
 volume - <string, {enum-one:concentrator,concentrator0,decodermeta,decoderpacket,index,logdecoder0,netwitness\_vg00}> volume group name  
 commit - <bool, optional> commit changes

## Response Output

```
Set /database/config/meta.dir to /var/netwitness/concentrator/metadb=14.08 TB;/var/netwitness/concentrator/metadb0==20.17 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb=1.41 TB;/var/netwitness/concentrator/sessiondb0==569.72 GB
```

|                               |                               |                                                                                                   |
|-------------------------------|-------------------------------|---------------------------------------------------------------------------------------------------|
| Change Service                | NWHOST1700 - Concentrator (C) | Explore                                                                                           |
| NWHOST1700 - Concentrator     | /database/config              | NWHOST1700 - Concentrator                                                                         |
| NWHOST1700 - Concentrator (C) | hash.algorithm                | none                                                                                              |
| concentrator                  | hash.databases                | session,meta                                                                                      |
| connections                   | hash.dir                      |                                                                                                   |
| database                      | manifest.dir                  |                                                                                                   |
| config                        | meta.compression              | none                                                                                              |
| stats                         | meta.compression.level        | 0                                                                                                 |
| deviceappliance               | meta.dir                      | /var/netwitness/concentrator/metadb=14.08 TB;/var/netwitness/concentrator/metadb0==20.17 TB       |
| index                         | meta.dir.cold                 |                                                                                                   |
| logs                          | meta.dir.warm                 |                                                                                                   |
| rest                          | meta.file.size                | auto                                                                                              |
| sdks                          | meta.files                    | auto                                                                                              |
| services                      | meta.free.space.min           | 132 GB                                                                                            |
| storedproc                    | meta.index.fidelity           | 4                                                                                                 |
| sys                           | meta.integrity.flush          | sync                                                                                              |
| users                         | meta.write.block.size         | 64 KB                                                                                             |
|                               | session.dir                   | /var/netwitness/concentrator/sessiondb=1.41 TB;/var/netwitness/concentrator/sessiondb0==569.72 GB |

- `service=logdecoder volume=logdecoder0 commit=1`

## Appendix E. Sample Storage Configuration Scenarios for 8 or 12-Drive PowerVault

---

This appendix illustrates the following example of how to configure storage on one non-encrypted 8 or 12-drive PowerVault external storage devices.

- [Configure Storage for Archiver using NW-PV-A/NW-PV-A-N](#)
- [Configure Storage for Decoder using NW-PV-B/NW-PV-B-N](#)
- [Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N](#)
- [Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N](#)
- [Configure Storage for Log Hybrid using NW-PV-A/NW-PV-A-N](#)
- [Configure Storage for Network Hybrid using NW-PV-A/NW-PV-A-N](#)
- [Configure Storage for Endpoint Log Hybrid using NW-PV-A/NW-PV-A-N](#)

### Configure Storage for Archiver using NW-PV-A/NW-PV-A-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for an Archiver physical host.

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>

## b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList ▼ Parameters:  Send

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 18:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 931.512 GB HDD x 2
 1.819 TB HDD x 2
 2.182 TB HDD x 3
Devices: sda /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:0:0
 sdb /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 249, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 10.692 TB HDD x 12
Devices:
```

- Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

```
controller=1 enclosure=251 scheme=archiver commit=1
```

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance  
 raidNew ▼ Parameters:  Send

Message Help

```
raidNew: allocate RAID devices in a drive shelf
security.roles: appliance.manage
parameters:
 controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
 enclosure - <uint32, optional, {enum-one:The value must be one of the following: 64|251}> Enclosure number of the shelf to clear. Required if the contr
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-decod
 allocate
```

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=251&scheme=archiver&commit=1

Output (or command manual help)

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r6 drives=251:0,251:1,251:2,251:3,251:4,251:5,251:6,251:7,251:8,251:9,251:10,251:11 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.66.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.
```

- Execute the `raidList` command to verify the new RAID array.

You should now see the following information.

In Use: TRUE

Devices: <device> (for example, sdc)

Properties for /appliance  
 raidList ▼ Parameters:

Message Help

raidList: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 18:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 1.09 TB HDD x 2
 2.182 TB HDD x 2
Devices: sda /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:0:0
 sdb /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 251, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: true
Drives: 7.277 TB HDD x 12
Devices: sdc /dev/disk/by-path/pci-0000:3b:00.0-scsi-0:2:0:0
```

- Execute the `partNew` command with the following parameters to create partitions and mount points in the `etc/fstab` file.

name=sdc service=archiver volume=archiver commit=1

Properties for /appliance  
 partNew ▼ Parameters: name=sdc service=archiver volume=archiver commit=1

Message Help

partNew: make partitions on a block device  
 security.roles: appliance.manage  
 parameters:  
 name - <string, (enum-one:The value must be one of the following: sdc)> block device name  
 service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage  
 volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|log-indexed-decoder|logindex)> volume to create

/appliance?msg=partNew&force-content-type=text/plain&name=sdc&service=archiver&volume=archiver&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVH 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f archiver0 /dev/sdc1
Volume group "archiver0" successfully created
/sbin/lvcreate -y -n database -l 100%FREE archiver0
Wiping xfs signature on /dev/archiver0/database.
Logical volume "database" created.
/sbin/mkfs.xfs /dev/archiver0/database
meta-data=/dev/archiver0/database isize=512 agcount=73, agsize=268435424 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=19533659136, imaxpct=1
 = sunit=32 swidth=256 blks
naming = bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/archiver/database0
/bin/mount /var/netwitness/archiver/database0
```

- Execute the `srvAlloc` command with the following parameters to allocate the space to the archiver service. This adds storage to the archiver service configuration and restarts the service every time it is executed.

service=archiver volume=archiver0 commit=1

## Configure Storage for Decoder using NW-PV-B/NW-PV-B-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Network Decoder physical host.

**Note:** The block device size depends on the PV type (drive count) and the drive size (8 TB or 12 TB or 16 TB).

1. Execute the `raidList` command.

- a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

```
In Use: FALSE
Devices: <empty>
```

- b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList Parameters:  Send

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 18:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 931.512 GB HDD x 2
 1.819 TB HDD x 2
 2.182 TB HDD x 3
Devices: sda /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:0:0
 sdb /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 249, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 10.692 TB HDD x 12
Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded. Parameters for the first enclosure:

`controller=1 enclosure=0 scheme=decoder-hotspare commit=1`

| Scheme                | Enclosure Type | Drives Required                | Allocation                                                                                                                                  |
|-----------------------|----------------|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| decoder               | External       | 8 or 12 or 15 HDDs             | 3x drives in RAID 5 for decodersmall, all remaining drives in RAID 5 for decoder                                                            |
| log decoder           | External       | 8 or 12 or 15 HDDs             | Same as decoder configuration                                                                                                               |
| decoder hot spare     | External       | 8 or 12 or 15 HDDs             | 2x drives in RAID 1 for decodersmall, 1 drive as hot spare, all remaining drives in RAID 5 for decoder                                      |
| log decoder hot spare | External       | 8 or 12 or 15 HDDs             | Same as decoder hot spare configuration                                                                                                     |
| archiver              | External       | 8 or 12 or 15 HDDs             | All drives in RAID 6 for archiver database volume                                                                                           |
| packet expansion      | External       | 8 or 12 or 15 HDDs             | All drives in RAID 6 for decoder volume. No drives allocated for decodersmall                                                               |
| network hybrid        | External       | 8 or 12 or 15 HDDs             | 3x drives in RAID 5 for meta expansion, all remaining drives in RAID 5 for packet expansion                                                 |
| network hybrid        | Internal       | S5 / S6 hybrid drive set       | 2x small HDD RAID 1 for decoder meta, 5x large HDD decoder, 3x large HDD concentrator, 2x SSD index                                         |
| log hybrid            | External       | 8 or 12 or 15 HDDs             | Half of the drives in RAID 5 for meta expansion, half the drives in RAID 5 for packet expansion                                             |
| log hybrid            | Internal       | S5 / S6 hybrid drive set       | 2x small HDD RAID 1 for decoder meta, 4x large HDD decoder, 4x large HDD concentrator, 2x SSD index                                         |
| endpoint hybrid       | Internal       | S5 / S6 hybrid drive set       | 2x small HDD RAID 1 for decoder meta, 4x large HDD RAID 10 for log decoder and endpoint, 4x large HDD RAID 5 for concentrator, 2x SSD index |
| log indexed decoder   | Internal       | S6E hybrid drive set           | 10x HDD RAID 6 for log decoder meta and packet, 2x SSD index                                                                                |
| concentrator          | External       | 2 or more SSDs, 4 or more HDDs | All SSDs in RAID 1 or RAID 5 for index, all HDDs in RAID 6 for meta                                                                         |

raidNew  Parameters controller=1 enclosure=249 scheme=decoder-hotspare commit=1

Message Help

parameters:  
 controller - <uint32, (enum-one:The value must be one of the following: 0,1)> Controller the shelf is attached to  
 enclosure - <uint32, optional, (enum-one:The value must be one of the following: 64|249)> Enclosure number of the shelf to clear. Required if the controller has more than one enclosure attached.  
 scheme - <string, (enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-decoder|endpoint-log-hybrid|packet-

## Response Output

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r1 drives=249:0,249:1 ra Strip=128
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86\_64

Controller = 1

Status = Success

Description = Add VD Succeeded.

```
/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=249:2,249:3,249:4,249:5,249:6,249:7,249:8,249:9,249:10 ra Strip=128
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86\_64

Controller = 1

Status = Success

Description = Add VD Succeeded.

Description = Add VD Succeeded.

```
/opt/MegaRAID/perccli/perccli64 /c1 /e249 /s11 add hotsparedrive
```

CLI Version = 007.1623.0000.0000 May 17, 2021

Operating system = Linux 3.10.0-1160.83.1.el7.x86\_64

Controller = 1

Status = Success

Description = Add Hot Spare Succeeded.

```
[root@netwitness ~]# lsblk
```

| NAME                      | MAJ:MIN | RM | SIZE  | RO | TYPE | MOUNTPOINT      |
|---------------------------|---------|----|-------|----|------|-----------------|
| sda                       | 8:0     | 0  | 931G  | 0  | disk |                 |
| └─sda1                    | 8:1     | 0  | 1G    | 0  | part | /boot           |
| └─sda2                    | 8:2     | 0  | 930G  | 0  | part |                 |
| └─netwitness_vg00-root    | 253:0   | 0  | 30G   | 0  | lvm  | /               |
| └─netwitness_vg00-swap    | 253:1   | 0  | 4G    | 0  | lvm  | [SWAP]          |
| └─netwitness_vg00-nwhome  | 253:2   | 0  | 2.7T  | 0  | lvm  | /var/netwitness |
| └─netwitness_vg00-varlog  | 253:3   | 0  | 10G   | 0  | lvm  | /var/log        |
| └─netwitness_vg00-usrhome | 253:4   | 0  | 10G   | 0  | lvm  | /home           |
| sdb                       | 8:16    | 0  | 1.8T  | 0  | disk |                 |
| └─sdb1                    | 8:17    | 0  | 1.8T  | 0  | part |                 |
| └─netwitness_vg00-nwhome  | 253:2   | 0  | 2.7T  | 0  | lvm  | /var/netwitness |
| sdc                       | 8:32    | 0  | 10.7T | 0  | disk |                 |
| sdd                       | 8:48    | 0  | 85.5T | 0  | disk |                 |

3. Execute the `partNew` command to create the **decodersmall** partition first (decoder dir, index, metadb, sessiondb) (First Enclosure, SDC, SDD) with the following parameters.

`name=sdc service=decoder volume=decodersmall commit=1`

partNew Parameters: `name=sdc service=decoder volume=decodersmall commit=1`

Send

Message Help

parameters:

name - <string, (enum-one:The value must be one of the following: sdc|sdd)> block device name

service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage

volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log|hybrid-log|invariant-decoder|logindex)> volume to create

Response Output

```
Volume group "decodersmall" successfully created
/sbin/lvcreate -y -n decoroot -L 10G decodersmall
Wiping xfs signature on /dev/decodersmall/decroot.
Logical volume "decoroot" created.
/sbin/mkfs.xfs /dev/decodersmall/decroot
meta-data=/dev/decodersmall/decroot isize=512 agcount=16, agsize=163840 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=2621440, imaxpct=25
= sunit=32 swidth=32 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=2560, version=2
= sectsz=512 sunit=32 blks, lazy-count=1
```

```
[root@s6coreappliance ~]# df -kh
Filesystem Size Used Avail Use% Mounted on
devtmpfs 63G 0 63G 0% /dev
tmpfs 63G 40K 63G 1% /dev/shm
tmpfs 63G 19M 63G 1% /run
tmpfs 63G 0 63G 0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G 6.3G 24G 21% /
/dev/sda1 1014M 91M 924M 9% /boot
/dev/mapper/netwitness_vg00-nwhome 2.7T 1.2G 2.7T 1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome 10G 33M 10G 1% /home
/dev/mapper/netwitness_vg00-varlog 10G 915M 9.1G 9% /var/log
tmpfs 13G 0 13G 0% /run/user/0
/dev/mapper/decodersmall-decoroot 10G 33M 10G 1% /var/netwitness/decoder
/dev/mapper/decodersmall-index 30G 33M 30G 1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb 600G 34M 600G 1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb 11T 34M 11T 1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb 86T 37M 86T 1% /var/netwitness/decoder/packetdb
```

4. Execute the `partNew` command to create the decoder volume (packetdb) (First Enclosure, SDC, SDD) with the following parameters.

`name=sdd service=decoder volume=decoder commit=1`

partNew Parameters: `name=sdd service=decoder volume=decoder commit=1`

Send

Message Help

parameters:

name - <string, (enum-one:The value must be one of the following: sdc|sdd)> block device name

service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage

volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log|hybrid-log|invariant-decoder|logindex)> volume to create

Response Output

```
Volume group "decoder" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder/packetdb
meta-data=/dev/decoder/packetdb isize=512 agcount=86, agsize=268435424 blks
= sectsz=512 attr=2, projid32bit=1
= crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=22960667648, imaxpct=1
= sunit=32 swidth=256 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
= sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
```

```
[root@s6coreappliance ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 1.8T 0 disk
├─sdb1 8:17 0 1.8T 0 part
└─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
sdc 8:32 0 10.7T 0 disk
├─sdc1 8:33 0 10.7T 0 part
│ ├─decodersmall-decoroot 253:5 0 10G 0 lvm /var/netwitness/decoder
│ ├─decodersmall-index 253:6 0 30G 0 lvm /var/netwitness/decoder/index
│ ├─decodersmall-sessiondb 253:7 0 600G 0 lvm /var/netwitness/decoder/sessiondb
│ └─decodersmall-metadb 253:8 0 10.1T 0 lvm /var/netwitness/decoder/metadb
sdd 8:48 0 85.5T 0 disk
├─sdd1 8:49 0 85.5T 0 part
└─decoder-packetdb 253:9 0 85.5T 0 lvm /var/netwitness/decoder/packetdb
```

Use `lsblk` and `df -kh` to confirm the block device sizes and disk allocation.

```
[root@s6coreappliance ~]# df -kh
Filesystem Size Used Avail Use% Mounted on
devtmpfs 63G 0 63G 0% /dev
tmpfs 63G 40K 63G 1% /dev/shm
tmpfs 63G 19M 63G 1% /run
tmpfs 63G 0 63G 0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G 6.3G 24G 21% /
/dev/sda1 1014M 91M 924M 9% /boot
/dev/mapper/netwitness_vg00-nwhome 2.7T 1.2G 2.7T 1% /var/netwitness
/dev/mapper/netwitness_vg00-usrhome 10G 33M 10G 1% /home
/dev/mapper/netwitness_vg00-varlog 10G 915M 9.1G 9% /var/log
tmpfs 13G 0 13G 0% /run/user/0
/dev/mapper/decodersmall-decoroot 10G 33M 10G 1% /var/netwitness/decoder
/dev/mapper/decodersmall-index 30G 33M 30G 1% /var/netwitness/decoder/index
/dev/mapper/decodersmall-sessiondb 600G 34M 600G 1% /var/netwitness/decoder/sessiondb
/dev/mapper/decodersmall-metadb 11T 34M 11T 1% /var/netwitness/decoder/metadb
/dev/mapper/decoder-packetdb 86T 37M 86T 1% /var/netwitness/decoder/packetdb
```

- Execute the `srvAlloc` command with the following parameters to add the storage information into the Service Configuration settings.

- `service=decoder volume=decodersmall commit=1`

srvAlloc

Parameters service=decoder volume=decodersmall commit=1

Send

Message Help

parameters:

service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, {enum-one:The value must be one of the following: decoder|decodersmall|netwitness\_vg00}> volume group name  
commit - <bool, optional> commit changes

Response Output

Set /database/config/meta.dir to /var/netwitness/decoder/metadb==9.56 TB  
Set /database/config/session.dir to /var/netwitness/decoder/sessiondb==569.72 GB  
Set /index/config/index.dir to /var/netwitness/decoder/index==28.49 GB

- `service=decoder volume=decoder commit=1`

|                                                                                                                                                                                                                                                                                                                                          |            |                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|-----------------------------------------|
| srvAlloc                                                                                                                                                                                                                                                                                                                                 | Parameters | service=decoder volume=decoder commit=1 |
| <div>Send</div>                                                                                                                                                                                                                                                                                                                          |            |                                         |
| Message Help                                                                                                                                                                                                                                                                                                                             |            |                                         |
| parameters:<br>service - <string, {enum-one:The value must be one of the following: archiver concentrator decoder logdecoder}> service that will use storage<br>volume - <string, {enum-one:The value must be one of the following: decoder decodersmall netwitness_vg00}> volume group name<br>commit - <bool, optional> commit changes |            |                                         |
| Response Output                                                                                                                                                                                                                                                                                                                          |            |                                         |
| Set /database/config/packet.dir to /var/netwitness/decoder/packetdb==81.26 TB                                                                                                                                                                                                                                                            |            |                                         |

## Configure Storage for Concentrator using NW-PV-C/NW-PV-C-N

The following scenario configures storage on one, non-encrypted, 8-Drive PowerVault for a Network Concentrator physical host.

1. Execute the `raidList` command.

- a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

In Use: FALSE

Devices: <empty>

- b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Output (or command manual help)

```
Controller 0 at PCI Address 18:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 1.09 TB HDD x 2
 2.182 TB HDD x 2
Devices: sda /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:0:0
 sdb /dev/disk/by-path/pci-0000:18:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 3b:00.0, Enclosure 251, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 1.746 TB SSD x 2
 7.277 TB HDD x 6
Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

```
controller=1 enclosure=251 scheme=concentrator commit=1
```

The following example illustrates what you should see after you create a RAID array.

```

Properties for /appliance
raidNew Parameters: controller=1 enclosure=251 scheme=concentrator commit=1 Send

Message Help
scheme - <string, (enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hy
expansion|decoder-hotspare|logdecoder-hotspare)> Type of RAID volumes to allocate
preferSecure - <bool, optional, (bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,o
drives and a controller with a security key set
commit - <bool, optional> commit changes

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=251&scheme=concentrator&commit=1

Output (or command manual help)

/opt/MegaRAID/percc11/percc1164 /c1 add vd r1 drives=251:0,251:1 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.66.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

/opt/MegaRAID/percc11/percc1164 /c1 add vd r6 drives=251:2,251:3,251:4,251:5,251:6,251:7 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.66.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

[root@conc95 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 1.1T 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 1.1T 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 3.2T 0 lvm /var/netwitness
│ └─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
└─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
│ └─netwitness_vg00-nwhome 253:2 0 3.2T 0 lvm /var/netwitness
sdc 8:32 0 1.8T 0 disk
sdd 8:48 0 29.1T 0 disk
[root@conc95 ~]#

```

3. Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the **concentrator** volume before **index** volume or it will fail.

```
name=sdd service=concentrator volume=concentrator commit=1
```

Properties for /appliance

partNew Parameters: name=sdd service=concentrator volume=concentrator commit=1

Message Help

partNew: make partitions on a block device  
security.roles: appliance.manage  
parameters:  
name - <string, (enum-one:The value must be one of the following: sdc|sdd)> block device name  
service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage  
volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersma  
logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex)> volume to create

/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f concentrator /dev/sdd1
Volume group "concentrator" successfully created
/sbin/lvcreate -y -n root -L 30G concentrator
Wiping xfs signature on /dev/concentrator/root.
Logical volume "root" created.
/sbin/mkfs.xfs /dev/concentrator/root
meta-data=/dev/concentrator/root isize=512 agcount=16, agsize=491488 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=7863808, imaxpct=25
 = sunit=32 swidth=128 blks
naming -version 2 bsize=4096 ascii-ci=0 ftype=1
```

4. Execute the partNew command with the following parameters to create an **index** on SSDs.

```
name=sdcc service=concentrator volume=index commit=1
```

Properties for /appliance

partNew Parameters: name=sdcc service=concentrator volume=index commit=1

Message Help

partNew: make partitions on a block device  
security.roles: appliance.manage  
parameters:  
name - <string, (enum-one:The value must be one of the following: sdc|sdd)> block device name  
service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage  
volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersma  
logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex)> volume to create

/appliance?msg=partNew&force-content-type=text/plain&name=sdcc&service=concentrator&volume=index&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f index /dev/sdc1
Volume group "index" successfully created
/sbin/lvcreate -y -n index -L 1000784 index
Wiping xfs signature on /dev/index/index.
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index isize=512 agcount=32, agsize=14646240 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=468679680, imaxpct=5
 = sunit=32 swidth=128 blks
naming -version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=228864, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/sbin/mkdir -p /var/netwitness/concentrator/index
/sbin/mount /var/netwitness/concentrator/index
```

```
[root@con95 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 1.1T 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 1.1T 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 3.2T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
│ └─netwitness_vg00-nwhome 253:2 0 3.2T 0 lvm /var/netwitness
sdc 8:32 0 1.8T 0 disk
├─sdcl 8:33 0 1.8T 0 part
│ └─index-index 253:8 0 1.8T 0 lvm /var/netwitness/concentrator/index
sdd 8:48 0 29.1T 0 disk
├─sdd1 8:49 0 29.1T 0 part
│ ├─concentrator-root 253:5 0 30G 0 lvm /var/netwitness/concentrator
│ ├─concentrator-sessiondb 253:6 0 2.9T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-metadb 253:7 0 26.2T 0 lvm
```

5. Execute the srvAlloc command with the following parameters.  
service=concentrator volume=concentrator commit=1

## Configure Storage for Concentrator using NW-PV-D/NW-PV-D-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Log Decoder Hybrid physical host.

1. Execute the `raidList` command.

a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

In Use: FALSE

Devices: <empty>

b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList Parameters:  Send

Message Help

raidList: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 02:00.0, Enclosure 32
Vendor: DP
Model: BP13G+EXP
In Use: true
Drives: 931.512 GB HDD x 2
 1.819 TB HDD x 2
Devices: sda /dev/disk/by-path/pci-0000:02:00.0-scsi-0:2:0:0
 sdb /dev/disk/by-path/pci-0000:02:00.0-scsi-0:2:1:0

Controller 1 at PCI Address 03:00.0, Enclosure 108
Vendor: DELL
Model: MD1400
In Use: false
Drives: 1.455 TB SSD x 3
 10.692 TB HDD x 9
Devices:

```

- Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=108 scheme=concentrator commit=1`

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance  
 raidNew Parameters:

Message Help

raidNew: allocate RAID devices in a drive shelf  
 security.roles: appliance.manage  
 parameters:  
 controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to  
 enclosure - <uint32, optional, {enum-one:The value must be one of the following: 32|108}> Enclosure number of the shelf to clear. Required if  
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-in  
 expansion|decoder-hotspare|logdecoder-hotspare}> Type of RAID volumes to allocate

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=108&scheme=concentrator&commit=1

Output (or command manual help)

```
/opt/MegaRAID/perccli/perccli164 /c1 add vd r5 drives=108:3,108:1,108:2 ra Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded

/opt/MegaRAID/perccli/perccli164 /c1 add vd r6 drives=108:3,108:4,108:5,108:6,108:7,108:8,108:9,108:10,108:11 ra Strip=128
Controller = 1
Status = Success
Description = Add VD Succeeded
```

```
[root@Concentrator132 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
│ └─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
└─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 1.8T 0 disk
├─sdb1 8:17 0 1.8T 0 part
└─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
sdc 8:32 0 2.9T 0 disk
sdd 8:48 0 74.9T 0 disk
```

- Execute the `partNew` command to create the **concentrator** partition first with the following parameters. You must create the concentrator volume before index volume or it will fail.

`name=sdd service=concentrator volume=concentrator commit=1`

```

Properties for /appliance
partNew Parameters: name=sdd service=concentrator volume=concentrator commit=1 Send

Message Help
security.roles: appliance.manage
parameters:
 name - <string, (enum-one:The value must be one of the following: sdc|sdd)> block device name
 service - <string, (enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder)> service that will use storage
 volume - <string, optional, (enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex)> volume to create
 commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f concentrator /dev/sdd1
Volume group "concentrator" successfully created
/sbin/lvcreate -y -n root -l 30G concentrator
Wiping xfs signature on /dev/concentrator/root.
Logical volume "root" created.
/sbin/mkfs.xfs /dev/concentrator/root
meta-data=/dev/concentrator/root isize=512 attr=4, agsize=1966080 blks
= sectsz=4096 attr=2, projid32bit=1
= crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=7864320, inapct=25
= sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=3840, version=2
= sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator
/bin/mount /var/netwitness/concentrator
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator

[root@Concentrator132 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ └─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ └─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
│ └─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 1.8T 0 disk
├─sdb1 8:17 0 1.8T 0 part
│ └─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
sdc 8:32 0 2.9T 0 disk
sdd 8:48 0 74.9T 0 disk
├─sdd1 8:49 0 74.9T 0 part
│ └─concentrator-root 253:5 0 30G 0 lvm /var/netwitness/concentrator
│ └─concentrator-sessiondb 253:6 0 7.5T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-metadb 253:7 0 67.3T 0 lvm /var/netwitness/concentrator/metadb

```

4. Execute the `partNew` command with the following parameters to create an **index** on SSDs.

```
name=sdc service=concentrator volume=index commit=1
```

```

Properties for /appliance
partNew Parameters: name=sdsc service=concentrator volume=index commit=1 Send

Message Help
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdc|sdd}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create
 commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdsc&service=concentrator&volume=index&commit=1

Output (or command manual help)

/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart lvm 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f index /dev/sdc1
Volume group "index" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index
Wiping xfs signature on /dev/index/index.
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index isize=512 agcount=4, agsize=195280640 blks
 = sectsz=4096 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=781122560, imaxpct=5
 = sunit=0 swidth=0 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=381407, version=2
 = sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index
/bin/mount /var/netwitness/concentrator/index

[root@Concentrator132 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 931G 0 disk
├─sda1 8:1 0 1G 0 part /boot
├─sda2 8:2 0 930G 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 1.8T 0 disk
├─sdb1 8:17 0 1.8T 0 part
│ └─netwitness_vg00-nwhome 253:2 0 2.7T 0 lvm /var/netwitness
sdc 8:32 0 2.9T 0 disk
├─sdc1 8:33 0 2.9T 0 part
│ └─index-index 253:8 0 2.9T 0 lvm /var/netwitness/concentrator/index
sdd 8:48 0 74.9T 0 disk
├─sdd1 8:49 0 74.9T 0 part
│ ├─concentrator-root 253:5 0 30G 0 lvm /var/netwitness/concentrator
│ ├─concentrator-sessiondb 253:6 0 7.5T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-metadb 253:7 0 67.3T 0 lvm /var/netwitness/concentrator/metadb

[root@Concentrator132 ~]#

[root@conc95 ~]# df -kh
Filesystem Size Used Avail Use% Mounted on
devtmpfs 63G 0 63G 0% /dev
tmpfs 63G 420K 63G 1% /dev/shm
tmpfs 63G 43M 63G 1% /run
tmpfs 63G 0 63G 0% /sys/fs/cgroup
/dev/mapper/netwitness_vg00-root 30G 3.6G 27G 12% /
/dev/sdal 1014M 122M 893M 12% /boot
/dev/mapper/netwitness_vg00-varlog 10G 156M 9.9G 2% /var/log
/dev/mapper/netwitness_vg00-usrhome 10G 33M 10G 1% /home
/dev/mapper/netwitness_vg00-nwhome 3.3T 494M 3.3T 1% /var/netwitness
/dev/mapper/concentrator-root 30G 61M 30G 1% /var/netwitness/concentrator
/dev/mapper/concentrator-sessiondb 3.0T 34M 3.0T 1% /var/netwitness/concentrator/sessiondb
/dev/mapper/concentrator-metadb 27T 34M 27T 1% /var/netwitness/concentrator/metadb
tmpfs 13G 0 13G 0% /run/user/0

```

- Execute the `srvAlloc` command with the following parameters.  
`service=concentrator volume=concentrator commit=1`

## Configure Storage for Log Hybrid using NW-PV-A/NW-PV-A-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Log Hybrid physical host.

1. Execute the `raidList` command

- Record the Controller Number, Enclosure Number, In Use, Drives, and Devices. You should see the following information.

In Use: FALSE

Devices: <empty>

- Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

```

Message Help
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

Controller 0 at PCI Address 3c:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 1.746 TB SSD x 2
 2.182 TB HDD x 2
 7.277 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:0:0
 sdd /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:1:0
 sde /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:2:0
 sdf /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:3:0
 sdg /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:4:0

Controller 1 at PCI Address af:00.0, Enclosure 72, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 10.692 TB HDD x 12
Devices:

```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=72 scheme=log-hybrid preferSecure=1 commit=1`

- The following example illustrates what you should see after you create a RAID array.

```

Properties for /appliance
raidNew Parameters: controller=1 enclosure=72 scheme=log-hybrid preferSecure=1 commit=1 Send

Message Help
parameters:
 controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
 enclosure - <uint32, optional, {enum-one:The value must be one of the following: 64,72}> Enclosure number of the shelf to clear. Required if the controller has
more than one enclosure attached.
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-
decoder|endpoint-log-hybrid|packet-expansion|decoder-hotspare|logdecoder-hotspare}> Type of RAID volumes to allocate
 preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off}> Prefer creation of a
secure array given compatible physical drives and a controller with a security key set

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=72&scheme=log-hybrid&preferSecure=1&commit=1

Output (or command manual help)

/opt/MegaRAID/percli/percli64 /c1 add vd r5 drives=72:0,72:1,72:2,72:3,72:4,72:5 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

/opt/MegaRAID/percli/percli64 /c1 add vd r5 drives=72:6,72:7,72:8,72:9,72:10,72:11 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

```

- Execute the `devlist` command to find the exact memory distribution between the newly created RAIDS,

in log-hybrid the distribution is exactly equal so any service can be installed in any of the

## RAIDS.

Properties for /appliance  
 devlist  Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sda: vendor=DELL model="PERC H740P Mini" size="2.18 TB" used=1
sdb: vendor=DELL model="PERC H840 Adp" size="53.46 TB" used=1
sdc: vendor=DELL model="PERC H840 Adp" size="53.46 TB" used=1
sdd: vendor=DELL model="PERC H740P Mini" size="7.28 TB" used=1
sde: vendor=DELL model="PERC H740P Mini" size="21.83 TB" used=1
sdf: vendor=DELL model="PERC H740P Mini" size="21.83 TB" used=1
sdg: vendor=DELL model="PERC H740P Mini" size="1.75 TB" used=1
```

3. Execute the partNew command with the following parameters to create partitions and mount points in the /etc/fstab file.

- a. name=sdb service=logdecoder volume=logdecoder commit=1

Properties for /appliance  
 partNew  Parameters: name=sdb service=logdecoder volume=logdecoder commit=1

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdb sdc}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, {optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-
meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}}> volume to create
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdb&service=logdecoder&volume=logdecoder&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdb mklabel gpt
/sbin/parted -s -a optimal /dev/sdb mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdb1
Physical volume "/dev/sdb1" successfully created.
/sbin/vgcreate -f logdecoder0 /dev/sdb1
Volume group "logdecoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE logdecoder0
Wiping xfs signature on /dev/logdecoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/logdecoder0/packetdb
meta-data=/dev/logdecoder0/packetdb isize=512 agcount=54, agsize=268435424 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=14350416896, imaxpct=1
 = sunit=32 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/logdecoder/packetdb
/bin/mount /var/netwitness/logdecoder/packetdb
```

- b. name=sdc service=concentrator volume=concentrator commit=1

Properties for /appliance  
partNew Parameters: name=sdc service=concentrator volume=concentrator commit=1 Send

Message Help

partNew: make partitions on a block device  
security.roles: appliance.manage  
parameters:  
name - <string, {enum-one:The value must be one of the following: sdb|sdc}> block\_device\_name  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create

/appliance?msg=partNew&force-content-type=text/plain&name=sdc&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdc1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512 agcount=32, agsize=44845024 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=1435040768, imaxpct=5
 = sunit=32 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb0
/bin/mount /var/netwitness/concentrator/sessiondb0
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to logdecoder and concentrator services.

- a. This adds storage to the logdecoder service configuration and restarts the service every time it is executed.

service=logdecoder volume=logdecoder0 commit=1

Properties for /appliance  
srvAlloc Parameters: service=logdecoder volume=logdecoder0 commit=1 Send

Message Help

srvAlloc: apply volume group storage to a service on this appliance  
security.roles: appliance.manage  
parameters:  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, {enum-one:The value must be one of the following: concentrator0|hybrid-concentrator|hybrid-logdecoder-meta|index|logdecoder0|logpacket|netwitness\_vg00}> volume group name  
commit - <bool, optional> commit changes

/appliance?msg=srvAlloc&force-content-type=text/plain&service=logdecoder&volume=logdecoder0&commit=1

Output (or command manual help)

```
Set /database/config/packet.dir to /var/netwitness/logdecoder/packetdb/packetdb=20.74 TB;/var/netwitness/logdecoder/packetdb0=50.78 TB
```

- b. This adds storage to concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance

Parameters:

Message Help

srvAlloc: apply volume group storage to a service on this appliance  
security.roles: appliance.manage  
parameters:  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, {enum-one:The value must be one of the following: concentrator0|hybrid-concentrator|hybrid-logdecoder-  
meta|index|logdecoder0|logpacket|netwitness\_vg00}> volume group name  
commit - <bool, optional> commit changes

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

```
Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==18.85 TB;/var/netwitness/concentrator/metadb==45.71 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==1.86 TB;/var/netwitness/concentrator/sessiondb==5.08 TB
```

- Execute the `lsblk` command in backed to see all the raids and partitions inside the service.

```
[root@loghybrid ~]# lsblk
```

| NAME                                | MAJ:MIN | RM | SIZE  | RO | TYPE | MOUNTPOINT                              |
|-------------------------------------|---------|----|-------|----|------|-----------------------------------------|
| sda                                 | 8:0     | 0  | 2.2T  | 0  | disk |                                         |
| └─sda1                              | 8:1     | 0  | 1M    | 0  | part |                                         |
| └─sda2                              | 8:2     | 0  | 1G    | 0  | part | /boot                                   |
| └─sda3                              | 8:3     | 0  | 2.2T  | 0  | part |                                         |
| └─netwitness_vg00-root              | 253:0   | 0  | 30G   | 0  | lvm  | /                                       |
| └─netwitness_vg00-swap              | 253:1   | 0  | 4G    | 0  | lvm  | [SWAP]                                  |
| └─netwitness_vg00-nwhome            | 253:5   | 0  | 2.1T  | 0  | lvm  | /var/netwitness                         |
| └─netwitness_vg00-varlog            | 253:6   | 0  | 10G   | 0  | lvm  | /var/log                                |
| └─netwitness_vg00-usrhome           | 253:7   | 0  | 10G   | 0  | lvm  | /home                                   |
| sdb                                 | 8:16    | 0  | 53.5T | 0  | disk |                                         |
| └─sdb1                              | 8:17    | 0  | 53.5T | 0  | part |                                         |
| └─logdecoder0-packetdb              | 253:2   | 0  | 53.5T | 0  | lvm  | /var/netwitness/logdecoder/packetdb0    |
| sdc                                 | 8:32    | 0  | 53.5T | 0  | disk |                                         |
| └─sdc1                              | 8:33    | 0  | 53.5T | 0  | part |                                         |
| └─concentrator0-sessiondb           | 253:3   | 0  | 5.4T  | 0  | lvm  | /var/netwitness/concentrator/sessiondb0 |
| └─concentrator0-metadb              | 253:4   | 0  | 48.1T | 0  | lvm  | /var/netwitness/concentrator/metadb0    |
| sdd                                 | 8:48    | 0  | 7.3T  | 0  | disk |                                         |
| └─sdd1                              | 8:49    | 0  | 7.3T  | 0  | part |                                         |
| └─hybrid--logdecoder--meta-decoroot | 253:8   | 0  | 7.3T  | 0  | lvm  | /var/netwitness/logdecoder              |
| sde                                 | 8:64    | 0  | 21.9T | 0  | disk |                                         |
| └─sde1                              | 8:65    | 0  | 21.9T | 0  | part |                                         |
| └─logpacket-packetdb                | 253:9   | 0  | 21.9T | 0  | lvm  | /var/netwitness/logdecoder/packetdb     |
| sdf                                 | 8:80    | 0  | 21.9T | 0  | disk |                                         |
| └─sdf1                              | 8:81    | 0  | 21.9T | 0  | part |                                         |
| └─hybrid--concentrator-root         | 253:10  | 0  | 21.9T | 0  | lvm  | /var/netwitness/concentrator            |
| sdg                                 | 8:96    | 0  | 1.8T  | 0  | disk |                                         |
| └─sdg1                              | 8:97    | 0  | 1.8T  | 0  | part |                                         |
| └─index-index                       | 253:11  | 0  | 1.8T  | 0  | lvm  | /var/netwitness/concentrator/index      |

## Configure Storage for Network Hybrid using NW-PV-A/NW-PV-A-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Network Hybrid physical host.

- Execute the `raidList` command
  - Record the Controller Number, Enclosure Number, In Use, Drives, and Devices. You should see the following information.  
In Use: FALSE  
Devices: <empty>
  - Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList ▼ Parameters:  Send

Message Help

raidList: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 3c:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 1.746 TB SSD x 2
 2.182 TB HDD x 2
 7.277 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:0:0
 sdd /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:1:0
 sde /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:2:0
 sdf /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:3:0
 sdg /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:4:0

Controller 1 at PCI Address af:00.0, Enclosure 72, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 10.692 TB HDD x 12
Devices:

```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=72 scheme=network-hybrid preferSecure=1 commit=1`

- a. The following example illustrates what you should see after you create a RAID array.

Properties for /appliance  
 raidNew ▼ Parameters:  Send

Message Help

scheme - <string, (enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-decoder|endpoint-log-hybrid|packet-expansion|decoder-hotspare|logdecoder-hotspare)> Type of RAID volumes to allocate  
 preferSecure - <bool, optional, (bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off)> Prefer creation of a secure array given compatible physical drives and a controller with a security key set  
 commit - <bool, optional> commit changes

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=72&scheme=network-hybrid&preferSecure=1&commit=1

Output (or command manual help)

```

/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=72:0,72:1,72:2 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

/opt/MegaRAID/perccli/perccli64 /c1 add vd r5 drives=72:3,72:4,72:5,72:6,72:7,72:8,72:9,72:10,72:11 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

```

- b. Execute `devlist` command to find the exact memory distribution between the newly created raids , in network hybrid the memory for 'sdc' raid is more than 'sdb' raid ,therefore packetdb will be installed in the raid with higher memory allocated that is 'sdc' .

Properties for /appliance

devlist Parameters:  Send

Message Help

devlist: list storage devices  
security.roles: appliance.manage

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sda: vendor=DELL model="PERC H740P Mini" size="2.18 TB" used=1
sdb: vendor=DELL model="PERC H840 Adp" size="21.38 TB" used=0
sdc: vendor=DELL model="PERC H840 Adp" size="85.54 TB" used=0
sdd: vendor=DELL model="PERC H740P Mini" size="7.28 TB" used=1
sde: vendor=DELL model="PERC H740P Mini" size="29.11 TB" used=1
sdf: vendor=DELL model="PERC H740P Mini" size="14.55 TB" used=1
sdg: vendor=DELL model="PERC H740P Mini" size="1.75 TB" used=1
```

- c. Execute the `lsblk` command to list all the raids that are newly created.

```
[root@networkhybrid ~]# lsblk
```

| NAME                            | MAJ:MIN | RM | SIZE  | RO | TYPE | MOUNTPOINT                         |
|---------------------------------|---------|----|-------|----|------|------------------------------------|
| sda                             | 8:0     | 0  | 2.2T  | 0  | disk |                                    |
| └sda1                           | 8:1     | 0  | 1M    | 0  | part |                                    |
| └sda2                           | 8:2     | 0  | 1G    | 0  | part | /boot                              |
| └sda3                           | 8:3     | 0  | 2.2T  | 0  | part |                                    |
| ├netwitness_vg00-root           | 253:0   | 0  | 30G   | 0  | lvm  | /                                  |
| ├netwitness_vg00-swap           | 253:1   | 0  | 4G    | 0  | lvm  | [SWAP]                             |
| ├netwitness_vg00-nwhome         | 253:5   | 0  | 2.1T  | 0  | lvm  | /var/netwitness                    |
| ├netwitness_vg00-varlog         | 253:6   | 0  | 10G   | 0  | lvm  | /var/log                           |
| └netwitness_vg00-usrhome        | 253:7   | 0  | 10G   | 0  | lvm  | /home                              |
| sdb                             | 8:16    | 0  | 21.4T | 0  | disk |                                    |
| sdc                             | 8:32    | 0  | 85.5T | 0  | disk |                                    |
| sdd                             | 8:48    | 0  | 7.3T  | 0  | disk |                                    |
| └sddl                           | 8:49    | 0  | 7.3T  | 0  | part |                                    |
| ├hybrid--decoder--meta-decoroot | 253:8   | 0  | 7.3T  | 0  | lvm  | /var/netwitness/decoder            |
| sde                             | 8:64    | 0  | 29.1T | 0  | disk |                                    |
| └sdel                           | 8:65    | 0  | 29.1T | 0  | part |                                    |
| ├packet-packetdb                | 253:9   | 0  | 29.1T | 0  | lvm  | /var/netwitness/decoder/packetdb   |
| sdf                             | 8:80    | 0  | 14.6T | 0  | disk |                                    |
| └sdfl                           | 8:81    | 0  | 14.6T | 0  | part |                                    |
| ├hybrid--concentrator-root      | 253:10  | 0  | 14.6T | 0  | lvm  | /var/netwitness/concentrator       |
| sdg                             | 8:96    | 0  | 1.8T  | 0  | disk |                                    |
| └sdgl                           | 8:97    | 0  | 1.8T  | 0  | part |                                    |
| ├index-index                    | 253:11  | 0  | 1.8T  | 0  | lvm  | /var/netwitness/concentrator/index |

3. Execute the `partNew` command with the following parameters to create partitions and mount points in the `/etc/fstab` file.

- a. name=sdb service=concentrator volume=concentrator commit=1

Properties for /appliance  
partNew Parameters: name=sdb service=concentrator volume=concentrator commit=1 Send

Message Help  
security.roles: appliance.manage  
parameters:  
name - <string, {enum-one:The value must be one of the following: sdb|sdc}> block device name  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create  
commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdb&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdb mklabel gpt
/sbin/parted -s -a optimal /dev/sdb mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdb1
Physical volume "/dev/sdb1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdb1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512 agcount=32, agsize=17938016 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=574016512, imaxpct=5
 = sunit=32 swidth=64 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=280288, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb0
/bin/mount /var/netwitness/concentrator/sessiondb0
```

- b. name=sdc service=decoder volume=decoder commit=1

Properties for /appliance  
partNew Parameters: name=sdc service=decoder volume=decoder commit=1 Send

Message Help  
security.roles: appliance.manage  
parameters:  
name - <string, {enum-one:The value must be one of the following: sdb|sdc}> block device name  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create  
commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdc&service=decoder&volume=decoder&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f decoder0 /dev/sdc1
Volume group "decoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder0
Wiping xfs signature on /dev/decoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder0/packetdb
meta-data=/dev/decoder0/packetdb isize=512 agcount=86, agsize=268435424 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=22960667648, imaxpct=1
 = sunit=32 swidth=256 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb0
/bin/mount /var/netwitness/decoder/packetdb0
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to decoder and concentrator services.

- a. This adds storage to the decoder service configuration and restarts the service every time it is executed.

service=decoder volume=decoder0 commit=1

Properties for /appliance  
 srvAlloc  Parameters:

Message Help

srvAlloc: apply volume group storage to a service on this appliance  
 security.roles: appliance.manage  
 parameters:  
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
 volume - <string, {enum-one:The value must be one of the following: concentrator0|decoder0|hybrid-concentrator|hybrid-decoder-meta|index|netwitness\_vg00|packet}>  
 volume group name  
 commit - <bool, optional> commit changes

/appliance?msg=srvAlloc&force-content-type=text/plain&service=decoder&volume=decoder0&commit=1

Output (or command manual help)

Set /database/config/packet.dir to /var/netwitness/decoder/packetdb/packetdb==27.65 TB; /var/netwitness/decoder/packetdb0==81.26 TB

- b. This adds storage to the concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance  
 srvAlloc  Parameters:

Message Help

srvAlloc: apply volume group storage to a service on this appliance  
 security.roles: appliance.manage  
 parameters:  
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
 volume - <string, {enum-one:The value must be one of the following: concentrator0|decoder0|hybrid-concentrator|hybrid-decoder-meta|index|netwitness\_vg00|packet}>  
 volume group name  
 commit - <bool, optional> commit changes

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==12.55 TB; /var/netwitness/concentrator/metadb0==18.28 TB  
 Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==1.24 TB; /var/netwitness/concentrator/sessiondb0==2.03 TB

5. Execute the `lsblk` command to list all the raids and partition in the service.

```
[root@networkhybrid ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 1M 0 part
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ └─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ └─netwitness_vg00-nwhome 253:5 0 2.1T 0 lvm /var/netwitness
│ └─netwitness_vg00-varlog 253:6 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:7 0 10G 0 lvm /home
sdb 8:16 0 21.4T 0 disk
├─sdb1 8:17 0 21.4T 0 part
│ └─concentrator0-sessiondb 253:2 0 2.1T 0 lvm /var/netwitness/concentrator/sessiondb0
│ └─concentrator0-metadb 253:3 0 19.3T 0 lvm /var/netwitness/concentrator/metadb0
sdc 8:32 0 85.5T 0 disk
├─sdc1 8:33 0 85.5T 0 part
│ └─decoder0-packetdb 253:4 0 85.5T 0 lvm /var/netwitness/decoder/packetdb0
sdd 8:48 0 7.3T 0 disk
├─sddl 8:49 0 7.3T 0 part
│ └─hybrid--decoder--meta-decoroot 253:8 0 7.3T 0 lvm /var/netwitness/decoder
sde 8:64 0 29.1T 0 disk
├─sde1 8:65 0 29.1T 0 part
│ └─packet-packetdb 253:9 0 29.1T 0 lvm /var/netwitness/decoder/packetdb
sdf 8:80 0 14.6T 0 disk
├─sdf1 8:81 0 14.6T 0 part
│ └─hybrid--concentrator-root 253:10 0 14.6T 0 lvm /var/netwitness/concentrator
sdg 8:96 0 1.8T 0 disk
├─sdgl 8:97 0 1.8T 0 part
│ └─index-index 253:11 0 1.8T 0 lvm /var/netwitness/concentrator/index
```

## Configure Storage for Endpoint Log Hybrid using NW-PV-A/NW-PV-A-N

The following scenario configures storage on one, non-encrypted, 12-Drive PowerVault for a Endpoint Log Hybrid physical host.

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices. You should see the following information.
 

```
In Use: FALSE
Devices: <empty>
```
  - b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList ▼ Parameters:  Send

Message Help

raidlist: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 3c:00.0, Enclosure 64, SCSI Channel 2
Vendor: DP
Model: BP14G+EXP
In Use: true
Drives: 1.746 TB SSD x 2
 2.182 TB HDD x 2
 7.277 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:0:0
 sdd /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:1:0
 sde /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:2:0
 sdf /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:3:0
 sdg /dev/disk/by-path/pci-0000:3c:00.0-scsi-0:2:4:0

Controller 1 at PCI Address af:00.0, Enclosure 72, SCSI Channel 2
Vendor: DELL
Model: MD1400
In Use: false
Drives: 10.692 TB HDD x 12
Devices:

```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=72 scheme=log-hybrid preferSecure=1 commit=1`

- a. The following example illustrates what you should see after you create a RAID array.

Properties for /appliance  
 raidNew ▼ Parameters: controller=1 enclosure=72 scheme=log-hybrid preferSecure=1 commit=1 Send

Message Help

Parameters:  
 controller - <uint32, {enum-one:The value must be one of the following: 0, 1}> Controller the shelf is attached to  
 enclosure - <uint32, optional, {enum-one:The value must be one of the following: 64, 72}> Enclosure number of the shelf to clear. Required if the controller has more than one enclosure attached.  
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-decoder|endpoint-log-hybrid|packet-expansion|decoder-hotspare|logdecoder-hotspare}> Type of RAID volumes to allocate  
 preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0, 1, yes, no, true, false, on, off}> Prefer creation of a secure array given compatible physical drives and a controller with a security key set

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=72&scheme=log-hybrid&preferSecure=1&commit=1

Output (or command manual help)

```

/opt/MegaRAID/percccli/percccli64 /c1 add vd r5 drives=72:0,72:1,72:2,72:3,72:4,72:5 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

/opt/MegaRAID/percccli/percccli64 /c1 add vd r5 drives=72:6,72:7,72:8,72:9,72:10,72:11 ra Strip=128
CLI Version = 007.1623.0000.0000 May 17, 2021
Operating system = Linux 3.10.0-1160.83.1.el7.x86_64
Controller = 1
Status = Success
Description = Add VD Succeeded.

```

- b. Execute the `devlist` command to find the exact memory distribution for the newly created raids, in endpoint-log-hybrid both are exactly equal, therefore any service can be installed in any of the raid.

Properties for /appliance

devlist  Parameters:

Message Help

devlist: list storage devices  
security.roles: appliance.manage

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sda: vendor=DELL model="PERC H740P Mini" size="2.18 TB" used=1
sdb: vendor=DELL model="PERC H840 Adp" size="53.46 TB" used=0
sdc: vendor=DELL model="PERC H840 Adp" size="53.46 TB" used=0
sdd: vendor=DELL model="PERC H740P Mini" size="7.28 TB" used=1
sde: vendor=DELL model="PERC H740P Mini" size="14.55 TB" used=1
sdf: vendor=DELL model="PERC H740P Mini" size="21.83 TB" used=1
sdg: vendor=DELL model="PERC H740P Mini" size="1.75 TB" used=1
```

- c. Execute the `lsblk` command at backend to list the newly created raids.

```
[root@endpointloghybrid ~]# lsblk
```

| NAME                                | MAJ:MIN | RM | SIZE  | RO | TYPE | MOUNTPOINT                          |
|-------------------------------------|---------|----|-------|----|------|-------------------------------------|
| sda                                 | 8:0     | 0  | 2.2T  | 0  | disk |                                     |
| └─sda1                              | 8:1     | 0  | 1M    | 0  | part |                                     |
| └─sda2                              | 8:2     | 0  | 1G    | 0  | part | /boot                               |
| └─sda3                              | 8:3     | 0  | 2.2T  | 0  | part |                                     |
| └─netwitness_vg00-root              | 253:0   | 0  | 30G   | 0  | lvm  | /                                   |
| └─netwitness_vg00-swap              | 253:1   | 0  | 4G    | 0  | lvm  | [SWAP]                              |
| └─netwitness_vg00-nwhome            | 253:2   | 0  | 2.1T  | 0  | lvm  | /var/netwitness                     |
| └─netwitness_vg00-varlog            | 253:3   | 0  | 10G   | 0  | lvm  | /var/log                            |
| └─netwitness_vg00-usrhome           | 253:4   | 0  | 10G   | 0  | lvm  | /home                               |
| sdb                                 | 8:16    | 0  | 53.5T | 0  | disk |                                     |
| sdc                                 | 8:32    | 0  | 53.5T | 0  | disk |                                     |
| sdd                                 | 8:48    | 0  | 7.3T  | 0  | disk |                                     |
| └─sdd1                              | 8:49    | 0  | 7.3T  | 0  | part |                                     |
| └─hybrid--logdecoder--meta-decoroot | 253:8   | 0  | 7.3T  | 0  | lvm  | /var/netwitness/logdecoder          |
| sde                                 | 8:64    | 0  | 14.6T | 0  | disk |                                     |
| └─sde1                              | 8:65    | 0  | 14.6T | 0  | part |                                     |
| └─endpoint--log--hybrid-mongo       | 253:9   | 0  | 7.3T  | 0  | lvm  | /var/netwitness/mongo               |
| └─endpoint--log--hybrid-packetdb    | 253:10  | 0  | 7.3T  | 0  | lvm  | /var/netwitness/logdecoder/packetdb |
| sdf                                 | 8:80    | 0  | 21.9T | 0  | disk |                                     |
| └─sdf1                              | 8:81    | 0  | 21.9T | 0  | part |                                     |
| └─hybrid--concentrator-root         | 253:11  | 0  | 21.9T | 0  | lvm  | /var/netwitness/concentrator        |
| sdg                                 | 8:96    | 0  | 1.8T  | 0  | disk |                                     |
| └─sdg1                              | 8:97    | 0  | 1.8T  | 0  | part |                                     |
| └─index-index                       | 253:12  | 0  | 1.8T  | 0  | lvm  | /var/netwitness/concentrator/index  |

3. Execute the `partNew` command with the following parameters to create partitions and mount points in the `/etc/fstab` file.

- a. name=sdb service=logdecoder volume=logdecoder commit=1

Properties for /appliance  
partNew Parameters: name=sdb service=logdecoder volume=logdecoder commit=1 Send

Message Help

name - <string, {enum-one:The value must be one of the following: sdb|sdc}> block device name  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decoder|small|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create  
commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdb&service=logdecoder&volume=logdecoder&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdb mklabel gpt
/sbin/parted -s -a optimal /dev/sdb mkpart LVN 0% 100%
/sbin/pvcreate -f /dev/sdb1
Physical volume "/dev/sdb1" successfully created.
/sbin/vgcreate -f logdecoder0 /dev/sdb1
Volume group "logdecoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE logdecoder0
Wiping xfs signature on /dev/logdecoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/logdecoder0/packetdb
meta-data=/dev/logdecoder0/packetdb isize=512 agcount=54, agsize=268435424 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=14350416896, imaxpct=1
 = sunit=32 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/logdecoder/packetdb
/bin/mount /var/netwitness/logdecoder/packetdb
```

- b. name=sdc service=concentrator volume=concentrator commit=1

Properties for /appliance  
partNew Parameters: name=sdc service=concentrator volume=concentrator commit=1 Send

Message Help

name - <string, {enum-one:The value must be one of the following: sdb|sdc}> block device name  
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage  
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decoder|small|decoder|packet|hybrid-decoder-meta|logdecodersmall|logdecoder|logpacket|hybrid-logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create  
commit - <bool, optional> commit changes

/appliance?msg=partNew&force-content-type=text/plain&name=sdc&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVN 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdc1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512 agcount=32, agsize=44845024 blks
 = sectsz=512 attr=2, projid32bit=1
 = crc=1 finobt=0, sparse=0
data = bsize=4096 blocks=1435040768, imaxpct=5
 = sunit=32 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0 ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=512 sunit=32 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb
/bin/mount /var/netwitness/concentrator/sessiondb
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to logdecoder and concentrator services.

- a. This adds storage to the logdecoder service configuration and restarts the service every time it is executed.

service=logdecoder volume=logdecoder0 commit=1

Properties for /appliance  
 srvAlloc ▾ Parameters: service=logdecoder volume=logdecoder0 commit=1

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, {enum-one:The value must be one of the following: concentrator0|endpoint-log-hybrid|hybrid-concentrator|hybrid-logdecoder-
meta|index|logdecoder0|netwitness_vg00}> volume group name
 commit - <bool, optional> commit changes

```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=logdecoder&volume=logdecoder0&commit=1

Output (or command manual help)

```

Set /database/config/packet.dir to /var/netwitness/logdecoder/packetdb/packetdb==6.91 TB;/var/netwitness/logdecoder/packetdb==50.78 TB

```

- b. This adds storage to the concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance  
 srvAlloc ▾ Parameters: service=concentrator volume=concentrator0 commit=1

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, {enum-one:The value must be one of the following: concentrator0|endpoint-log-hybrid|hybrid-concentrator|hybrid-logdecoder-
meta|index|logdecoder0|netwitness_vg00}> volume group name
 commit - <bool, optional> commit changes

```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

```

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==18.85 TB;/var/netwitness/concentrator/metadb==45.71 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==1.86 TB;/var/netwitness/concentrator/sessiondb==5.08 TB

```

5. Execute `lsblk` command to list all the raids and partitions in the service.

```

[root@endpointloghybrid ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 1M 0 part
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ └─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ └─netwitness_vg00-nwhome 253:2 0 2.1T 0 lvm /var/netwitness
│ └─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 53.5T 0 disk
├─sdb1 8:17 0 53.5T 0 part
│ └─logdecoder0-packetdb 253:5 0 53.5T 0 lvm /var/netwitness/logdecoder/packetdb0
sdc 8:32 0 53.5T 0 disk
├─sdc1 8:33 0 53.5T 0 part
│ └─concentrator0-sessiondb 253:6 0 5.4T 0 lvm /var/netwitness/concentrator/sessiondb0
│ └─concentrator0-metadb 253:7 0 48.1T 0 lvm /var/netwitness/concentrator/metadb0
sdd 8:48 0 7.3T 0 disk
├─sdd1 8:49 0 7.3T 0 part
│ └─hybrid--logdecoder--meta-decoroot 253:8 0 7.3T 0 lvm /var/netwitness/logdecoder
sde 8:64 0 14.6T 0 disk
├─sde1 8:65 0 14.6T 0 part
│ └─endpoint--log--hybrid-mongo 253:9 0 7.3T 0 lvm /var/netwitness/mongo
│ └─endpoint--log--hybrid-packetdb 253:10 0 7.3T 0 lvm /var/netwitness/logdecoder/packetdb
sdf 8:80 0 21.9T 0 disk
├─sdf1 8:81 0 21.9T 0 part
│ └─hybrid--concentrator-root 253:11 0 21.9T 0 lvm /var/netwitness/concentrator
sdg 8:96 0 1.8T 0 disk
├─sdg1 8:97 0 1.8T 0 part
│ └─index-index 253:12 0 1.8T 0 lvm /var/netwitness/concentrator/index

```

## Appendix F. Sample Storage Configuration Scenarios for S7 Physical Hosts with 12-Drive PowerVault MD2412

This appendix illustrates the following examples of how to configure storage on unencrypted MD2412 PowerVault (12 \* 16 TB Self Encrypting Drives (SEDs)) external storage devices attached to S7 physical hosts.

- [Configure Storage for Archiver using MD2412](#)
- [Configure Storage for Decoder using MD2412](#)
- [Configure Storage for Concentrator using MD2412](#)
- [Configure Storage for Log Hybrid using MD2412](#)
- [Configure Storage for Network Hybrid using MD2412](#)
- [Configure Storage for Endpoint Log Hybrid using MD2412](#)

**Note:**

The [Enable Encryption on Series 6E or Series 7 Hosts](#) topic describes the process to turn on encryption on NW servers and configured storage devices (MD2412).

If the user chooses to enable or turn on encryption before configuring storage (MD2412), the security on the controllers (internal and external) must be enabled and controller security key must be set. To enable the security on the controllers (internal and external) and set the controller security key, refer to *Chapter 8 ->Security Key and RAID Management ->Create a security key in the Dell PowerEdge RAID Controller 12 User's Guide*.

Once the security key is set, to enable encryption on the external storage, use `preferSecure=1` when executing `raidNew` command to create RAID.

Alternatively, configure the external storage with `preferSecure=0` when creating RAID and then follow the steps in [Enable Encryption on Series 6E or Series 7 Hosts](#) to turn on encryption.

### Configure Storage for Archiver using MD2412

The following scenario configures storage on one, unencrypted, 12-Drive MD2412 PowerVault for an Archiver physical host.

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>

## b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance

raidList  Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 2.182 TB HDD x 4
Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 2
Vendor: DELL
Model: DSES Enclosure
In Use: false
Drives: 14.552 TB HDD x 12
Devices:
```

- Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

```
controller=1 enclosure=386 scheme=archiver preferSecure=0 commit=1
```

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance  
 raidNew Parameters: controller=1 enclosure=386 scheme=archiver preferSecure=0 commit=1

Message Help

raidNew: allocate RAID devices in a drive shelf  
 security.roles: appliance.manage  
 parameters:  
   controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to  
   enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386}> Enclosure number of the shelf to clear. Required  
   scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=386&scheme=archiver&preferSecure=0&commit=1

Output (or command manual help)

```
/opt/MegaRAID/perccli2/perccli2 /c1 add vd r6 drives=386:0,386:1,386:2,386:3,386:4,386:5,386:6,386:7,386:8,386:9,386:10,386:11 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 1
Status = Success
Description = Add VD succeeded.
```

VD Information :

```
=====
```

| VDID | VDSize      | Status  | ErrType | ErrCd | Msg |
|------|-------------|---------|---------|-------|-----|
| 1    | 145.522 TiB | Success | -       | -     | -   |

- Execute the `raidList` command to verify the new RAID array.

In Use: TRUE

Devices: <device> (for example, sdc)

Properties for /appliance

raidList ▾ Parameters: 

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

</appliance?msg=raidList&force-content-type=text/plain>

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0
```

4. Execute the `devlist` command to find the exact memory distribution for the newly created raids.

Properties for /appliance

devlist  Parameters: 

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&amp;force-content-type=text/plain

Output (or command manual help)

```
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="145.52 TB" used=0
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

- Execute the `lsblk` command at backend to list the newly created raids.

```
[root@S7Archiver ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ ├─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:1 0 10G 0 lvm /var/log
│ ├─netwitness_vg00-usrhome 253:2 0 10G 0 lvm /home
│ ├─netwitness_vg00-root 253:3 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:4 0 4G 0 lvm [SWAP]
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
│ └─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 145.5T 0 disk
└─sdc1 8:33 0 72.8T 0 part
```

- Execute the `partNew` command with the following parameters to create partitions and mount points as required.

```
name=sdc service=archiver volume=archiver commit=1
```

Properties for /appliance  
 partNew Parameters: name=sdsc service=archiver volume=archiver commit=1 Send

Message Help

```

partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdc}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that v
 volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|f
logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create

```

/appliance?msg=partNew&force-content-type=text/plain&name=sdsc&service=archiver&volume=archiver&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f archiver0 /dev/sdc1
Volume group "archiver0" successfully created
/sbin/lvcreate -y -n database -l 100%FREE archiver0
Logical volume "database" created.
/sbin/mkfs.xfs /dev/archiver0/database
meta-data=/dev/archiver0/database isize=512 agcount=146, agsize=268435440 blks
 = sectsz=4096 attr=2, projid32bit=1
 = crc=1 finobt=1, sparse=1, rmapbt=0
 = reflink=1 bigtime=0 inobtcount=0
data = bsize=4096 blocks=39063387136, imaxpct=1
 = sunit=16 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/archiver/database0
/bin/mount /var/netwitness/archiver/database0

```

- Execute `devlist` command to verify partitions created.

Properties for /appliance  
 devlist Parameters:

Message Help

```

devlist: list storage devices
security.roles: appliance.manage

```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```

sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="145.52 TB" used=1
sda: vendor=DELL model=RAID size="2.18 TB" used=1

```

8. Execute the `srvAlloc` command with the following parameters to allocate the space to archiver service.

This adds storage to the archiver service configuration and restarts the service every time it is executed.

```
service=archiver volume=archiver0 commit=1
```

Properties for /appliance

`srvAlloc` ▾ Parameters: `service=archiver volume=archiver0 commit=1`

Message Help

```
srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|
 volume - <string, {enum-one:The value must be one of the following: archiver0|netwitness_vg00}>
 commit - <bool, optional> commit changes
```

`/appliance?msg=srvAlloc&force-content-type=text/plain&service=archiver&volume=archiver0&commit=1`

Output (or command manual help)

9. Execute the `srvList` command to verify the services allocated on the appliance.

Properties for /appliance

`srvList` ▾ Parameters:

Message Help

```
srvList: list of core services on this appliance
security.roles: appliance.manage
```

`/appliance?msg=srvList&force-content-type=text/plain`

Output (or command manual help)

```
localhost:56008: type=archiver
mounts.dir=/var/netwitness/archiver/database0
```

10. Execute `lsblk` command to list all the raids and partitions in the service.

```
[root@S7Archiver ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ ├─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:1 0 10G 0 lvm /var/log
│ ├─netwitness_vg00-usrhome 253:2 0 10G 0 lvm /home
│ ├─netwitness_vg00-root 253:3 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:4 0 4G 0 lvm [SWAP]
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
│ └─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 145.5T 0 disk
├─sdc1 8:33 0 145.5T 0 part
│ └─archiver0-database 253:5 0 145.5T 0 lvm /var/netwitness/archiver/database0
```

## Configure Storage for Decoder using MD2412

The following scenario configures storage on one, unencrypted, 12-Drive MD2412 PowerVault for a Network Decoder physical host..

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>

## b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance

raidList  Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 2
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

```
controller=1 enclosure=302 scheme=decoder preferSecure=0 commit=1
```

Properties for /appliance

raidNew  Parameters: controller=1 enclosure=302 scheme=decoder preferSecure=0 commit=1

Message Help

scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to be created  
 preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false}>  
 commit - <bool, optional> commit changes

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=302&scheme=decoder&preferSecure=0&commit=1

Output (or command manual help)

```
/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=302:0,302:1,302:2 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 1
Status = Success
Description = Add VD succeeded.

VD Information :
=====

VDID VDSIZE Status ErrType ErrCd Msg

 3 29.104 TiB Success - - -

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=302:3,302:4,302:5,302:6,302:7,302:8,302:9,302:10,302:11 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
```

- Execute the `raidList` command to verify the new RAID array.

You should now see the following information.

In Use: TRUE

Devices: <devices> (for example, sdc, sdd)

Properties for /appliance  
 raidList ▾ Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:2:0

Controller 1 at PCI Address 8b:00.0, Enclosure 302, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:3:0
 sdd /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:4:0
```

4. Execute the `devlist` command to find the exact memory distribution for the newly created raids..

Properties for /appliance  
 devlist ▾ Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdd: vendor=DELL model=RAID size="116.42 TB" used=0
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="29.1 TB" used=0
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

5. Execute `lsblk` command at the backend to verify block devices created on the host.

```
[root@S7Decoder ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ ├─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:1 0 10G 0 lvm /var/log
│ ├─netwitness_vg00-usrhome 253:2 0 10G 0 lvm /home
│ ├─netwitness_vg00-root 253:3 0 30G 0 lvm /
│ └─netwitness_vg00-swap 253:4 0 4G 0 lvm [SWAP]
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
│ └─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 29.1T 0 disk
sdd 8:48 0 116.4T 0 disk
```

- Execute the `partNew` command to create the `decodersmall` partition first (decoder dir, index, metadb, sessiondb) on the smaller block device (`sd`) with the following parameters.

```
name=sd service=decoder volume=decodersmall commit=1
```

Properties for /appliance

Parameters:

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdd|sd}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decc
 volume - <string, optional, {enum-one:The value must be one of the following: concentrator|inde
 logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> vc
```

/appliance?msg=partNew&force-content-type=text/plain&name=sd&service=decoder&volume=decodersmall&commit=

Output (or command manual help)

```
/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f decodersmall /dev/sdc1
Volume group "decodersmall" successfully created
/sbin/lvcreate -y -n decoroot -L 20G decodersmall
Wiping xfs signature on /dev/decodersmall/decoroot.
Logical volume "decoroot" created.
/sbin/mkfs.xfs /dev/decodersmall/decoroot
meta-data=/dev/decodersmall/decoroot isize=512 agcount=16, agsize=327664 blks
 = sectsz=4096 attr=2, projid32bit=1
 = crc=1 finobt=1, sparse=1, rmapbt=0
 = reflink=1 bigtime=0 inobtcount=0
data = bsize=4096 blocks=5242624, imaxpct=25
 = sunit=16 swidth=32 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=2560, version=2
 = sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder
/bin/mount /var/netwitness/decoder
```

- Execute the `partNew` command to create the decoder volume (`packetdb`) on the bigger block device of the two created (`sdd`) with the following parameters.

```
name=sdd service=decoder volume=decoder commit=1
```

Properties for /appliance

`partNew` Parameters: `name=sdd service=decoder volume=decoder commit=1`

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdd|sdc}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|
 volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|
logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volu
```

`/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=decoder&volume=decoder&commit=1`

Output (or command manual help)

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f decoder /dev/sdd1
Volume group "decoder" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder
Wiping xfs signature on /dev/decoder/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder/packetdb
meta-data=/dev/decoder/packetdb isize=512 agcount=117, agsize=268435440 blks
 = sectsz=4096 attr=2, projid32bit=1
 = crc=1 finobt=1, sparse=1, rmapbt=0
 = reflink=1 bigtime=0 inobtcount=0
data = bsize=4096 blocks=31250709504, imaxpct=1
 = sunit=16 swidth=128 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=521728, version=2
 = sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb
/bin/mount /var/netwitness/decoder/packetdb
```

- Execute `devlist` command to verify partitions created.

Properties for /appliance  
 devlist ▾ Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdd: vendor=DELL model=RAID size="116.42 TB" used=1
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="29.1 TB" used=1
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

9. Execute the `srvAlloc` command with the following parameters to add the storage information into the Service Configuration settings. Service is restarted each time the command is executed.

```
service=decoder volume=decodersmall commit=1
```

Properties for /appliance  
 srvAlloc ▾ Parameters:

Message Help

```
srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|deco
 volume - <string, {enum-one:The value must be one of the following: decoder|decodersmall|netwit
 commit - <bool, optional> commit changes
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=decoder&volume=decodersmall&commit=1

Output (or command manual help)

```
Set /database/config/meta.dir to /var/netwitness/decoder/metadb==27.04 TB
Set /database/config/session.dir to /var/netwitness/decoder/sessiondb==569.72 GB
Set /index/config/index.dir to /var/netwitness/decoder/index==28.49 GB
```

```
service=decoder volume=decoder commit=1
```

Properties for /appliance

 Parameters: 

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|dec
 volume - <string, {enum-one:The value must be one of the following: decoder|decodersmall|netw
 commit - <bool, optional> commit changes

```

Output (or command manual help)

```
Set /database/config/packet.dir to /var/netwitness/decoder/packetdb==110.6 TB
```

10. Execute the `srvList` command to verify the services allocated on the appliance.

Properties for /appliance

 Parameters: 

Message Help

```

srvList: list of core services on this appliance
security.roles: appliance.manage

```

Output (or command manual help)

```

localhost:56004: type=decoder
packet.dir="/var/netwitness/decoder/packetdb==110.6 TB"
meta.dir="/var/netwitness/decoder/metadb==27.04 TB"
session.dir="/var/netwitness/decoder/sessiondb==569.72 GB"
index.dir="/var/netwitness/decoder/index==28.49 GB"

```

11. Execute `lsblk` command to list all the raids and partitions in the service.

```
[root@S7Decoder ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
└─sda3 8:3 0 2.2T 0 part
 ├─netwitness_vg00-nwhome 253:5 0 4.3T 0 lvm /var/netwitness
 ├─netwitness_vg00-varlog 253:6 0 10G 0 lvm /var/log
 ├─netwitness_vg00-usrhome 253:7 0 10G 0 lvm /home
 ├─netwitness_vg00-root 253:8 0 30G 0 lvm /
 └─netwitness_vg00-swap 253:9 0 4G 0 lvm [SWAP]
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
└─netwitness_vg00-nwhome 253:5 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 29.1T 0 disk
├─sdc1 8:33 0 29.1T 0 part
│ ├─decodersmall-decoroot 253:1 0 20G 0 lvm /var/netwitness/decoder
│ ├─decodersmall-index 253:2 0 30G 0 lvm /var/netwitness/decoder/index
│ ├─decodersmall-sessiondb 253:3 0 600G 0 lvm /var/netwitness/decoder/sessiondb
│ └─decodersmall-metadb 253:4 0 28.5T 0 lvm /var/netwitness/decoder/metadb
└─sdd 8:48 0 116.4T 0 disk
 ├─sdd1 8:49 0 116.4T 0 part
 └─decoder-packetdb 253:0 0 116.4T 0 lvm /var/netwitness/decoder/packetdb
```

## Configure Storage for Concentrator using MD2412

The following scenarios configure storage on one or more unencrypted, 12-Drive MD2412 PowerVault for a Concentrator physical host.

**Note:** All 12 drives on PowerVault MD2412 are HDDs. This is also used as storage for concentrator service whereas older PowerVault MD1400 for concentrator storage has mixed drive configuration with HDDs and SSDs. Since MD2412 has no support for SSDs and the concentrator service requires atleast 2 SSDs for index database, SSDs are to be installed on the concentrator physical host in slots 4 through 9 with a minimum of 3 drives.

Supported SSD configurations on the S7 concentrator host -

- 3 drives - slots 4,5 and 6 - RAID 5
- 6 drives - slots 4,5,6,7,8 and 9 - RAID 6
- 3 Drives + 3 Drives - RAID 5, initially 3 drives are installed followed by extension with another 3 drives.

The SSD drive packs are available in two sizes - 3 \* 3.2 TB or 3 \* 7.68 TB. In the following examples, the drive pack of 3 \* 3.2 TB is used. When a larger drive pack (3 \* 7.68 TB) is used, the virtual drives are larger in size.

### Scenario 1:

3 SSD drives are added to the Concentrator physical host. Host is connected to one unencrypted 12-Drive MD2412 PowerVault.

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>

## b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance  
 raidList Parameters:

Message Help

raidList: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:2:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 2
Vendor: DELL
Model: DSES Enclosure
In Use: false
Drives: 14.552 TB HDD x 12
Devices:

```

SSD configuration on this host is as seen below. 3 SSDs are added to slots 4,5 and 6 on the concentrator host for configuring index.

To view details about the SSDs installed on controller 0, run the following command on the concentrator:

```
/opt/MegaRAID/perccli2/perccli2 /c0/eall/sall show
```

Partial output of the above command is as shown below.

```

CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-513.24.1.el8_9.x86_64
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.

Drive Information :
=====

```

| EID:Slr | PID | State | Status | DG | Size      | Intf | Med | SED_Type   | SeSz | Model        | Sp | LU/NS | Count | Alt-EID |
|---------|-----|-------|--------|----|-----------|------|-----|------------|------|--------------|----|-------|-------|---------|
| 286:0   | 275 | Conf  | Online | 0  | 2.182 TiB | SAS  | HDD | Enterprise | 512B | CL2400MM0149 | U  |       | 1     | -       |
| 286:1   | 276 | Conf  | Online | 0  | 2.182 TiB | SAS  | HDD | Enterprise | 512B | CL2400MM0149 | U  |       | 1     | -       |
| 286:2   | 277 | Conf  | Online | 1  | 2.182 TiB | SAS  | HDD | Enterprise | 512B | CL2400MM0149 | U  |       | 1     | -       |
| 286:3   | 278 | Conf  | Online | 1  | 2.182 TiB | SAS  | HDD | Enterprise | 512B | CL2400MM0149 | U  |       | 1     | -       |
| 286:4   | 279 | UConf | Good   | -  | 2.910 TiB | SAS  | SSD | Enterprise | 512B | KPM7WVUG3T20 | U  |       | 1     | -       |
| 286:5   | 280 | UConf | Good   | -  | 2.910 TiB | SAS  | SSD | Enterprise | 512B | KPM7WVUG3T20 | U  |       | 1     | -       |
| 286:6   | 281 | UConf | Good   | -  | 2.910 TiB | SAS  | SSD | Enterprise | 512B | KPM7WVUG3T20 | U  |       | 1     | -       |

- Execute the `raidNew` command with `scheme=concentrator` for controller (1) with HDDs and `scheme=concentrator-metakit` for controller (0) with SSDs.

To configure RAID for concentrator scheme, execute `raidNew` with the following parameters using the controller number and the enclosure number you just recorded from `raidList`.

controller=1 enclosure=386 scheme=concentrator preferSecure=0 commit=1

Properties for /appliance

raidNew ▼ Parameters: controller=1 enclosure=386 scheme=concentrator preferSecure=0 commit=1 Send

Message Help

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386}> Enclosure number of the shelf to clear. Required if scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-index|metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate  
preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off}> Prefer commit - <bool, optional> commit changes

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=386&scheme=concentrator&preferSecure=0&commit=1

Output (or command manual help)

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r6 drives=386:0,386:1,386:2,386:3,386:4,386:5,386:6,386:7,386:8,386:9,386:10,386:11 nora Strip=64  
CLI Version = 008.0004.0000.0022 Apr 28, 2023  
Operating system = Linux4.18.0-513.24.1.el8\_9.x86\_64  
Controller = 1  
Status = Success  
Description = Add VD succeeded.  
  
VD Information :  
\*\*\*\*\*  

| VDID | VDSize      | Status  | ErrType | ErrCd | Msg |
|------|-------------|---------|---------|-------|-----|
| 1    | 145.522 TiB | Success | -       | -     | -   |

3. Execute the raidList command to verify the concentrator RAID created.

Properties for /appliance

raidList ▼

Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

`/appliance?msg=raidList&force-content-type=text/plain`

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:2:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0
```

- Execute the `raidNew` command with the following parameters using the controller number and the enclosure number associated with SSDs added to the concentrator host. This creates the RAID for index.

```
controller=0 enclosure=286 scheme=concentrator-metakit preferSecure=0
commit=1
```

Properties for /appliance

raidNew Parameters: controller=0 enclosure=286 scheme=concentrator-metakit preferSecure=0 commit=1

Message Help

```
enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386}> Enclosure number of the shelf
scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybr
expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volume
preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,f
drives and a controller with a security key set
commit - <bool, optional> commit changes
```

/appliance?msg=raidNew&force-content-type=text/plain&controller=0&enclosure=286&scheme=concentrator-metakit&preferSecure=0&commit=1

Output (or command manual help)

```
/opt/MegaRAID/percccli2/percccli2 /c0 add vd r5 drives=286:4,286:5,286:6 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-513.24.1.el8_9.x86_64
Controller = 0
Status = Success
Description = Add VD succeeded.

VD Information :
=====

VDID VDSIZE Status ErrType ErrCd Msg

1 5.820 TiB Success - - -

```

- Execute the `raidList` command to verify the index RAID created.

Properties for /appliance

raidList ▾ Parameters: 

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
 Devices: sdd /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:2:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:4:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0
```

6. Execute `devlist` to verify block devices created.

Properties for /appliance

devlist ▾

Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&amp;force-content-type=text/plain

Output (or command manual help)

```
sdd: vendor=DELL model=RAID size="5.82 TB" used=0
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="145.52 TB" used=0
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

The same can be verified using `lsblk`.

```
[root@S7Concentrator ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
└─sda3 8:3 0 2.2T 0 part
 ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
 ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
 ├─netwitness_vg00-nwhome 253:7 0 4.3T 0 lvm /var/netwitness
 ├─netwitness_vg00-varlog 253:14 0 10G 0 lvm /var/log
 └─netwitness_vg00-usrhome 253:17 0 10G 0 lvm /home
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
└─netwitness_vg00-nwhome 253:7 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 145.5T 0 disk
└─sdc1 8:33 0 29.1T 0 part
sdd 8:48 0 5.8T 0 disk
```

- Execute the `partNew` command to create the concentrator partition first with the following parameters. You must create the concentrator volume before index volume or it will fail. `sdc` is associated with HDDs on PowerVault 2412 and is to be used for concentrator volume.

```
name=sdc service=concentrator volume=concentrator commit=1
```

Properties for /appliance

partNew Parameters:

Message Help

```

partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
 name - <string, {enum-one:The value must be one of the following: sdd|sdc}> block device name
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}>
 volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create

```

/appliance?msg=partNew&force-content-type=text/plain&name=sdc&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f concentrator /dev/sdc1
Volume group "concentrator" successfully created
/sbin/lvcreate -y -n root -L 30G concentrator
Logical volume "root" created.
/sbin/mkfs.xfs /dev/concentrator/root
meta-data=/dev/concentrator/root isize=512 agcount=16, agsize=491504 blks
 = sectsz=4096 attr=2, projid32bit=1
 = crc=1 finobt=1, sparse=1, rmapbt=0
 = reflink=1 bigtime=0 inobtcount=0
data = bsize=4096 blocks=7864064, imaxpct=25
 = sunit=16 swidth=160 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=3839, version=2
 = sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator
/bin/mount /var/netwitness/concentrator
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator

```

8. Execute the `partNew` command with the following parameters to create an index on SSDs. `sdd` is associated with the SSDs added and is to be used for index volume.

```
name=sdd service=concentrator volume=index commit=1
```

Properties for /appliance

partNew ▾ Parameters:

Message Help

partNew: make partitions on a block device  
 security.roles: appliance.manage  
 parameters:  
   name - <string, {enum-one:The value must be one of the following: sdd|sdc}> block device name  
   service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create

Output (or command manual help)

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f index /dev/sdd1
Volume group "index" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index isize=512 agcount=32, agsize=48828368 blks
= sectsz=4096 attr=2, projid32bit=1
= crc=1 finobt=1, sparse=1, rmapbt=0
= reflink=1 bigtime=0 inobtcount=0
data = bsize=4096 blocks=1562507776, imaxpct=5
= sunit=16 swidth=32 blks
naming =version 2 bsize=4096 ascii-ci=0, ftype=1
log =internal log bsize=4096 blocks=521728, version=2
= sectsz=4096 sunit=1 blks, lazy-count=1
realtime =none extsz=4096 blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index
/bin/mount /var/netwitness/concentrator/index
```

- Run `devlist` to verify block devices have been configured.

## Properties for /appliance

devlist ▾

Parameters:

## Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&amp;force-content-type=text/plain

## Output (or command manual help)

```
sdd: vendor=DELL model=RAID size="5.82 TB" used=1
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="145.52 TB" used=1
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

10. Execute the `srvAlloc` command with the following parameters.

```
service=concentrator volume=concentrator commit=1
```

## Properties for /appliance

srvAlloc ▾

Parameters: service=concentrator volume=concentrator commit=1

## Message Help

```
srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|
 volume - <string, {enum-one:The value must be one of the following: concentrator|index|netw
 commit - <bool, optional> commit changes
```

/appliance?msg=srvAlloc&amp;force-content-type=text/plain&amp;service=concentrator&amp;volume=concentrator&amp;commit=1

## Output (or command manual help)

```
Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==124.39 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==13.82 TB
```

```
service=concentrator volume=index commit=1
```

Properties for /appliance

 Parameters: 

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|index|netwitness}
 volume - <string, {enum-one:The value must be one of the following: concentrator|index|netwitness}
 commit - <bool, optional> commit changes

```

Output (or command manual help)

```
Set /index/config/index.dir to /var/netwitness/concentrator/index==5.53 TB
```

11. Execute `srvList` to verify storage has been assigned as required.

Properties for /appliance

 Parameters: 

Message Help

```

srvList: list of core services on this appliance
security.roles: appliance.manage

```

Output (or command manual help)

```

localhost:56005: type=concentrator
meta.dir="/var/netwitness/concentrator/metadb==124.39 TB"
session.dir="/var/netwitness/concentrator/sessiondb==13.82 TB"
index.dir="/var/netwitness/concentrator/index==5.53 TB"

```

12. Verify partitions and mounts on `lsblk`.

```

[root@S7Concentrator ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 2.2T 0 part
│ ├─netwitness_vg00-root 253:0 0 30G 0 lvm /
│ ├─netwitness_vg00-swap 253:1 0 4G 0 lvm [SWAP]
│ ├─netwitness_vg00-nwhome 253:2 0 4.3T 0 lvm /var/netwitness
│ ├─netwitness_vg00-varlog 253:3 0 10G 0 lvm /var/log
│ └─netwitness_vg00-usrhome 253:4 0 10G 0 lvm /home
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
└─netwitness_vg00-nwhome 253:2 0 4.3T 0 lvm /var/netwitness
sdc 8:32 0 145.5T 0 disk
├─sdc1 8:33 0 145.5T 0 part
│ ├─concentrator-root 253:5 0 30G 0 lvm /var/netwitness/concentrator
│ ├─concentrator-sessiondb 253:6 0 14.6T 0 lvm /var/netwitness/concentrator/sessiondb
│ └─concentrator-metadb 253:7 0 131T 0 lvm /var/netwitness/concentrator/metadb
sdd 8:48 0 5.8T 0 disk
├─sdd1 8:49 0 5.8T 0 part
└─index-index 253:8 0 5.8T 0 lvm /var/netwitness/concentrator/index

```

## Scenario 2:

3 SSD Drives + 3 SSD Drives are installed on the Concentrator physical host. This means that initially 3 SSD drives are installed, followed by extension with another 3 drives.

Host is connected to two unencrypted 12-Drive MD2412 PowerVaults.

The following scenario configures storage as per the scenario described above.

1. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>
  - b. Verify the Drive Count, Size, and Vendor.  
The following example illustrates what you should see before you create a RAID array.

Properties for /appliance

raidList ▾ Parameters:

Message Help

```
raidlist: list drive shelves attached to this appliance
security.roles: appliance.manage
```

</appliance?msg=raidList&force-content-type=text/plain>

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 2
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 2
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:
```

SSD configuration on this host is as seen below. 3 SSDs from the first drive pack are added to slots 4,5 and 6 on the concentrator host for configuring index.

To view details about the SSDs installed on controller 0, run the following command on the concentrator:

```
/opt/MegaRAID/perccli2/perccli2 /c0/eall/sall show
```

Partial output of the above command is as shown below.

```

CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.

Drive Information :
=====

EID:SlT PID State Status DG Size Intf Med SED_Type SeSz Model Sp LU/NS Count Alt-EID

286:0 275 Conf Online 0 2.182 TiB SAS HDD Enterprise 512B CL2400MM0149 U 1 -
286:1 276 Conf Online 0 2.182 TiB SAS HDD Enterprise 512B CL2400MM0149 U 1 -
286:2 277 Conf Online 1 2.182 TiB SAS HDD Enterprise 512B CL2400MM0149 U 1 -
286:3 278 Conf Online 1 2.182 TiB SAS HDD Enterprise 512B CL2400MM0149 U 1 -
286:4 279 UConf Good - 2.910 TiB SAS SSD Enterprise 512B KPM7WVUG3T20 U 1 -
286:5 280 UConf Good - 2.910 TiB SAS SSD Enterprise 512B KPM7WVUG3T20 U 1 -
286:6 281 UConf Good - 2.910 TiB SAS SSD Enterprise 512B KPM7WVUG3T20 U 1 -

```

2. Execute the `raidNew` command with `scheme=concentrator` for controller (1) with HDDs and `scheme=concentrator-metakit` for controller (0) with SSDs.

To configure RAID for concentrator scheme, execute `raidNew` with the following parameters using the controller number and the enclosure number you just recorded from `raidList`.

`controller=1 enclosure=386 scheme=concentrator preferSecure=0 commit=1`

This creates concentrator RAID using the HDDs in PowerVault MD2412.

Properties for /appliance

raidNew Parameters:

Message Help

raidNew: allocate RAID devices in a drive shelf

security.roles: appliance.manage

parameters:

- controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
- enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386|400}> Enclosure number of the shelf to clear. Required if attached.
- scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-ir-expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=386&scheme=concentrator&preferSecure=0&commit=1

Output (or command manual help)

```

/opt/MegaRAID/percc112/percc112 /c1 add vd r6 drives=386:0,386:1,386:2,386:3,386:4,386:5,386:6,386:7,386:8,386:9,386:10,386:11 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 1
Status = Success
Description = Add VD succeeded.

VD Information :
=====

VDID VDSIZE Status ErrType ErrCd Msg

1 145.522 TiB Success - - -

```

3. Execute the `raidList` command to verify the concentrator RAID created.

Properties for /appliance

raidList ▼ Parameters:

Message Help

raidList: list drive shelves attached to this appliance  
security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
 Devices: sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:

```

- To configure RAID for index scheme, execute `raidNew` with the following parameters using the controller number and the enclosure number associated with SSDs added to the concentrator host. This creates the RAID for index.

```

controller=0 enclosure=286 scheme=concentrator-metakit preferSecure=0
commit=1

```

Properties for /appliance

raidNew ▾ Parameters:

Message Help

```

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386|400}> Enclosure number of the
attached.
scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hy
expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volu
preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true
drives and a controller with a security key set
commit - <bool, optional> commit changes

```

/appliance?msg=raidNew&force-content-type=text/plain&controller=0&enclosure=286&scheme=concentrator-metakit&preferSecure=0&commit=1

Output (or command manual help)

```

/opt/MegaRAID/perccli2/perccli2 /c0 add vd r5 drives=286:4,286:5,286:6 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 0
Status = Success
Description = Add VD succeeded.

VD Information :
=====

VDID VDSize Status ErrType ErrCd Msg

 1 5.820 TiB Success - -

```

5. Execute `raidList` command to verify the index RAID created.

Properties for /appliance

raidList  Parameters:

Message Help

raidList: list drive shelves attached to this appliance  
security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 3
 Devices: sdd /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:

```

6. Add the second drive pack with another 3 SSDs.

SSD configuration on this host is as seen below. 3 SSDs from the second drive pack are added to slots 7,8 and 9 on the concentrator host for configuring index.

To view details about the SSDs installed on controller 0, run the following command on the concentrator:

```
/opt/MegaRAID/percccli2/percccli2 /c0/eall/sall show
```

Partial output of the above command is as shown below.

```

CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.

```

Drive Information :

```
=====
```

| EID:Slr | PID | State | Status | DG | Size  | Intf | Med | SED_Type | SeSz       | Model | Sp           | LU/NS | Count | Alt-EID |
|---------|-----|-------|--------|----|-------|------|-----|----------|------------|-------|--------------|-------|-------|---------|
| -----   |     |       |        |    |       |      |     |          |            |       |              |       |       |         |
| 286:0   | 275 | Conf  | Online | 1  | 2.182 | TiB  | SAS | HDD      | Enterprise | 512B  | CL2400MM0149 | U     | 1     | -       |
| 286:1   | 276 | Conf  | Online | 1  | 2.182 | TiB  | SAS | HDD      | Enterprise | 512B  | CL2400MM0149 | U     | 1     | -       |
| 286:2   | 277 | Conf  | Online | 2  | 2.182 | TiB  | SAS | HDD      | Enterprise | 512B  | CL2400MM0149 | U     | 1     | -       |
| 286:3   | 278 | Conf  | Online | 2  | 2.182 | TiB  | SAS | HDD      | Enterprise | 512B  | CL2400MM0149 | U     | 1     | -       |
| 286:4   | 279 | Conf  | Online | 0  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| 286:5   | 280 | Conf  | Online | 0  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| 286:6   | 281 | Conf  | Online | 0  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| 286:7   | 282 | UConf | Good   | -  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| 286:8   | 283 | UConf | Good   | -  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| 286:9   | 284 | UConf | Good   | -  | 2.910 | TiB  | SAS | SSD      | Enterprise | 512B  | KPM7WVUG3T20 | U     | 1     | -       |
| -----   |     |       |        |    |       |      |     |          |            |       |              |       |       |         |

7. Configure RAID on the second PowerVault. Execute the `raidList` command.
  - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.  
You should see the following information.  
In Use: FALSE  
Devices: <empty>
  - b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance

raidList  Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 6
 Devices: sdd /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: false
 Drives: 14.552 TB HDD x 12
 Devices:
```

8. Execute the `raidNew` command with `scheme=concentrator` for controller (1) with HDDs and `scheme=concentrator-metakit` for controller (0) with SSDs.

To configure RAID for concentrator scheme, execute `raidNew` with the following parameters using the controller number and the enclosure number you just recorded from `raidList`.

This creates concentrator RAID using the HDDs in PowerVault MD2412.

```
controller=1 enclosure=400 scheme=concentrator preferSecure=0 commit=1
```

Properties for /appliance  
raidNew Parameters: controller=1 enclosure=400 scheme=concentrator preferSecure=0 commit=1

Message Help  

```

security:roles: appliance.manage
parameters:
 controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
 enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386|400}> Enclosure number of the shelf to clear. Required if
 attached.
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexec
 expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate
 preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true,false,on,off}> Prefer crea

```

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=400&scheme=concentrator&preferSecure=0&commit=1

Output (or command manual help)

```

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r6 drives=400:0,400:1,400:2,400:3,400:4,400:5,400:6,400:7,400:8,400:9,400:10,400:11 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 1
Status = Success
Description = Add VD succeeded.

VD Information :
=====
VDID VDSIZE Status ErrType ErrCd Msg

 2 145.522 TiB Success - - -
=====

```

9. Execute the `raidList` command to verify the concentrator RAID created.

Properties for /appliance

raidList ▾ Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&amp;force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 6
 Devices: sdd /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sde /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:2:0
```

10. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number associated with SSDs added to the concentrator host. This creates the RAID for index.

```
controller=0 enclosure=286 scheme=concentrator-metakit preferSecure=0
commit=1
```

Properties for /appliance

raidNew  Parameters:

Message Help

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 286|386|400}> Enclosure number of the attached.  
 scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-expansion|decoder-hotspare|logdecoder-hotspare|decoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volume  
 preferSecure - <bool, optional, {bool:The value must be one of the following acceptable boolean values: 0,1,yes,no,true}> drives and a controller with a security key set  
 commit - <bool, optional> commit changes

/appliance?msg=raidNew&force-content-type=text/plain&controller=0&enclosure=286&scheme=concentrator-metakit&preferSecure=0&commit=1

Output (or command manual help)

```
/opt/MegaRAID/percccli2/percccli2 /c0 add vd r5 drives=286:7,286:8,286:9 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-553.8.1.el8_10.x86_64
Controller = 0
Status = Success
Description = Add VD succeeded.
```

VD Information :

=====

| VDID | VDSIZE    | Status  | ErrType | ErrCd | Msg |
|------|-----------|---------|---------|-------|-----|
| 2    | 5.820 TiB | Success | -       | -     | -   |

11. Execute the `raidList` command to verify the index RAID created.

Properties for /appliance

raidList ▾

Parameters:

Message Help

raidList: list drive shelves attached to this appliance  
 security.roles: appliance.manage

/appliance?msg=raidList&amp;force-content-type=text/plain

Output (or command manual help)

```
Controller 0 at PCI Address 4a:00.0, Enclosure 286, SCSI Channel 0
 Vendor: DP
 Model: BP_PSV
 In Use: true
 Drives: 2.182 TB HDD x 4
 2.91 TB SSD x 6
 Devices: sdd /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:1:0
 sdf /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:2:0
 sda /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:3:0
 sdb /dev/disk/by-path/pci-0000:4a:00.0-scsi-0:0:5:0

Controller 1 at PCI Address 8b:00.0, Enclosure 386, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
 Devices: sdc /dev/disk/by-path/pci-0000:8b:00.0-scsi-0:0:1:0

Controller 1 at PCI Address 8b:00.0, Enclosure 400, SCSI Channel 0
 Vendor: DELL
 Model: DSES Enclosure
 In Use: true
 Drives: 14.552 TB HDD x 12
```

12. Execute `devlist` to verify block devices created.

Properties for /appliance

devlist  Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="5.82 TB" used=0
sdd: vendor=DELL model=RAID size="5.82 TB" used=0
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sde: vendor=DELL model=RAID size="145.52 TB" used=0
sdc: vendor=DELL model=RAID size="145.52 TB" used=0
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

13. The same can be verified using `lsblk`.

```
[root@S7Concentrator ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 2.2T 0 disk
├─sda1 8:1 0 256M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
└─sda3 8:3 0 2.2T 0 part
 ├─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
 ├─netwitness_vg00-varlog 253:1 0 10G 0 lvm /var/log
 ├─netwitness_vg00-usrhome 253:2 0 10G 0 lvm /home
 ├─netwitness_vg00-root 253:3 0 30G 0 lvm /
 └─netwitness_vg00-swap 253:4 0 4G 0 lvm [SWAP]
sdb 8:16 0 2.2T 0 disk
├─sdb1 8:17 0 2.2T 0 part
└─┬─netwitness_vg00-nwhome 253:0 0 4.3T 0 lvm /var/netwitness
 └─
sdc 8:32 0 145.5T 0 disk
sdd 8:48 0 5.8T 0 disk
sde 8:64 0 145.5T 0 disk
sdf 8:80 0 5.8T 0 disk
```

14. Execute the `partNew` command to create the concentrator partition first with the following parameters. You must create the concentrator volume before index volume or it will fail.
- sdc is associated with HDDs on PowerVault MD2412 and is to be used for concentrator volume.
- ```
name=sdc service=concentrator volume=concentrator commit=1
```

Properties for /appliance

partNew Parameters: name=sdd service=concentrator volume=concentrator commit=1

Send

Message Help

```

partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdf|sdd|sde|sdc}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create

```

/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdc mklabel gpt
/sbin/parted -s -a optimal /dev/sdc mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdc1
Physical volume "/dev/sdc1" successfully created.
/sbin/vgcreate -f concentrator /dev/sdc1
Volume group "concentrator" successfully created
/sbin/lvcreate -y -n root -L 30G concentrator
Wiping xfs signature on /dev/concentrator/root.
Logical volume "root" created.
/sbin/mkfs.xfs /dev/concentrator/root
meta-data=/dev/concentrator/root isize=512    agcount=16, agsize=491504 blks
          =                       sectsz=4096   attr=2, projid32bit=1
          =                       crc=1         finobt=1, sparse=1, rmapbt=0
          =                       reflink=1      bigtime=0 inobtcount=0
data      =                       bsize=4096   blocks=7864064, imaxpct=25
          =                       sunit=16      swidth=160 blks
naming    =version 2              bsize=4096   ascii-ci=0, ftype=1
log       =internal log          bsize=4096   blocks=3839, version=2
          =                       sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none                  extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator
/bin/mount /var/netwitness/concentrator
/sbin/lvcreate -v -n sessiondb -l 10%FREE concentrator

```

- Execute the `partNew` command with the following parameters to create an index on SSDs. `sdd` is associated with the SSDs added and is to be used for index volume.

```
name=sdd service=concentrator volume=index commit=1
```

Properties for /appliance

partNew Parameters: name=sdd service=concentrator volume=index commit=1

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdf|sdd|sde}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}>
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volu
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdd&service=concentrator&volume=index&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdd mklabel gpt
/sbin/parted -s -a optimal /dev/sdd mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdd1
Physical volume "/dev/sdd1" successfully created.
/sbin/vgcreate -f index /dev/sdd1
Volume group "index" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index
Wiping xfs signature on /dev/index/index.
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index/index
meta-data=/dev/index/index      isize=512    agcount=32, agsize=48828368 blks
=                               sectsz=4096   attr=2, projid32bit=1
=                               crc=1        finobt=1, sparse=1, rmapbt=0
=                               reflink=1     bigtime=0 inobtcount=0
data      =                     bsize=4096   blocks=1562507776, imaxpct=5
=                               sunit=16      swidth=32 blks
naming    =version 2           bsize=4096   ascii-ci=0, ftype=1
log        =internal log      bsize=4096   blocks=521728, version=2
=                               sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none               extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index
/bin/mount /var/netwitness/concentrator/index
```

16. Run `devlist` to verify block devices have been configured.

Properties for /appliance

devlist Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="5.82 TB" used=0
sdd: vendor=DELL model=RAID size="5.82 TB" used=1
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sde: vendor=DELL model=RAID size="145.52 TB" used=0
sdc: vendor=DELL model=RAID size="145.52 TB" used=1
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

17. Verify partitions and mounts on lsblk.

```
[root@S7Concentrator ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda                                8:0    0   2.2T  0 disk
├─sda1                            8:1    0   256M  0 part  /boot/efi
├─sda2                            8:2    0    1G    0 part  /boot
├─sda3                            8:3    0   2.2T  0 part
│   ├─netwitness_vg00-nwhome      253:0    0   4.3T  0 lvm    /var/netwitness
│   ├─netwitness_vg00-varlog      253:1    0   10G    0 lvm    /var/log
│   ├─netwitness_vg00-usrhome     253:2    0   10G    0 lvm    /home
│   ├─netwitness_vg00-root        253:3    0    30G    0 lvm    /
│   └─netwitness_vg00-swap        253:4    0    4G    0 lvm    [SWAP]
sdb                                8:16    0   2.2T  0 disk
├─sdb1                            8:17    0   2.2T  0 part
│   └─netwitness_vg00-nwhome      253:0    0   4.3T  0 lvm    /var/netwitness
sdc                                8:32    0  145.5T  0 disk
├─sdc1                            8:33    0  145.5T  0 part
│   ├─concentrator-root          253:5    0    30G    0 lvm    /var/netwitness/concentrator
│   ├─concentrator-sessiondb     253:6    0   14.6T  0 lvm    /var/netwitness/concentrator/sessiondb
│   └─concentrator-metadb        253:7    0   131T   0 lvm    /var/netwitness/concentrator/metadb
sdd                                8:48    0   5.8T  0 disk
├─sdd1                            8:49    0   5.8T  0 part
│   └─index-index               253:8    0   5.8T  0 lvm    /var/netwitness/concentrator/index
sde                                8:64    0  145.5T  0 disk
sdf                                8:80    0   5.8T  0 disk
```

18. Execute the `partNew` command to create the concentrator partition on the second PowerVault with the following parameters. You must create the concentrator volume before index volume or it will fail.

`sde` is associated with HDDs on PowerVault MD2412 and is to be used for concentrator volume.

`name=sde service=concentrator volume=concentrator commit=1`

Properties for /appliance

partNew Parameters:

Message Help

```
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdf|sdd|sde|sdc}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> se
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|deco
  logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create
  commit - <bool, optional> commit changes
  start - <string, optional> The start of the partition. Can be in sectors (number ends with 's'), or percentage (nu
```

/appliance?msg=partNew&force-content-type=text/plain&name=sde&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sde mklabel gpt
/sbin/parted -s -a optimal /dev/sde mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sde1
Physical volume "/dev/sde1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sde1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512    agcount=32, agsize=122073056 blks
/bin/mkdir -p /var/netwitness/concentrator/sessiondb0
/bin/mount /var/netwitness/concentrator/sessiondb0
/sbin/lvcreate -y -n metadb -l 100%FREE concentrator0
Logical volume "metadb" created.
/sbin/mkfs.xfs /dev/concentrator0/metadb
meta-data=/dev/concentrator0/metadb isize=512    agcount=131, agsize=268435440 blks
/bin/mkdir -p /var/netwitness/concentrator/metadb0
/bin/mount /var/netwitness/concentrator/metadb0
```

19. Execute the `partNew` command to create the index partition on SSDs associated with the second SSD drive-pack. `sdf` is associated with these SSDs and is to be used for index volume.

`name=sdf service=concentrator volume=index commit=1`

Properties for /appliance

partNew Parameters:

Message Help

parameters:

- name - <string, {enum-one:The value must be one of the following: sdf|sdd|sde}> block device name
- service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service
- volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|logdecoder-meta|archiver|hybrid-concentrator|endpoint-log-hybrid|log-indexed-decoder|logindex}> volume to create
- commit - <bool, optional> commit changes
- start - <string, optional> The start of the partition. Can be in sectors (number ends with 's'), or percentage (number ends with '%')
- end - <string, optional> The end of the partition. Can be in sectors (number ends with 's'), or percentage (number ends with '%')

Output (or command manual help)

```
/sbin/parted -s /dev/sdf mklabel gpt
/sbin/parted -s -a optimal /dev/sdf mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdf1
Physical volume "/dev/sdf1" successfully created.
/sbin/vgcreate -f index0 /dev/sdf1
Volume group "index0" successfully created
/sbin/lvcreate -y -n index -l 100%FREE index0
Wiping xfs signature on /dev/index0/index.
Logical volume "index" created.
/sbin/mkfs.xfs /dev/index0/index
meta-data=/dev/index0/index      isize=512    agcount=32, agsize=48828368 blks
=                               sectsz=4096   attr=2, projid32bit=1
=                               crc=1         finobt=1, sparse=1, rmapbt=0
=                               reflink=1      bigtime=0 inobtcount=0
data      =                       bsize=4096   blocks=1562507776, imaxpct=5
=                               sunit=16        swidth=32 blks
naming    =version 2              bsize=4096   ascii-ci=0, ftype=1
log       =internal log          bsize=4096   blocks=521728, version=2
=                               sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none                  extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/index0
/bin/mount /var/netwitness/concentrator/index0
```

20. Run `devlist` to verify block devices have been configured.

Properties for /appliance

devlist ▾

Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="5.82 TB" used=1
sdd: vendor=DELL model=RAID size="5.82 TB" used=1
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sde: vendor=DELL model=RAID size="145.52 TB" used=1
sdc: vendor=DELL model=RAID size="145.52 TB" used=1
sda: vendor=DELL model=RAID size="2.18 TB" used=1
```

21. Verify partitions and mounts on lsblk.

```
[root@S7Concentrator ~]# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	2.2T	0	disk	
├─sda1	8:1	0	256M	0	part	/boot/efi
├─sda2	8:2	0	1G	0	part	/boot
└─sda3	8:3	0	2.2T	0	part	
├─netwitness_vg00-nwhome	253:0	0	4.3T	0	lvm	/var/netwitness
├─netwitness_vg00-varlog	253:1	0	10G	0	lvm	/var/log
├─netwitness_vg00-usrhome	253:2	0	10G	0	lvm	/home
├─netwitness_vg00-root	253:3	0	30G	0	lvm	/
└─netwitness_vg00-swap	253:4	0	4G	0	lvm	[SWAP]
sdb	8:16	0	2.2T	0	disk	
├─sdb1	8:17	0	2.2T	0	part	
└─netwitness_vg00-nwhome	253:0	0	4.3T	0	lvm	/var/netwitness
sdc	8:32	0	145.5T	0	disk	
├─sdc1	8:33	0	145.5T	0	part	
├─concentrator-root	253:5	0	30G	0	lvm	/var/netwitness/concentrator
├─concentrator-sessiondb	253:6	0	14.6T	0	lvm	/var/netwitness/concentrator/sessiondb
└─concentrator-metadb	253:7	0	131T	0	lvm	/var/netwitness/concentrator/metadb
sdd	8:48	0	5.8T	0	disk	
├─sdd1	8:49	0	5.8T	0	part	
└─index-index	253:8	0	5.8T	0	lvm	/var/netwitness/concentrator/index
sde	8:64	0	145.5T	0	disk	
├─sde1	8:65	0	145.5T	0	part	
├─concentrator0-sessiondb	253:9	0	14.6T	0	lvm	/var/netwitness/concentrator/sessiondb0
└─concentrator0-metadb	253:10	0	131T	0	lvm	/var/netwitness/concentrator/metadb0
sdf	8:80	0	5.8T	0	disk	
├─sdf1	8:81	0	5.8T	0	part	
└─index0-index	253:11	0	5.8T	0	lvm	/var/netwitness/concentrator/index0

22. Execute the `srvAlloc` commands with the following parameters.

```
service=concentrator volume=concentrator commit=1
```

Properties for /appliance
 srvAlloc ▾ Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}>
  volume - <string, {enum-one:The value must be one of the following: concentrator|concentrator0|index|index0|netw
  commit - <bool, optional> commit changes
  
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==124.39 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==13.82 TB
  
```

```
service=concentrator volume=index commit=1
```

Properties for /appliance
 srvAlloc ▾ Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> se
  volume - <string, {enum-one:The value must be one of the following: concentrator|concentrator0|index|index0|netw
  commit - <bool, optional> commit changes
  
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=index&commit=1

Output (or command manual help)

```

Set /index/config/index.dir to /var/netwitness/concentrator/index==5.53 TB
  
```

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance
 srvAlloc ▾ Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
  volume - <string, {enum-one:The value must be one of the following: concentrator|concentrator0|index|index0|netwitness_vg00}> volume group name
  commit - <bool, optional> commit changes
  
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

```

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==124.39 TB;/var/netwitness/concentrator/metadb0==124.42 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==13.82 TB;/var/netwitness/concentrator/sessiondb0==13.82 TB
  
```

```
service=concentrator volume=index0 commit=1
```

Properties for /appliance
 srvAlloc Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use stor-
  volume - <string, {enum-one:The value must be one of the following: concentrator|concentrator0|index|index0|netwitness_vg00}> volume group
  commit - <bool, optional> commit changes
  
```

Output (or command manual help)

```

Set /index/config/index.dir to /var/netwitness/concentrator/index==5.53 TB;/var/netwitness/concentrator/index0==5.53 TB
  
```

23. Execute `srvList` to verify storage has been assigned as required.

Properties for /appliance
 srvList Parameters:

Message Help

```

srvList: list of core services on this appliance
security.roles: appliance.manage
  
```

Output (or command manual help)

```

localhost:56005: type=concentrator
meta.dir="/var/netwitness/concentrator/metadb==124.39 TB;/var/netwitness/concentrator/metadb0==124.42 TB"
session.dir="/var/netwitness/concentrator/sessiondb==13.82 TB;/var/netwitness/concentrator/sessiondb0==13.82 TB"
index.dir="/var/netwitness/concentrator/index==5.53 TB;/var/netwitness/concentrator/index0==5.53 TB"
  
```

24. Verify partitions and mounts on `lsblk`.

```

[root@S7Concentrator ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                 8:0      0   2.2T  0 disk
├─sda1                             8:1      0   256M  0 part /boot/efi
├─sda2                             8:2      0    1G   0 part /boot
├─sda3                             8:3      0   2.2T  0 part
│   ├─netwitness_vg00-nwhome        253:0      0   4.3T  0 lvm  /var/netwitness
│   ├─netwitness_vg00-varlog        253:1      0   10G   0 lvm  /var/log
│   ├─netwitness_vg00-usrhome       253:2      0   10G   0 lvm  /home
│   ├─netwitness_vg00-root          253:3      0   30G   0 lvm  /
│   └─netwitness_vg00-swap          253:4      0    4G   0 lvm  [SWAP]
sdb                                 8:16     0   2.2T  0 disk
├─sdb1                             8:17     0   2.2T  0 part
│   └─netwitness_vg00-nwhome        253:0      0   4.3T  0 lvm  /var/netwitness
sdc                                 8:32     0  145.5T  0 disk
├─sdc1                             8:33     0  145.5T  0 part
│   ├─concentrator-root             253:5      0   30G   0 lvm  /var/netwitness/concentrator
│   ├─concentrator-sessiondb        253:6      0  14.6T  0 lvm  /var/netwitness/concentrator/sessiondb
│   └─concentrator-metadb           253:7      0  131T   0 lvm  /var/netwitness/concentrator/metadb
sdd                                 8:48     0   5.8T  0 disk
├─sdd1                             8:49     0   5.8T  0 part
│   └─index-index                  253:8      0   5.8T  0 lvm  /var/netwitness/concentrator/index
sde                                 8:64     0  145.5T  0 disk
├─sde1                             8:65     0  145.5T  0 part
│   ├─concentrator0-sessiondb       253:9      0  14.6T  0 lvm  /var/netwitness/concentrator/sessiondb0
│   └─concentrator0-metadb          253:10     0  131T   0 lvm  /var/netwitness/concentrator/metadb0
sdf                                 8:80     0   5.8T  0 disk
├─sdf1                             8:81     0   5.8T  0 part
│   └─index0-index                 253:11     0   5.8T  0 lvm  /var/netwitness/concentrator/index0

```

Configure Storage for Log Hybrid using MD2412

The following scenario configures storage on one unencrypted, 12-Drive MD2412 PowerVault for a Log Hybrid physical host.

1. Execute the `raidList` command.

- a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

In Use: FALSE

Devices: <empty>

- b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance
 raidList Parameters: Send

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Model: BP_PSV
In Use: true
Drives: 2.182 TB HDD x 2
        14.552 TB HDD x 10
Devices: sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
        sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
        sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
        sdf /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 1 at PCI Address b5:00.0, Enclosure 288, SCSI Channel 2
Vendor: DELL
Model: DSES Enclosure
In Use: false
Drives: 14.552 TB HDD x 12
Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=288 scheme=log-hybrid commit=1`

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance

raidNew

Parameters: controller=1 enclosure=288 scheme=log-hybrid commit=1

Send

Message Help

raidNew: allocate RAID devices in a drive shelf

security.roles: appliance.manage

parameters:

controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 288|294|288}> Enclosure number of the shelf

scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=288&scheme=log-hybrid&commit=1

Output (or command manual help)

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=288:0,288:1,288:2,288:3,288:4,288:5 nora Strip=64

CLI Version = 008.0004.0000.0022 Apr 28, 2023

Operating system = Linux4.18.0-553.16.1.el8_10.x86_64

Controller = 1

Status = Success

Description = Add VD succeeded.

VD Information :

=====

VDID	VDSIZE	Status	ErrType	ErrCd	Msg
1	72.760 TiB	Success	-	-	-

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=288:6,288:7,288:8,288:9,288:10,288:11 nora Strip=64

CLI Version = 008.0004.0000.0022 Apr 28, 2023

Operating system = Linux4.18.0-553.16.1.el8_10.x86_64

Execute the raidList command to verify the RAID created.

Properties for /appliance

raidList ▾

Parameters:

Send

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
In Use: true
Drives: 2.182 TB HDD x 2
        14.552 TB HDD x 10
Devices: sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
        sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
        sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
        sdf /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 1 at PCI Address b5:00.0, Enclosure 288, SCSI Channel 0
Vendor: DELL
Model: DSES Enclosure
In Use: true
Drives: 14.552 TB HDD x 12
Devices: sdg /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:1:0
        sdh /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:2:0
```

Execute the `devlist` command to find the exact memory distribution between the newly created RAIDS. In Log Hybrid, the distribution is exactly equal so any service can be installed in any of the RAIDS.

Properties for /appliance

devlist ▾

Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="2.18 TB" used=1
sdd: vendor=DELL model=RAID size="43.66 TB" used=1
sdb: vendor=DELL model=RAID size="14.55 TB" used=1
sdg: vendor=DELL model=RAID size="72.76 TB" used=0
sde: vendor=DELL model=RAID size="3.49 TB" used=1
sdc: vendor=DELL model=RAID size="43.66 TB" used=1
sdh: vendor=DELL model=RAID size="72.76 TB" used=0
```

Execute `lsblk` to verify the block devices created.

```
[root@S7LogHybrid ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                8:0      1 14.3G  0 disk
├─sda1                            8:1      1 14.3G  0 part
├─sdb                             8:16     0 14.6T  0 disk
│ └─sdb1                         8:17     0 14.6T  0 part
│   └─hybrid--logdecoder--meta-decoroot 253:8     0 14.6T  0 lvm /var/netwitness/logdecoder
├─sdc                             8:32     0 43.7T  0 disk
│ └─sdc1                         8:33     0 43.7T  0 part
│   └─logpacket-packetdb         253:7     0 43.7T  0 lvm /var/netwitness/logdecoder/packetdb
├─sdd                             8:48     0 43.7T  0 disk
│ └─sdd1                         8:49     0 43.7T  0 part
│   └─hybrid--concentrator-root 253:6     0 43.7T  0 lvm /var/netwitness/concentrator
├─sde                             8:64     0 3.5T   0 disk
│ └─sde1                         8:65     0 3.5T   0 part
│   └─index-index               253:5     0 3.5T   0 lvm /var/netwitness/concentrator/index
├─sdf                             8:80     0 2.2T   0 disk
│ ├──sdf1                       8:81     0 256M   0 part /boot/efi
│ ├──sdf2                       8:82     0 1G     0 part /boot
│ └─sdf3                       8:83     0 2.2T   0 part
│   ├──netwitness_vg00-nwhome 253:0     0 2.1T   0 lvm /var/netwitness
│   ├──netwitness_vg00-varlog 253:1     0 10G    0 lvm /var/log
│   ├──netwitness_vg00-usrhome 253:2     0 10G    0 lvm /home
│   ├──netwitness_vg00-root   253:3     0 30G    0 lvm /
│   └─netwitness_vg00-swap    253:4     0 4G     0 lvm [SWAP]
└─sdg                             8:96     0 72.8T  0 disk
    └─sdh                         8:112    0 72.8T  0 disk
```

3. Execute the `partNew` command with the following parameters to create partitions and mount points in the `/etc/fstab` file.

name=sdg service=logdecoder volume=logdecoder commit=1

Properties for /appliance

partNew Parameters: name=sdg service=logdecoder volume=logdecoder commit=1

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdg|sdh}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> servi
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder
hybrid|log-indexed-decoder|logindex}> volume to create
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdg&service=logdecoder&volume=logdecoder&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdg mklabel gpt
/sbin/parted -s -a optimal /dev/sdg mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdg1
Physical volume "/dev/sdg1" successfully created.
/sbin/vgcreate -f logdecoder0 /dev/sdg1
Volume group "logdecoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE logdecoder0
Wiping xfs signature on /dev/logdecoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/logdecoder0/packetdb
meta-data=/dev/logdecoder0/packetdb isize=512    agcount=73, agsize=268435424 blks
        =                               sectsz=4096  attr=2, projid32bit=1
        =                               crc=1        finobt=1, sparse=1, rmapbt=0
        =                               reflink=1    bigtime=0  inobtcount=0
data      =                               bsize=4096  blocks=19531693056, imaxpct=1
        =                               sunit=16     swidth=80  blks
naming    =version 2                       bsize=4096  ascii-ci=0, ftype=1
log        =internal log                   bsize=4096  blocks=521728, version=2
        =                               sectsz=4096  sunit=1 blks, lazy-count=1
realtime   =none                          extsz=4096  blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/logdecoder/packetdb
/bin/mount /var/netwitness/logdecoder/packetdb
```

name=sdh service=concentrator volume=concentrator commit=1

Properties for /appliance

partNew Parameters: name=sdh service=concentrator volume=concentrator commit=1

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdg|sdh}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use st
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-deco
hybrid|log-indexed-decoder|logindex}> volume to create
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdh&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdh mklabel gpt
/sbin/parted -s -a optimal /dev/sdh mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdh1
Physical volume "/dev/sdh1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdh1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512    agcount=32, agsize=61036512 blks
        =                               sectsz=4096  attr=2, projid32bit=1
        =                               crc=1        finobt=1, sparse=1, rmapbt=0
        =                               reflink=1    bigtime=0  inobtcount=0
data      =                               bsize=4096  blocks=1953168384, imaxpct=5
        =                               sunit=16     swidth=80  blks
naming    =version 2                       bsize=4096  ascii-ci=0, ftype=1
log        =internal log                   bsize=4096  blocks=521728, version=2
        =                               sectsz=4096  sunit=1 blks, lazy-count=1
realtime   =none                          extsz=4096  blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb
/bin/mount /var/netwitness/concentrator/sessiondb
/sbin/lvcreate -y -n metadb -l 100%FREE concentrator0
```

Execute `devlist` to verify partitions created.

Properties for /appliance
 devlist Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="2.18 TB" used=1
sdd: vendor=DELL model=RAID size="43.66 TB" used=1
sdb: vendor=DELL model=RAID size="14.55 TB" used=1
sdg: vendor=DELL model=RAID size="72.76 TB" used=1
sde: vendor=DELL model=RAID size="3.49 TB" used=1
sdc: vendor=DELL model=RAID size="43.66 TB" used=1
sdh: vendor=DELL model=RAID size="72.76 TB" used=1
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to logdecoder and concentrator services.

- a. This adds storage to the logdecoder service configuration and restarts the service every time it is executed.

```
service=logdecoder volume=logdecoder0 commit=1
```

Properties for /appliance
 srvAlloc Parameters: Send

Message Help

```
srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use stora
  volume - <string, {enum-one:The value must be one of the following: concentrator0|hybrid-concentrator|hybrid-logdecoder-meta|index|logdecod
  commit - <bool, optional> commit changes
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=logdecoder&volume=logdecoder0&commit=1

Output (or command manual help)

```
Set /database/config/packet.dir to /var/netwitness/logdecoder/packetdb/packetdb==41.47 TB;/var/netwitness/logdecoder/packetdb0==69.12 TB
```

- b. This adds storage to concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance

srvAlloc Parameters: service=concentrator volume=concentrator0 commit=1

Message Help

srvAlloc: apply volume group storage to a service on this appliance
 security.roles: appliance.manage
 parameters:
 service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
 volume - <string, {enum-one:The value must be one of the following: concentrator0|hybrid-concentrator|hybrid-logdecoder-meta|index|logdecoder}>
 commit - <bool, optional> commit changes

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==37.71 TB;/var/netwitness/concentrator/metadb0==62.21 TB
 Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==3.73 TB;/var/netwitness/concentrator/sessiondb0==6.91 TB

- c. Execute the `srvList` command to verify the services allocated on the appliance.

Properties for /appliance

srvList Parameters:

Message Help

srvList: list of core services on this appliance
 security.roles: appliance.manage

/appliance?msg=srvList&force-content-type=text/plain

Output (or command manual help)

localhost:56002: type=logdecoder
 packet.dir="/var/netwitness/logdecoder/packetdb/packetdb==41.47 TB;/var/netwitness/logdecoder/packetdb0==69.12 TB"
 meta.dir="/var/netwitness/logdecoder/metadb==12 TB"
 session.dir="/var/netwitness/logdecoder/sessiondb==1.24 TB"
 index.dir="/var/netwitness/logdecoder/index==565.04 GB"
 localhost:56005: type=concentrator
 meta.dir="/var/netwitness/concentrator/metadb==37.71 TB;/var/netwitness/concentrator/metadb0==62.21 TB"
 session.dir="/var/netwitness/concentrator/sessiondb==3.73 TB;/var/netwitness/concentrator/sessiondb0==6.91 TB"
 index.dir="/var/netwitness/concentrator/index==3.32 TB"

5. Execute the `lsblk` command in backend to see all the raids and partitions inside the service.

```

[root@S7LogHybrid ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      1 14.3G  0 disk
├─sda1                              8:1      1 14.3G  0 part
├─sdb                                8:16     0 14.6T  0 disk
├─┌sdb1                            8:17     0 14.6T  0 part
│ └─hybrid--logdecoder--meta-decoroot 253:8     0 14.6T  0 lvm  /var/netwitness/logdecoder
├─sdc                                8:32     0 43.7T  0 disk
├─┌sdc1                            8:33     0 43.7T  0 part
│ └─logpacket-packetdb              253:7     0 43.7T  0 lvm  /var/netwitness/logdecoder/packetdb
├─sdd                                8:48     0 43.7T  0 disk
├─┌sdd1                            8:49     0 43.7T  0 part
│ └─hybrid--concentrator-root        253:6     0 43.7T  0 lvm  /var/netwitness/concentrator
├─sde                                8:64     0  3.5T  0 disk
├─┌sde1                            8:65     0  3.5T  0 part
│ └─index-index                    253:5     0  3.5T  0 lvm  /var/netwitness/concentrator/index
├─sdf                                8:80     0  2.2T  0 disk
├─┌sdf1                            8:81     0 256M  0 part  /boot/efi
│ └─sdf2                            8:82     0  1G    0 part  /boot
├─┌sdf3                            8:83     0  2.2T  0 part
│ └─netwitness_vg00-nwhome          253:0     0  2.1T  0 lvm  /var/netwitness
│   └─netwitness_vg00-varlog        253:1     0  10G   0 lvm  /var/log
│     └─netwitness_vg00-usrhome      253:2     0  10G   0 lvm  /home
│       └─netwitness_vg00-root       253:3     0  30G   0 lvm  /
│         └─netwitness_vg00-swap     253:4     0  4G    0 lvm  [SWAP]
├─sdg                                8:96     0 72.8T  0 disk
├─┌sdg1                            8:97     0 72.8T  0 part
│ └─logdecoder0-packetdb            253:9     0 72.8T  0 lvm  /var/netwitness/logdecoder/packetdb0
├─sdh                                8:112    0 72.8T  0 disk
├─┌sdh1                            8:113    0 72.8T  0 part
│ └─concentrator0-sessiondb         253:10    0  7.3T  0 lvm  /var/netwitness/concentrator/sessiondb0
└─└concentrator0-metadb             253:11    0 65.5T  0 lvm  /var/netwitness/concentrator/metadb0

```

Configure Storage for Network Hybrid using MD2412

The following scenario configures storage on an unencrypted, 12-Drive MD2412 PowerVault for a Network Hybrid physical host.

1. Execute the `raidList` command.

- a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.

You should see the following information.

In Use: FALSE

Devices: <empty>

- b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance
raidList ▼ Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 2.182 TB HDD x 2
        14.552 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
         sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
         sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
         sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 1 at PCI Address b5:00.0, Enclosure 288, SCSI Channel 2
Vendor: DELL
Model: DSES Enclosure
In Use: false
Drives: 14.552 TB HDD x 12
Devices:
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

`controller=1 enclosure=288 scheme=network-hybrid preferSecure=0 commit=1`

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance
 raidNew Parameters: controller=1 enclosure=288 scheme=network-hybrid preferSecure=0 commit=1

Message Help

```
raidNew: allocate RAID devices in a drive shelf
security.roles: appliance.manage
parameters:
  controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to
  enclosure - <uint32, optional, {enum-one:The value must be one of the following: 288|294|288}> Enclosure number of the shelf to clear. Required if
  scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|log-indexed-d
  metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate
```

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=288&scheme=network-hybrid&preferSecure=0&commit=1

Output (or command manual help)

```
/opt/MegaRAID/perccli2/perccli2 /c1 add vd r5 drives=288:0,288:1,288:2 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-513.24.1.el8_9.x86_64
Controller = 1
Status = Success
Description = Add VD succeeded.

VD Information :
=====
VDID VDSIZE Status ErrType ErrCd Msg
-----
1 29.104 TiB Success - - -

/opt/MegaRAID/perccli2/perccli2 /c1 add vd r5 drives=288:3,288:4,288:5,288:6,288:7,288:8,288:9,288:10,288:11 nora Strip=64
CLI Version = 008.0004.0000.0022 Apr 28, 2023
Operating system = Linux4.18.0-513.24.1.el8_9.x86_64
Controller = 1
```

Following is the response from raidList after creating RAID array.

Properties for /appliance
 raidList Parameters:

Message Help

```
raidList: list drive shelves attached to this appliance
security.roles: appliance.manage
```

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```
Drives: 2.102 TB HDD x 2
14.552 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 1 at PCI Address b5:00.0, Enclosure 288, SCSI Channel 0
Vendor: DELL
Model: DSES Enclosure
In Use: true
Drives: 14.552 TB HDD x 12
Devices: sdf /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:1:0
sdh /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:2:0
```

Execute `devlist` command to find the exact memory distribution between the newly created RAIDs. In Network Hybrid, the memory for 'sdh' block device is more than 'sdf' block device, therefore packetdb will be installed in the device with higher memory allocated which is 'sdh'.

Properties for /appliance
 devlist ▼ Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="29.1 TB" used=0
sdd: vendor=DELL model=RAID size="14.55 TB" used=1
sdb: vendor=DELL model=RAID size="2.18 TB" used=1
sde: vendor=DELL model=RAID size="3.49 TB" used=1
sdc: vendor=DELL model=RAID size="29.1 TB" used=1
sda: vendor=DELL model=RAID size="58.21 TB" used=1
sdh: vendor=DELL model=RAID size="116.42 TB" used=0
```

- Execute the `partNew` command with the following parameters to create partitions and mount points in the `/etc/fstab` file.

`name=sdf service=concentrator volume=concentrator commit=1`

Properties for /appliance
 partNew ▼ Parameters:

Message Help

```
partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdf|sdh}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that w.
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|h
  hybrid|log-indexed-decoder|logindex}> volume to create
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdf&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```
/sbin/parted -s /dev/sdf mklabel gpt
/sbin/parted -s -a optimal /dev/sdf mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdf1
Physical volume "/dev/sdf1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdf1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512    agcount=32, agsize=24414576 blks
        =                       sectsz=4096   attr=2, projid32bit=1
        =                       crc=1          finobt=1, sparse=1, rmapbt=0
        =                       reflink=1       bigtime=0 inobtcount=0
data      =                       bsize=4096   blocks=781266432, imaxpct=5
        =                       sunit=16      swidth=32 blks
naming    =version 2                  bsize=4096   ascii-ci=0, ftype=1
log       =internal log              bsize=4096   blocks=381477, version=2
        =                       sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none                      extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb0
/bin/mount /var/netwitness/concentrator/sessiondb0
```

`name=sdh service=decoder volume=decoder commit=1`

Properties for /appliance
 partNew Parameters: Send

Message Help

```

partNew: make partitions on a block device
security.roles: appliance.manage
parameters:
  name - <string, {enum-one:The value must be one of the following: sdf|sdh}> block device name
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service
  volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|hybrid|log-indexed-decoder|logindex}> volume to create
  
```

/appliance?msg=partNew&force-content-type=text/plain&name=sdh&service=decoder&volume=decoder&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdh mklabel gpt
/sbin/parted -s -a optimal /dev/sdh mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdh1
Physical volume "/dev/sdh1" successfully created.
/sbin/vgcreate -f decoder0 /dev/sdh1
Volume group "decoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE decoder0
Wiping xfs signature on /dev/decoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/decoder0/packetdb
meta-data=/dev/decoder0/packetdb isize=512    agcount=117, agsize=268435440 blks
       =                               sectsz=4096   attr=2, projid32bit=1
       =                               crc=1        finobt=1, sparse=1, rmapbt=0
       =                               reflink=1     bigtime=0 inobtcount=0
data    =                               bsize=4096   blocks=31250709504, imaxpct=1
       =                               sunit=16     swidth=128 blks
naming  =version 2                       bsize=4096   ascii-ci=0, ftype=1
log      =internal log                   bsize=4096   blocks=521728, version=2
       =                               sectsz=4096   sunit=1 blks, lazy-count=1
realtime =none                          extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/decoder/packetdb0
/bin/mount /var/netwitness/decoder/packetdb0
  
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to decoder and concentrator services.

This adds storage to the decoder service configuration and restarts the service every time it is executed.

```
service=decoder volume=decoder0 commit=1
```

Properties for /appliance
 srvAlloc Parameters: Send

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
  volume - <string, {enum-one:The value must be one of the following: concentrator0|decoder0|hybrid-concentrator|hybrid-decoder-meta|index|netwitness_vg00|packet}> vc
  commit - <bool, optional> commit changes
  
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=decoder&volume=decoder0&commit=1

Output (or command manual help)

```

Set /database/config/packet.dir to /var/netwitness/decoder/packetdb/packetdb==27.65 TB;/var/netwitness/decoder/packetdb0==110.6 TB
  
```

This adds storage to the concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance
 srvAlloc Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
  volume - <string, {enum-one:The value must be one of the following: concentrator0|decoder0|hybrid-concentrator|hybrid-decoder-meta|index|netwitness_
  commit - <bool, optional> commit changes
  
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

```

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==12.55 TB;/var/netwitness/concentrator/metadb0==24.88 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==1.24 TB;/var/netwitness/concentrator/sessiondb0==2.76 TB
  
```

Execute the `srvList` command to verify the services allocated on the appliance.

Properties for /appliance
 srvList Parameters:

Message Help

```

srvList: list of core services on this appliance
security.roles: appliance.manage
  
```

/appliance?msg=srvList&force-content-type=text/plain

Output (or command manual help)

```

localhost:56004: type=decoder
packet.dir="/var/netwitness/decoder/packetdb/packetdb==27.65 TB;/var/netwitness/decoder/packetdb0==110.6 TB"
meta.dir="/var/netwitness/decoder/metadb==48.08 TB"
session.dir="/var/netwitness/decoder/sessiondb==4.97 TB"
index.dir="/var/netwitness/decoder/index==2.21 TB"
localhost:56005: type=concentrator
meta.dir="/var/netwitness/concentrator/metadb==12.55 TB;/var/netwitness/concentrator/metadb0==24.88 TB"
session.dir="/var/netwitness/concentrator/sessiondb==1.24 TB;/var/netwitness/concentrator/sessiondb0==2.76 TB"
index.dir="/var/netwitness/concentrator/index==3.32 TB"
  
```

5. Execute `lsblk` command to list all the raids and partitions in the service.

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	58.2T	0	disk	
└─sda1	8:1	0	58.2T	0	part	
└─┬─hybrid--decoder--meta-decoroot	253:2	0	58.2T	0	lvm	/var/netwitness/decoder
sdb	8:16	0	2.2T	0	disk	
└─sdb1	8:17	0	256M	0	part	/boot/efi
└─sdb2	8:18	0	1G	0	part	/boot
└─sdb3	8:19	0	2.2T	0	part	
└─┬─netwitness_vg00-root	253:0	0	30G	0	lvm	/
└─┬─netwitness_vg00-swap	253:1	0	4G	0	lvm	[SWAP]
└─┬─netwitness_vg00-nwhome	253:3	0	2.1T	0	lvm	/var/netwitness
└─┬─netwitness_vg00-varlog	253:4	0	10G	0	lvm	/var/log
└─┬─netwitness_vg00-usrhome	253:5	0	10G	0	lvm	/home
sdc	8:32	0	29.1T	0	disk	
└─sdc1	8:33	0	29.1T	0	part	
└─┬─packet-packetdb	253:6	0	29.1T	0	lvm	/var/netwitness/decoder/packetdb
sdd	8:48	0	14.6T	0	disk	
└─sdd1	8:49	0	14.6T	0	part	
└─┬─hybrid--concentrator-root	253:7	0	14.6T	0	lvm	/var/netwitness/concentrator
sde	8:64	0	3.5T	0	disk	
└─sde1	8:65	0	3.5T	0	part	
└─┬─index-index	253:8	0	3.5T	0	lvm	/var/netwitness/concentrator/index
sdf	8:80	0	29.1T	0	disk	
└─sdf1	8:81	0	29.1T	0	part	
└─┬─concentrator0-sessiondb	253:9	0	2.9T	0	lvm	/var/netwitness/concentrator/sessiondb0
└─┬─concentrator0-metadb	253:10	0	26.2T	0	lvm	/var/netwitness/concentrator/metadb0
sdg	8:96	1	14.3G	0	disk	
└─sda1	8:97	1	14.3G	0	part	
sdh	8:112	0	116.4T	0	disk	
└─sdh1	8:113	0	116.4T	0	part	
└─┬─decoder0-packetdb	253:11	0	116.4T	0	lvm	/var/netwitness/decoder/packetdb0

Configure Storage for Endpoint Log Hybrid using MD2412

The following scenario configures storage on one unencrypted, 12-Drive MD2412 PowerVault for an Endpoint Log Hybrid physical host.

1. Execute the `raidList` command.
 - a. Record the Controller Number, Enclosure Number, In Use, Drives, and Devices.
You should see the following information.
In Use: FALSE
Devices: <empty>

b. Verify the Drive Count, Size, and Vendor.

The following example illustrates what you should see before you create a RAID array.

Properties for /appliance
 raidList Parameters:

Message Help

raidList: list drive shelves attached to this appliance
 security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

In Use: true
Drives: 2.182 TB HDD x 2
        14.552 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
        sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
        sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
        sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 1 at PCI Address b5:00.0, Enclosure 302, SCSI Channel 2
Vendor: DELL
Model: DSES Enclosure
In Use: false
Drives: 14.552 TB HDD x 12
Devices:
  
```

2. Execute the `raidNew` command with the following parameters using the controller number and the enclosure number you just recorded.

In this case, the security key has been set on the controller PERCH965e. Hence, `preferSecure=1` parameter is used to enable encryption. If the security key is not set on the controller, parameter `preferSecure=0` can be used to complete configuring RAID on the controller.

```
controller=1 enclosure=302 scheme=log-hybrid preferSecure=1 commit=1
```

The following example illustrates what you should see after you create a RAID array.

Properties for /appliance

raidNew

Parameters: controller=1 enclosure=302 scheme=log-hybrid preferSecure=1 commit=1

Send

Message Help

raidNew: allocate RAID devices in a drive shelf

security.roles: appliance.manage

parameters:

controller - <uint32, {enum-one:The value must be one of the following: 0,1}> Controller the shelf is attached to

enclosure - <uint32, optional, {enum-one:The value must be one of the following: 288|294|302}> Enclosure number of the shelf to clear.

scheme - <string, {enum-one:The value must be one of the following: decoder|logdecoder|concentrator|archiver|network-hybrid|log-hybrid|logdecoder-metakit|logdecoder-metakit|concentrator-metakit}> Type of RAID volumes to allocate

/appliance?msg=raidNew&force-content-type=text/plain&controller=1&enclosure=302&scheme=log-hybrid&preferSecure=1&commit=1

Output (or command manual help)

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=302:0,302:1,302:2,302:3,302:4,302:5 nora Strip=64

CLI Version = 008.0004.0000.0022 Apr 28, 2023

Operating system = Linux4.18.0-513.24.1.el8_9.x86_64

Controller = 1

Status = Success

Description = Add VD succeeded.

VD Information :

VDID	VDSize	Status	ErrType	ErrCd	Msg
1	72.760 TiB	Success	-	-	-

/opt/MegaRAID/percccli2/percccli2 /c1 add vd r5 drives=302:6,302:7,302:8,302:9,302:10,302:11 nora Strip=64

CLI Version = 008.0004.0000.0022 Apr 28, 2023

Operating system = Linux4.18.0-513.24.1.el8_9.x86_64

Following is the response from raidList after creating RAID array.

Properties for /appliance

raidList Parameters:

Message Help

raidList: list drive shelves attached to this appliance
security.roles: appliance.manage

/appliance?msg=raidList&force-content-type=text/plain

Output (or command manual help)

```

Drives: 2.182 TB HDD x 2
        14.552 TB HDD x 10
Devices: sda /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:1:0
        sdb /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:2:0
        sdc /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:3:0
        sde /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:5:0

Controller 0 at PCI Address 37:00.0, Enclosure 294, SCSI Channel 0
Vendor: DP
Model: BP_PSV
In Use: true
Drives: 3.492 TB SSD x 2
Devices: sdd /dev/disk/by-path/pci-0000:37:00.0-scsi-0:0:4:0

Controller 1 at PCI Address b5:00.0, Enclosure 302, SCSI Channel 0
Vendor: DELL
Model: DSES Enclosure
In Use: true
Drives: 14.552 TB HDD x 12
Devices: sdf /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:1:0
        sdg /dev/disk/by-path/pci-0000:b5:00.0-scsi-0:0:2:0

```

Execute the `devlist` command to find the exact memory distribution for the newly created raids. In Endpoint Log Hybrid, both are exactly equal. Hence, any service can be installed in any of the raid.

Properties for /appliance

devlist ▼ Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="72.76 TB" used=0
sdd: vendor=DELL model=RAID size="3.49 TB" used=1
sdb: vendor=DELL model=RAID size="43.66 TB" used=1
sdg: vendor=DELL model=RAID size="72.76 TB" used=0
sde: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="14.55 TB" used=1
sda: vendor=DELL model=RAID size="29.1 TB" used=1
```

Execute the `lsblk` command at backend to list the newly created raids on the Endpoint Log Hybrid appliance.

```
[root@S7EndpointLogHybrid ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                 8:0      0 29.1T 0 disk
├─sda1                             8:1      0 29.1T 0 part
│   └─hybrid--logdecoder--meta-decoroot 253:8    0 29.1T 0 lvm  /var/netwitness/logdecoder
sdb                                 8:16     0 43.7T 0 disk
├─sdb1                             8:17     0 43.7T 0 part
│   └─endpoint--log--hybrid-mongo      253:5    0 21.8T 0 lvm  /var/netwitness/mongo
│       └─endpoint--log--hybrid-packetdb 253:7    0 21.8T 0 lvm  /var/netwitness/logdecoder/packetdb
sdc                                 8:32     0 14.6T 0 disk
├─sdc1                             8:33     0 14.6T 0 part
│   └─hybrid--concentrator-root        253:9    0 14.6T 0 lvm  /var/netwitness/concentrator
sdd                                 8:48     0  3.5T 0 disk
├─sdd1                             8:49     0  3.5T 0 part
│   └─index-index                     253:6    0  3.5T 0 lvm  /var/netwitness/concentrator/index
sde                                 8:64     0  2.2T 0 disk
├─sde1                             8:65     0 256M 0 part /boot/efi
├─sde2                             8:66     0  1G  0 part /boot
└─sde3                             8:67     0  2.2T 0 part
    ├──netwitness_vg00-root          253:0    0  30G 0 lvm  /
    ├──netwitness_vg00-swap          253:1    0   4G 0 lvm  [SWAP]
    ├──netwitness_vg00-nwhome        253:2    0  2.1T 0 lvm  /var/netwitness
    ├──netwitness_vg00-varlog        253:3    0  10G 0 lvm  /var/log
    └─netwitness_vg00-usrhome         253:4    0  10G 0 lvm  /home
sdf                                 8:80     0 72.8T 0 disk
sdg                                 8:96     0 72.8T 0 disk
```

- Execute the `partNew` command with the following parameters to create partitions and mount points as required.

```
name=sdf service=logdecoder volume=logdecoder commit=1
```

Properties for /appliance

partNew Parameters: name=sdf service=logdecoder volume=logdecoder commit=1 Send

Message Help

```

name - <string, {enum-one:The value must be one of the following: sdf|sdg}> block device name
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storage
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|decoder|packet|hybrid-decoder-me
hybrid|log-indexed-decoder|logindex}> volume to create
commit - <bool, optional> commit changes

```

/appliance?msg=partNew&force-content-type=text/plain&name=sdf&service=logdecoder&volume=logdecoder&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdf mklabel gpt
/sbin/parted -s -a optimal /dev/sdf mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdf1
Physical volume "/dev/sdf1" successfully created.
/sbin/vgcreate -f logdecoder0 /dev/sdf1
Volume group "logdecoder0" successfully created
/sbin/lvcreate -y -n packetdb -l 100%FREE logdecoder0
Wiping xfs signature on /dev/logdecoder0/packetdb.
Logical volume "packetdb" created.
/sbin/mkfs.xfs /dev/logdecoder0/packetdb
meta-data=/dev/logdecoder0/packetdb isize=512    agcount=73, agsize=268435424 blks
       =                                           sectsz=4096   attr=2, projid32bit=1
       =                                           crc=1        finobt=1, sparse=1, rmapbt=0
       =                                           reflink=1    bigtime=0 inobtcount=0
data      =                                           bsize=4096   blocks=19531693056, imaxpct=1
       =                                           sunit=16     swidth=80 blks
naming    =version 2                  bsize=4096   ascii-ci=0, ftype=1
log        =internal log              bsize=4096   blocks=521728, version=2
       =                                           sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none                      extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/logdecoder/packetdb0
/bin/mount /var/netwitness/logdecoder/packetdb0

```

name=sdg service=concentrator volume= concentrator commit=1

Properties for /appliance

partNew Parameters: name=sdg service=concentrator volume=concentrator commit=1 Send

Message Help

```

name - <string, {enum-one:The value must be one of the following: sdf|sdg}> block device name
service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}>
volume - <string, optional, {enum-one:The value must be one of the following: concentrator|index|decodersmall|d
hybrid|log-indexed-decoder|logindex}> volume to create
commit - <bool, optional> commit changes

```

/appliance?msg=partNew&force-content-type=text/plain&name=sdg&service=concentrator&volume=concentrator&commit=1

Output (or command manual help)

```

/sbin/parted -s /dev/sdg mklabel gpt
/sbin/parted -s -a optimal /dev/sdg mkpart LVM 0% 100%
/sbin/pvcreate -f /dev/sdg1
Physical volume "/dev/sdg1" successfully created.
/sbin/vgcreate -f concentrator0 /dev/sdg1
Volume group "concentrator0" successfully created
/sbin/lvcreate -y -n sessiondb -l 10%FREE concentrator0
Wiping xfs signature on /dev/concentrator0/sessiondb.
Logical volume "sessiondb" created.
/sbin/mkfs.xfs /dev/concentrator0/sessiondb
meta-data=/dev/concentrator0/sessiondb isize=512    agcount=32, agsize=61036512 blks
       =                                           sectsz=4096   attr=2, projid32bit=1
       =                                           crc=1        finobt=1, sparse=1, rmapbt=0
       =                                           reflink=1    bigtime=0 inobtcount=0
data      =                                           bsize=4096   blocks=1953168384, imaxpct=5
       =                                           sunit=16     swidth=80 blks
naming    =version 2                  bsize=4096   ascii-ci=0, ftype=1
log        =internal log              bsize=4096   blocks=521728, version=2
       =                                           sectsz=4096   sunit=1 blks, lazy-count=1
realtime  =none                      extsz=4096   blocks=0, rtextents=0
/bin/mkdir -p /var/netwitness/concentrator/sessiondb0
/bin/mount /var/netwitness/concentrator/sessiondb0

```

Execute the `devlist` command to verify partitions created.

Properties for /appliance
 devlist ▼ Parameters:

Message Help

```
devlist: list storage devices
security.roles: appliance.manage
```

/appliance?msg=devlist&force-content-type=text/plain

Output (or command manual help)

```
sdf: vendor=DELL model=RAID size="72.76 TB" used=1
sdd: vendor=DELL model=RAID size="3.49 TB" used=1
sdb: vendor=DELL model=RAID size="43.66 TB" used=1
sdg: vendor=DELL model=RAID size="72.76 TB" used=1
sde: vendor=DELL model=RAID size="2.18 TB" used=1
sdc: vendor=DELL model=RAID size="14.55 TB" used=1
sda: vendor=DELL model=RAID size="29.1 TB" used=1
```

4. Execute the `srvAlloc` command with the following parameters to allocate the space to logdecoder and concentrator services.

This adds storage to the logdecoder service configuration and restarts the service every time it is executed.

```
service=logdecoder volume=logdecoder0 commit=1
```

Properties for /appliance
 srvAlloc ▼ Parameters:

Message Help

```
srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use storag
  volume - <string, {enum-one:The value must be one of the following: concentrator0|endpoint-log-hybrid|hybrid-concentrator|hybrid-logdecoder-
  commit - <bool, optional> commit changes
```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=logdecoder&volume=logdecoder0&commit=1

Output (or command manual help)

```
Set /database/config/packet.dir to /var/netwitness/logdecoder/packetdb/packetdb=20.74 TB;/var/netwitness/logdecoder/packetdb0=69.12 TB
```

This adds storage to the concentrator service configuration and restarts the service every time it is executed.

```
service=concentrator volume=concentrator0 commit=1
```

Properties for /appliance

srvAlloc Parameters:

Message Help

```

srvAlloc: apply volume group storage to a service on this appliance
security.roles: appliance.manage
parameters:
  service - <string, {enum-one:The value must be one of the following: archiver|concentrator|decoder|logdecoder}> service that will use stor
  volume - <string, {enum-one:The value must be one of the following: concentrator0|endpoint-log-hybrid|hybrid-concentrator|hybrid-logdecode
  commit - <bool, optional> commit changes

```

/appliance?msg=srvAlloc&force-content-type=text/plain&service=concentrator&volume=concentrator0&commit=1

Output (or command manual help)

```

Set /database/config/meta.dir to /var/netwitness/concentrator/metadb==12.55 TB;/var/netwitness/concentrator/metadb0==62.21 TB
Set /database/config/session.dir to /var/netwitness/concentrator/sessiondb==1.24 TB;/var/netwitness/concentrator/sessiondb0==6.91 TB

```

Execute the `srvList` command to verify the services allocated on the appliance.

Properties for /appliance

srvList Parameters:

Message Help

```

srvList: list of core services on this appliance
security.roles: appliance.manage

```

/appliance?msg=srvList&force-content-type=text/plain

Output (or command manual help)

```

localhost:56002: type=logdecoder
packet.dir="/var/netwitness/logdecoder/packetdb/packetdb==20.74 TB;/var/netwitness/logdecoder/packetdb0==69.12 TB"
meta.dir="/var/netwitness/logdecoder/metadb==24.03 TB"
session.dir="/var/netwitness/logdecoder/sessiondb==2.49 TB"
index.dir="/var/netwitness/logdecoder/index==1.1 TB"
localhost:56005: type=concentrator
meta.dir="/var/netwitness/concentrator/metadb==12.55 TB;/var/netwitness/concentrator/metadb0==62.21 TB"
session.dir="/var/netwitness/concentrator/sessiondb==1.24 TB;/var/netwitness/concentrator/sessiondb0==6.91 TB"
index.dir="/var/netwitness/concentrator/index==3.32 TB"

```

5. Execute `lsblk` command to list all the raids and partitions in the service.

```

[root@S7EndpointLogHybrid ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0      0  29.1T  0 disk
├─sda1                              8:1      0  29.1T  0 part
│   └─hybrid--logdecoder--meta-decoroot 253:8    0  29.1T  0 lvm  /var/netwitness/logdecoder
sdb                                  8:16     0  43.7T  0 disk
├─sdb1                              8:17     0  43.7T  0 part
│   ├──endpoint--log--hybrid-mongo      253:5    0  21.8T  0 lvm  /var/netwitness/mongo
│   └─endpoint--log--hybrid-packetdb    253:7    0  21.8T  0 lvm  /var/netwitness/logdecoder/packetdb
sdc                                  8:32     0  14.6T  0 disk
├─sdc1                              8:33     0  14.6T  0 part
│   └─hybrid--concentrator-root         253:9    0  14.6T  0 lvm  /var/netwitness/concentrator
sdd                                  8:48     0   3.5T  0 disk
├─sdd1                              8:49     0   3.5T  0 part
│   └─index-index                      253:6    0   3.5T  0 lvm  /var/netwitness/concentrator/index
sde                                  8:64     0   2.2T  0 disk
├─sde1                              8:65     0  256M  0 part /boot/efi
├─sde2                              8:66     0    1G  0 part /boot
├─sde3                              8:67     0   2.2T  0 part
│   ├──netwitness_vg00-root             253:0    0   30G  0 lvm  /
│   ├──netwitness_vg00-swap             253:1    0    4G  0 lvm  [SWAP]
│   ├──netwitness_vg00-nwhome           253:2    0   2.1T  0 lvm  /var/netwitness
│   ├──netwitness_vg00-varlog           253:3    0   10G  0 lvm  /var/log
│   └─netwitness_vg00-usrhome           253:4    0   10G  0 lvm  /home
sdf                                  8:80     0  72.8T  0 disk
├─sdf1                              8:81     0  72.8T  0 part
│   └─logdecoder0-packetdb              253:10   0  72.8T  0 lvm  /var/netwitness/logdecoder/packetdb0
sdg                                  8:96     0  72.8T  0 disk
├─sdg1                              8:97     0  72.8T  0 part
│   ├──concentrator0-sessiondb          253:11   0   7.3T  0 lvm  /var/netwitness/concentrator/sessiondb0
│   └─concentrator0-metadb              253:12   0  65.5T  0 lvm  /var/netwitness/concentrator/metadb0

```