



NVIDIA System Management User Guide

Release 20.09

NVIDIA

Sep 23, 2023

Contents

1	Introduction	1
1.1	Verifying NVSM API Services	2
1.2	Configurable DSHM Features	3
1.2.1	Health Monitor Alerts	3
1.2.2	DSHM Alert List	3
1.2.2.1	Recommended Actions	9
1.2.3	Health Monitor Policies	10
1.3	Verifying the Installation	10
1.3.1	Verifying DSHM Services	10
1.3.2	Verifying NVSM Services	11
2	Release Notes	13
2.1	NVSM 20.09.20 Release	13
2.1.1	Bug Fixes	13
2.1.2	Known Issues	13
3	Using the NVSM CLI	15
3.1	Using the NVSM CLI Interactively	15
3.2	Using the NVSM CLI Non-Interactively	16
3.3	Getting Help	17
3.3.1	nvsm “man” Page	17
3.3.2	nvsm –help Flag	17
3.3.3	Help for NVSM CLI Commands	17
3.4	Examining System Health	18
3.4.1	List of Basic Commands	18
3.4.2	Show Health	20
3.4.3	Dump Health	20
3.4.4	Show Storage	21
3.4.4.1	Show Storage Alerts	22
3.4.4.2	Show Storage Drives	23
3.4.4.3	Show Storage Volumes	24
3.4.5	Show GPUs	25
3.4.5.1	Showing Individual GPUs	26
3.4.5.2	Identifying GPU Health Incidents	27
3.4.6	Show Processors	28
3.4.6.1	Show Processor Alerts	29
3.4.7	Show Memory	30
3.4.7.1	Show Memory Alerts	31
3.4.8	Show Fans and Temperature	32
3.4.8.1	Show Thermal Alerts	33
3.4.8.2	Show Fans	34
3.4.8.3	Show Temperatures	35
3.4.9	Show Power Supplies	36

3.4.9.1	Show Power Alerts	37
3.5	Show Network Adapters	38
3.5.1	Show Network Ports	38
3.5.2	Show Network Device Functions	39
3.5.3	Show Network Interfaces	39
3.6	System Monitoring Configuration	39
3.6.1	Configuring Email Alerts	39
3.6.2	Understanding System Monitoring Policies	40
3.6.2.1	Global Monitoring Policy	41
3.6.2.2	Memory Monitoring Policy	42
3.6.2.3	Processor Monitoring Policy	43
3.6.2.4	Storage Monitoring Policy	43
3.6.2.5	Thermal Monitoring Policy	44
3.6.2.6	Power Monitoring Policy	45
3.6.2.7	PCIe Monitoring Policy	45
3.6.2.8	GPU Monitoring Policy	46
3.6.2.9	Network Adaptor Monitoring Policies	47
3.7	Performing System Management Tasks	49
3.7.1	Rebuilding a RAID 1 Array	49
3.7.1.1	Viewing a Healthy RAID Volume	49
3.7.1.2	Viewing a Degraded RAID Volume	49
3.7.1.3	Rebuilding the RAID 1 Volume	50
3.7.2	Setting MaxQ/MaxP on DGX-2 Systems	51
3.7.2.1	MaxQ	51
3.7.2.2	MaxP	51
3.7.3	Configuring Support for Custom Drive Partitioning	51
4	Configuring NVSM Security	53
4.1	Overview of NVSM Security	53
4.2	What You Need to Configure NVSM Security	53
4.3	How to Configure NVSM Security	54
5	NVSM Call Home	55
5.1	NVSM Call Home Overview	55
5.1.1	Policy-enabled “automatic” Mode	56
5.1.2	Policy-enabled Offline Mode	56
5.1.3	On-demand Mode	56
5.2	Using NVSM Call Home	57
5.2.1	Prerequisites for Using NVSM Call Home in Automatic or On-Demand Mode	57
5.2.1.1	Enabling Ports	57
5.2.1.2	Enabling Access	57
5.2.1.3	Validating NVSM Call Home Readiness	57
5.2.2	Using NVSM Call Home in Automatic Mode	58
5.2.3	Using NVSM On-Demand Mode	59
5.2.4	Using NVSM Offline Call Home	60
6	Third-Party Licenses	63
6.1	mattn/go-sqlite3	63
7	Notices	65
7.1	Notice	65
7.2	Trademarks	66
7.3	Copyright	66


Chapter 1. Introduction

NVIDIA® System Management (NVSM) is a software framework for monitoring NVIDIA DGX™ nodes in a data center.

- ▶ For DGX Servers, it includes active health monitoring, system alerts, and log generation.
- ▶ For DGX Station, is it limited to using the CLI to check the health of the system and obtain diagnostic information.

NVSM can be used as a standalone utility from the command line by system administrators.

The following is a high level diagram of the NVSM framework, showing the NVSM API services at the heart of the framework, the DGX System Health Monitors (DSHM) responsible for monitoring the health of key system components, and the NVSM CLI for user control.



Note: DSHM functionality is not supported on DGX Station.

1.1. Verifying NVSM API Services

`nvsm_api_gateway` is part of the DGX OS image and is launched by `systemd` when DGX boots. The following are the services running under NVSM-APIS.

`nvsm_api_gateway-plugin-environment`

`nvsm_api_gateway-mqtt`

`nvsm_api_gateway-plugin-memory`

`nvsm_api_gateway`

`nvsm_api_gateway-selwatcher`

You can verify if each NVSM-APIS service is up and running using the `'systemctl'` command. For example, the following command verifies the memory service.

```
$ sudo systemctl status nvsm_api_gateway-plugin-memory
```

You can also view all the NVSM-APIS services and their status with the following command.

```
$ sudo systemctl status -all nvsm_api_gateway*
```

1.2. Configurable DSHM Features

DSHM contains the following features that you can configure using the NVSM CLI:

- ▶ Health Monitor Alerts
- ▶ Health Monitor Policies

1.2.1. Health Monitor Alerts

Alerts are events of significance that require attention. When a health monitor detects such an event in the subsystem that it monitors, it generates an alert to inform the user. The default behavior is to log the alerts in persistent storage as well as to send an E-mail notification to registered users. Refer to the section [Using the NVSM CLI](#) for details about configuring users for receiving alert E-mail notifications.

Each alert has a 'state'. An active alert can be in a 'critical' or 'warning' state. Here, 'critical' implies an event that needs immediate action, and 'warning' implies an event that needs user attention. When the alerting condition is removed, the alert state changes to 'cleared'. Details of how to view the generated alerts recorded in the database are available in the section [Using the NVSM CLI](#).

1.2.2. DSHM Alert List

The following table describes each DSHM alert ID.

Message and details	Alert ID	Component ID	Severity	Recommended Action
Unsupported drive configuration in {} (Drive(s) missing or Drive placement not in recommended slots. Drives found at these slots: {}. Supported Configurations are: {})	NV-DRIVE-01	<drive slot>	Critical	See Recommended Action A below.
System entered degraded mode, drive in {} is reporting an error. (Media errors detected in drive.)	NV-DRIVE-02	<drive slot>	Warning	See Recommended Action A below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
System entered degraded mode, drive in {} is reporting an error. (IO errors detected in drive.)	NV-DRIVE-03	<drive slot>	Warning	See Recommended Action A below.
System entered degraded mode, drive in {} is reporting an error. (NVMe controller failure detected in drive.)	NV-DRIVE-04	<drive slot>	Critical	See Recommended Action A below.
System entered degraded mode, drive in {} is reporting an error. (Available capacity percentage below critical threshold for drive.)	NV-DRIVE-05	<drive slot>	Critical	See Recommended Action A below.
System entered degraded mode, drive in {} is reporting an error.	NV-DRIVE-06	<drive slot>	Critical	See Recommended Action A below.
System entered degraded mode, drive in {} is reporting an error. (Drive used percentage above critical threshold for drive.)	NV-DRIVE-07	<drive slot>	Warning	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. ({}RAID-0/RAID-1/Partition} is corrupted.)	NV-VOL-01	NA	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Power supply module has failed.)	NV-PSU-01	<PSU#> where # is the PSU number.	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (operating temperature exceeds the thermal specifications of the component.)	NV-PSU-02	<PSU#> where # is the PSU number.	Warning	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Input to the PSU is missing)	NV-PSU-03	<PSU#> where # is the PSU number.	Critical	See Recommended Action A below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
System entered degraded mode, {} is reporting an error. (Input voltage is out of range for the Power Supply Module) (Input voltage is out of range for the Power Supply Module)	NV-PSU-04	<PSU#> where # is the PSU number.	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (PSU is missing)	NV-PSU-05	<PSU#> where # is the PSU number.	Warning	See Recommended Action A below.
Failures in Power supply modules have been detected. (System is operating in degraded performance mode.)	NV-PSU-06		Warning	Rectify the issues observed on the PSUs. Then see Recommended Action A below.
Failures in Power supply modules have been detected. (System is in power failed state)	NV-PSU-07		Critical	Rectify the issues observed on the PSUs. Then see Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Operating temperature exceeds the thermal specifications of the component.)	NV-PDB-01	<PDB#> where # is the PDB number	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Fan speed reading has fallen below the expected speed setting.)	NV-FAN-01	<FAN#_F> or <FAN#_R> where # is the fan module number. F is for front fan. R is for rear fan.	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Fan readings are inaccessible.)	NV-FAN-02	<FAN#_F> or <FAN#_R> where # is the fan module number. F is for front fan. R is for rear fan.	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (An unrecoverable CPU Internal error has occurred.)	NV-CPU-01	<CPU#> where # is the CPU socket number (CPU0 or CPU1)	Critical	See Recommended Action A below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
System entered degraded mode, {} is reporting an error. (CPU Thermtrip has occurred, processor socket temperature exceeded the thermal specifications of the component.)	NV-CPU-02	<CPU#> where # is the CPU socket number (CPU0 or CPU1)	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Processor socket temperature exceeded the thermal specifications of the component.)	NV-CPU-03		Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Processor socket temperature exceeded the thermal specifications of the component.)	NV-CPU-04		Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Uncorrectable error is reported).	NV-DIMM-01	<CPU#_DIMM_@\$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Correctable errors reported exceeds the configured threshold.)	NV-DIMM-02	<CPU#_DIMM_@\$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Warning	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (Unrecoverable error is observed on the DIMM, specific details of the error are unavailable.)	NV-DIMM-03	<CPU#_DIMM_@\$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting an error. (DIMM presence is not expected in this slot, please verify the DIMM details.)	NV-DIMM-04			See Recommended Action A below.
System entered degraded mode, GPU is reporting an error (Critical error has been reported by the GPU.)	NV-GPU-01		Critical	See Recommended Action A below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
GPU{} power Limits are not configured correctly (Expected limits (Power: 200000W, CLock: 1597MHz), Actual limits (Power: 200000W, Clock: 1163MHz).)	NV-GPU-02		Critical	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Link speed degradation observed between { BDF1, BDF2}, expected link speed is {} actual link speed is {})	NV-PCI-01		Warning	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Link width degradation observed between {BDF1, BDF2},, expected link width is {} actual link width is {})	NV-PCI-02		Warning	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Correctable errors reported on {BDF}).)	NV-PCI-03		Warning	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (UnCorrectable errors reported on {BDF})	NV-PCI-04		Critical	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Device is missing on {BDF})	NV-PCI-05		Critical	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Device Error Reporting is disabled on {BDF} for {})	NV-PCI-06		Critical	See Recommended Action A below.
System entered degraded mode, {ID} is reporting an error. (Device is disabled on {BDF})	NV-PCI-07		Critical	See Recommended Action A below.
System entered degraded mode, Storge Controller {} is reporting an error. (Storage controller state not optimal.)	NV-CONTROLLER-01		Warning	See Recommended Action A below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
System entered degraded mode, Storage Controller {} is reporting an error. (Storage controller port connectivity issues detected.)	NV-CONTROLLER-02		Warning	See Recommended Action A below.
System entered degraded mode, Storage Controller {} is reporting an error. (Storage controller phy link speed not set to optimal.)	NV-CONTROLLER-03		Warning	See Recommended Action A below.
Link error on {}. (Network Link is down)	NV-NET-01		Warning	See Recommended Action B below.
Network traffic errors observed on {}. (Rx collision rate of {}, has crossed threshold value of {} on {}network port.)	NV-NET-02		Warning	See Recommended Action B below.
Network traffic errors observed on {}. (Tx collision rate of {}, has crossed threshold value of {} on {}network port.)	NV-NET-03		Warning	See Recommended Action B below.
Network traffic errors observed on {}. (CRC error rate of {}, has crossed threshold value of {} on {}network port.)	NV-NET-04		Critical	See Recommended Action B below.
{} is reporting an error. ({}Network port is disabled.)	NV-NET-05		Critical	See Recommended Action B below.
Ethernet interface error on port {}. ({}Ethernet health check failing with Online NVRAM test failure.)	NV-ETH-01		Critical	See Recommended Action B below.
Ethernet interface configuration error on {}. (MAC address is missing on the Ethernet interface of {}.)	NV-ETH-02		Critical	See Recommended Action B below.
IB driver error. (HCA self test reports IB driver initialization failure.)	NV-IB-01		Critical	See Recommended Action C below.

continues on next page

Table 1 – continued from previous page

Message and details	Alert ID	Component ID	Severity	Recommended Action
Counter errors on IB port {}. ({}HCA self test on IB port reports counter error.)	NV-IB-02		Critical	See Recommended Action B below.
Configuration error on IB port {}. (GUID is missing on {}HCA.)	NV-IB-03		Critical	See Recommended Action D below.
System entered degraded mode, {} is reporting a fatal error (Critical error has been reported by the NVSwitch Id {} with SXID error {})	NV-NVSWITCH-01		Critical	See Recommended Action A below.
System entered degraded mode, {} is reporting a non fatal error (Critical error has been reported by the NVSwitch Id {} with SXID error {})	NV-NVSWITCH-02		Warning	See Recommended Action A below.

1.2.2.1 Recommended Actions

(A)

1. Run 'sudo nvsm dump health'
2. Open a case with NVIDIA Enterprise Support at this address <https://nvid.nvidia.com/dashboard/>
3. Attach this notification and the nvsysinfo log file from /tmp/nvsm-health- <hostname>-<timestamp>.tar.xz

(B)

1. Check the physical link connection
2. Open a case with NVIDIA Enterprise Support at <https://nvid.nvidia.com/dashboard/>

(C)

1. Check OFED installation troubleshooting
2. Open a case with NVIDIA Enterprise Support at this address <https://nvid.nvidia.com/dashboard/>

(D)

1. Check the status of the Subnet Manager
2. Open a case with NVIDIA Enterprise Support at this address <https://nvid.nvidia.com/dashboard/>

1.2.3. Health Monitor Policies

Users can tune certain aspects of health monitor behavior using health monitor policies. This includes details such as email related configuration for alert notification, selectively disabling devices to be monitored, etc. Details of the supported policies and how to configure them using the CLI are provided in the section *Using the NVSM CLI*.

1.3. Verifying the Installation

Before using NVSM, you can verify the installation to make sure all the services are present.

1.3.1. Verifying DSHM Services

Health monitors are part of the DGX OS image and launched by systemd when DGX boots. You can verify if all the DSHM services are up and running using the `systemctl` command. Below is an example of verifying whether the environmental DSHM service is functional.

```
$ sudo systemctl status nvsm-plugin-monitor-environment
```

```
nvsm-plugin-monitor-environment.service - Environmental DSHM service.
  Loaded: loaded (/user/lib/systemd/sysest/nvsm-plugin-monitor-environment.service;
  →enabled; vendor preset; enabled)
  Active: active (running) since Tues 2018-09-11 15:12:06 PDT: 3h 1min ago
  Main PID: 2540 (plugin-monitor-environment)
  Tasks: 1 (limit 12287)
  CGroup: /system.slice/nvsm-plugin-monitor-environment.service
          |_2540 /user/bin/python /usr/bin/plugin-monitor-environment
```

Other modules can be verified using similar commands:

To verify the storage module:

```
$ sudo systemctl status nvsm-plugin-monitor-storage
```

To verify the system module:

```
$ sudo systemctl status nvsm-plugin-monitor-system
```

To verify the environment module

```
$ sudo systemctl status nvsm-plugin-monitor-environment
```

1.3.2. Verifying NVSM Services

NVSM is part of the DGX OS image and is launched by systemd when DGX boots. The following are the services running under NVSM.

`nvsm-plugin-environment`

`nvsm-mqtt`

`nvms-plugin-memory`

`nvsm`

`nvsm-selwatcher`

You can verify if each NVSM service is up and running using the 'systemctl' command. For example, the following command verifies the memory service.

```
$ sudo systemctl status nvsm-plugin-memory
```

You can also view all the NVSM services and their status with the following command.

```
$ sudo systemctl status -all nvsm*
```

Chapter 2. Release Notes

2.1. NVSM 20.09.20 Release

NVIDIA® System Management (NVSM) is a software framework for monitoring NVIDIA DGX™ nodes in a data center. It includes active health monitoring, system alerts, and log generation. See the [NVSM User Guide](#) for more information.

NVSM Version 20.09.20 is released in March 2021.

NVSM Version 20.09.20 is provided as part of the DGX OS 5.0 software.

2.1.1. Bug Fixes

- ▶ [DGX A100]: On systems with encrypted rootfs, the RAID 1 array cannot be rebuilt using NVSM.

2.1.2. Known Issues

- ▶ [DGX A100]: `nvsm dump health` can take up to 25 minutes to complete, depending on the size of the log file.
- ▶ [DGX A100]: NVSM Enumerates NVSwitches as 8-13 Instead of 0-5
- ▶ NVSM fails to show CPU information on non-English locales
- ▶ [DGX-2]: NVSM does not detect downgraded GPU PCIe Link
- ▶ [DGX-2 KVM]: `nvidia-vm vmshow` command does not work for running VMs.
- ▶ `configure_raid_array.py` script cannot recreate RAID array after re-inserting a known good SSD.

Chapter 3. Using the NVSM CLI

NVIDIA DGX-2 servers running DGX OS version 4.0.1 or later should come with NVSM pre-installed.

NVSM CLI communicates with the privileged NVSM API server, so NVSM CLI requires superuser privileges to run. All examples given in this guide are prefixed with the “sudo” command.

3.1. Using the NVSM CLI Interactively

Starting an interactive session

The command “sudo nvsm” will start an NVSM CLI interactive session.

```
user@dgx-2:~$ sudo nvsm
[sudo] password for user:
nvsm->
```

Once at the “nvsm->” prompt, the user can enter NVSM CLI commands to view and manage the DGX system.

Example command

One such command is “show fans”, which prints the state of all fans known to NVSM.

```
nvsm-> show fans
/chassis/localhost/thermal/fans/FAN10_F
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = FAN10_F
  MemberId = 19
  ReadingUnits = RPM
  LowerThresholdNonCritical = 5046.000
  Reading = 9802 RPM
  LowerThresholdCritical = 3596.000
  ...
/chassis/localhost/thermal/fans/PDB_FAN4
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = PDB_FAN4
  MemberId = 23
  ReadingUnits = RPM
  LowerThresholdNonCritical = 11900.000
```

(continues on next page)

(continued from previous page)

```

    Reading = 14076 RPM
    LowerThresholdCritical = 10744.000
nvsm->

```

Leaving an interactive session

To leave the NVSM CLI interactive session, use the “exit” command.

```

nvsm-> exit
user@dgx2:~$

```

3.2. Using the NVSM CLI Non-Interactively

Any NVSM CLI command can be invoked from the system shell, without starting an NVSM CLI interactive session. To do this, simply append the desired NVSM CLI command to the “sudo nvsm” command. The “show fans” command given above can be invoked directly from the system shell as follows.

```

user@dgx2:~$ sudo nvsm show fans
/chassis/localhost/thermal/fans/FAN10_F
Properties:
    Status_State = Enabled
    Status_Health = OK
    Name = FAN10_F
    MemberId = 19
    ReadingUnits = RPM
    LowerThresholdNonCritical = 5046.000
    Reading = 9802 RPM
    LowerThresholdCritical = 3596.000
...
/chassis/localhost/thermal/fans/PDB_FAN4
Properties:
    Status_State = Enabled
    Status_Health = OK
    Name = PDB_FAN4
    MemberId = 23
    ReadingUnits = RPM
    LowerThresholdNonCritical = 11900.000
    Reading = 14076 RPM
    LowerThresholdCritical = 10744.000
user@dgx2:~$

```

The output of some NVSM commands can be too large to fit on one screen, it is sometimes useful to pipe this output to a paging utility such as “less”.

```

user@dgx2:~$ sudo nvsm show fans | less

```

Throughout this chapter, examples are given for both interactive and non-interactive NVSM CLI use cases. Note that these interactive and non-interactive examples are interchangeable.

3.3. Getting Help

Apart from the NVSM CLI User Guide (this document), there are many sources for finding additional help for NVSM CLI and the related NVSM tools.

3.3.1. nvsm “man” Page

A man page for NVSM CLI is included on DGX systems with NVSM installed. The user can view this man page by invoking the “man nvsm” command.

```
user@dgx2:~$ man nvsm
```

3.3.2. nvsm -help Flag

By passing the -help flag, the nvsm command itself will print a short description of the command line arguments it recognizes. These arguments affect the behavior of the NVSM CLI interactive session, such as inclusion of color or log messages.

```
user@dgx2:~$ nvsm --help
usage: nvsm [-h] [--color WHEN] [-i] [--] [<command>...]
NVIDIA System Management interface
optional arguments:
  -h, --help            show this help message and exit
  --color WHEN          Control colorization of output. Possible
                        values for WHEN are "always", "never", or
                        "auto". Default value is "auto".
  -i, --interactive     When this option is given, run in
                        interactive mode. The default is
                        automatic.
  --log-level {debug,info,warning,error,critical}
                        Set the output logging level. Default is
                        'warning'.
```

3.3.3. Help for NVSM CLI Commands

Each NVSM command within the NVSM CLI interactive session, such as show, set, and exit, recognizes a “-help” flag that describes the NVSM command and its arguments.

```
user@dgx2:~$ sudo nvsm
nvsm-> exit -help
usage: exit [-help]

Leave the NVSM shell.

optional arguments:
  -help, -h  show this help message and exit
```

3.4. Examining System Health

The most basic functionality of NVSM CLI is examination of system state. NVSM CLI provides a “show” command for this purpose.

Because NVSM CLI is modeled after the SMASH CLP, the output of the NVSM CLI “show” command should be familiar to users of BMC command line interfaces.

3.4.1. List of Basic Commands

The following table lists the basic commands (primarily “show”). Detailed use of these commands are explained in subsequent sections of the document.

Note: On DGX Station, the following are the only commands supported.

- ▶ `nvsm show health`
- ▶ `nvsm dump health`

Global Commands	Descriptions
<code>\$ sudo nvsm show alerts</code>	Displays warnings and critical alerts for all subsystems
<code>\$ sudo nvsm show policy</code>	Displays alert policies for subsystems
<code>\$ sudo nvsm show versions</code>	Displays system version properties

Health Commands	Descriptions
<code>\$ sudo nvsm show health</code>	Displays overall system health
<code>\$ sudo nvsm dump health</code>	Generates a health report file

Storage Commands	Descriptions
<code>\$ sudo nvsm show storage</code>	Displays all storage-related information
<code>\$ sudo nvsm show drives</code>	Displays the storage drives
<code>\$ sudo nvsm show volumes</code>	Displays the storage volumes

GPU Commands	Descriptions
\$ sudo nvsm show gpus	Displays informatin for all GPUs in the system.

Processor Commands	Descriptions
\$ sudo nvsm show processors	Displays information for all CPUs in the system
\$ sudo nvsm show cpus	Alias for “show processors”

Memory Commands	Descriptions
\$ sudo nvsm show memory	Displays information for all installed DIMMs
\$ sudo nvsm show dimms	Alias for “show memory”

Thermal Commands	Descriptions
\$ sudo nvsm show fans	Displays information for all the fans in the system.
\$ sudo nvsm show temperatures	Displays temperature information for all sensors in the system
\$ sudo nvsm show temps	Alias for “show temperatures”

Network Commands	Descriptions
\$ sudo nvsm show networkadapters	Displays information for the physical network adapters
\$ sudo nvsm show networkinterfaces	Displays information for the logical network interfaces
\$ sudo nvsm show networkports	Displays information for the network ports of a given network adapter
\$ sudo nvsm show networkdevicefunctions	Displays information for the PCIe functions for a given network adapter

Power Commands	Descriptions
\$ sudo nvsm show power	Displays information for all power supply units (PSUs) in the system.
\$ sudo nvsm show powermode	Display the current system power mode
\$ sudo nvsm show psus	Alias for “show power”

NVSwitch Commands	Descriptions
\$ sudo nvsm show nvswitches	Displays information for all the NVSwitch interconnects in the system.

Firmware Commands	Descriptions
\$ sudo nvsm show firmware	Guides you through the steps of selecting a firmware update container on your local DGX system, and running it to show the firmware versions installed on the system. This requires that you have already loaded the container onto the DGX system.
\$ sudo nvsm update firmware	Guides you through the steps of selecting a firmware update container on your local DGX system, and running it to update the firmware on the system. This requires that you have already loaded the container onto the DGX system.

3.4.2. Show Health

The “show health” command can be used to quickly assess overall system health.

```
user@dgx-2:~$ sudo nvsm show health
```

Example output:

```
...
Checks
-----Verify installed DIMM memory sticks.....
HealthyNumber of logical CPU cores [96].....
HealthyGPU link speed [0000:39:00.0][8GT/s].....
HealthyGPU link width [0000:39:00.0][x16].....
Healthy
...
Health Summary
-----
205 out of 205 checks are Healthy
Overall system status is Healthy
```

If any system health problems are found, this will be reflected in the health summary at the bottom of the “show health” output”. Detailed information on health checks performed will appear above.

3.4.3. Dump Health

The “dump health” command produces a health report file suitable for attaching to support tickets.

```
user@dgx-2:~$ sudo nvsm dump health
```

Example output:

```
Writing output to /tmp/nvsm-health-dgx-1-20180907085048.tar.xzDone.
```

The file produced by “dump health” is a familiar compressed tar archive, and its contents can be examined by using the “tar” command as shown in the following example.

```
user@dgx-2:~$ cd /tmp
user@dgx-2:/tmp$ sudo tar xlf nvsm-health-dgx-1-20180907085048.tar.xz
user@dgx-2:/tmp$ sudo ls ./nvsm-health-dgx-1-20180907085048
date          java          nvsysinfo_commands  sos_reports
df            last          nvsysinfo_log.txt   sos_strings
dmidecode     lib           proc                sys
etc           lsb-release  ps                  uname
free         lsmod        pstree              uptime
hostname     lsof         route              usr
initctl      lspci        run                 var
installed-debs mount        sos_commands       version.txt
ip_addr      netstat      sos_logs            vgdisplay
```

3.4.4. Show Storage

NVSM CLI provides a “show storage” command to view all storage-related information. This command can be invoked from the command line as follows.

```
user@dgx-2:~$ sudo nvsm show storage
```

Alternatively, the “show drives” and “show volumes” NVSM commands will show the storage drives or storage volumes respectively.

```
user@dgx-2:~$ sudo nvsm show drives
...
user@dgx-2:~$ sudo nvsm show volumes
...
```

Within an NVSM CLI interactive session, the CLI targets related to storage are located under the /systems/localhost/storage/1 target.

```
user@dgx2:~$ sudo nvsm
nvsm-> cd /systems/localhost/storage/
nvsm(/systems/localhost/storage/)-> show
```

Example output:

```
/systems/localhost/storage/
Properties:
  DriveCount = 10
  Volumes = [ md0, md1, nvme0n1p1, nvme1n1p1 ]
Targets:
  alerts
  drives
  policy
  volumes
Verbs:
  cd
  show
```

3.4.4.1 Show Storage Alerts

Storage alerts are generated when the DSHM monitoring daemon detects a storage-related problem and attempts to alert the user (via email or otherwise). Past storage alerts can be viewed within an NVSM CLI interactive session under the `/systems/localhost/storage/1/alerts` target.

```
user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/storage/alerts
nvsm(/systems/localhost/storage/alerts)-> show
```

Example output:

```
/systems/localhost/storage/alerts
Targets:
  alert0
  alert1
Verbs:
  cd
  show
```

In this example listing, there appear to be two storage alerts associated with this system. The contents of these alerts can be viewed with the “show” command.

For example:

```
nvsm(/systems/localhost/storage/alerts)-> show alert1
```

```
/systems/localhost/storage/alerts/alert1
Properties:
  system_name = dgx-2
  message_details = EFI System Partition 1 is corrupted
nvme0n1p1
  component_id = nvme0n1p1
  description = Storage sub-system is reporting an error
  event_time = 2018-07-14 12:51:19
  recommended_action =
    1. Please run nvsysinfo
    2. Please open a case with NVIDIA Enterprise Support at this address https://
↪nvid.nvidia.com/enterpriselogin
    3. Attach this notification and the nvsysinfo log file from /tmp/nvsysinfo-
↪XYZ*
  alert_id = NV-VOL-03
  system_serial = productserial
  message = System entered degraded mode, storage sub-system is reporting an error
  severity = Warning
Verbs:
  cd
  show
```

The message seen in this alert suggests a possible EFI partition corruption, which is an error condition that might adversely affect this system’s ability to boot. Note that the text seen here reflects the exact message that the user would have seen when this alert was generated.

Possible categories for storage alerts are given in the table below.

Alert ID	Severity	Details
NV-DRIVE-01	Critical	Drive missing
NV-DRIVE-02	Warning	Media errors detected in drive
NV-DRIVE-03	Warning	IO errors detected in drive
NV-DRIVE-04	Critical	NVMe controller failure detected in drive
NV-DRIVE-05	Warning	Available spare block percentage is below critical threshold of ten percent
NV-DRIVE-06	Warning	NVM subsystem usage exceeded ninety percent
NV-DRIVE-07	Warning	System has unsupported drive
NV-VOL-01	Critical	RAID-0 corruption observed
NV-VOL-02	Critical	RAID-1 corruption observed
NV-VOL-03	Warning	EFI System Partition 1 corruption observed
NV-VOL-04	Warning	EFI System Partition 2 corruption observed

3.4.4.2 Show Storage Drives

Within an NVSM CLI interactive session, each storage drive on the system is represented by a target under the `/systems/localhost/storage/drives` target. A listing of drives can be obtained as follows.

```
user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/storage/drives
nvsm(/systems/localhost/storage/drives)-> show
```

Example output:

```
/systems/localhost/storage/drives
Targets:
  nvme0n1
  nvme1n1
  nvme2n1
  nvme3n1
  nvme4n1
  nvme5n1
  nvme6n1
  nvme7n1
  nvme8n1
  nvme9n1
Verbs:
  cd
  show
```

Details for any particular drive can be viewed with the “show” command.

For example:

```
nvsm(/systems/localhost/storage/drives)-> show nvme2n1
```

```

/systems/localhost/storage/drives/nvme2n1
Properties:
  Capacity = 3840755982336
  BlockSizeBytes = 7501476528
  SerialNumber = 18141C244707
  PartNumber = N/A
  Model = Micron_9200_MTFDHAL3T8TCT
  Revision = 100007C0
  Manufacturer = Micron Technology Inc
  Status_State = Enabled
  Status_Health = OK
  Name = Non-Volatile Memory Express
  MediaType = SSD
  IndicatorLED = N/A
  EncryptionStatus = N/A
  HotSpareType = N/A
  Protocol = NVMe
  NegotiatedSpeedsGbs = 0
  Id = 2
Verbs:
  cd
  show

```

3.4.4.3 Show Storage Volumes

Within an NVSM CLI interactive session, each storage volume on the system is represented by a target under the `/systems/localhost/storage/volumes` target. A listing of volumes can be obtained as follows.

```

user@dgx-2:~$ sudo nvsm

nvsmnvsm-> cd /systems/localhost/storage/volumes
nvsm(/systems/localhost/storage/volumes)-> show

```

Example output:

```

/systems/localhost/storage/volumes
Targets:
  md0
  md1
  nvme0n1p1
  nvme1n1p1
Verbs:
  cd
  show

```

Details for any particular volume can be viewed with the “show” command.

For example:

```

nvsm(/systems/localhost/storage/volumes)-> show md0

```

```

/systems/localhost/storage/volumes/md0P
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = md0

```

(continues on next page)

(continued from previous page)

```

Encrypted = False
VolumeType = RAID-1
Drives = [ nvme0n1, nvme1n1 ]
CapacityBytes = 893.6G
Id = md0
Verbs:
  cd
  show

```

3.4.5. Show GPUs

Information for all GPUs installed on the system can be viewed invoking the “show gpus” command as follows.

```
user@dgx-2:~$ sudo nvsm show gpus
```

Within an NVSM CLI interactive session, the same information can be accessed under the /systems/localhost/gpus CLI target.

```

user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/gpus
nvsm(/systems/localhost/gpus)-> show

```

Example output:

```

/systems/localhost/gpus
Targets:
  0
  1
  2
  3
  4
  5
  6
  7
  8
  9
 10
 11
 12
 13
 14
 15
Verbs:
  cd
  show

```

Details for any particular GPU can also be viewed with the “show” command.

For example:

```

nvsm(/systems/localhost/gpus)-> show 6
/systems/localhost/gpus/6
Properties:

```

(continues on next page)

(continued from previous page)

```

Inventory_ModelName = Tesla V100-SXM3-32GB
Inventory_UUID = GPU-4c653056-0d6e-df7d-19c0-4663d6745b97
Inventory_SerialNumber = 0332318503073
Inventory_PCIEDeviceId = 1DB810DE
Inventory_PCIESubSystemId = 12AB10DE
Inventory_BrandName = Tesla
Inventory_PartNumber = 699-2G504-0200-000
Verbs:
  cd
  show

```

3.4.5.1 Showing Individual GPUs

Details for any particular GPU can also be viewed with the “show” command.

For example:

```

nvsm(/systems/localhost/gpus)-> show GPU6
/systems/localhost/gpus/GPU6
Properties:
  Inventory_ModelName = Tesla V100-SXM3-32GB
  Inventory_UUID = GPU-4c653056-0d6e-df7d-19c0-4663d6745b97
  Inventory_SerialNumber = 0332318503073
  Inventory_PCIEDeviceId = 1DB810DE
  Inventory_PCIESubSystemId = 12AB10DE
  Inventory_BrandName = Tesla
  Inventory_PartNumber = 699-2G504-0200-000
  Specifications_MaxPCIEGen = 3
  Specifications_MaxPCIELinkWidth = 16x
  Specifications_MaxSpeeds_GraphicsClock = 1597 MHz
  Specifications_MaxSpeeds_MemClock = 958 MHz
  Specifications_MaxSpeeds_SMClock = 1597 MHz
  Specifications_MaxSpeeds_VideoClock = 1432 MHz
  Connections_PCIEGen = 3
  Connections_PCIELinkWidth = 16x
  Connections_PCIELocation = 00000000:34:00.0
  Power_PowerDraw = 50.95 W
  Stats_ErrorStats_ECCMode = Enabled
  Stats_FrameBufferMemoryUsage_Free = 32510 MiB
  Stats_FrameBufferMemoryUsage_Total = 32510 MiB
  Stats_FrameBufferMemoryUsage_Used = 0 MiB
  Stats_PCIERxThroughput = 0 KB/s
  Stats_PCIEtxThroughput = 0 KB/s
  Stats_PerformanceState = P0
  Stats_UtilDecoder = 0 %
  Stats_UtilEncoder = 0 %
  Stats_UtilGPU = 0 %
  Stats_UtilMemory = 0 %
  Status_Health = OK
Verbs:
  cd
  show

```

3.4.5.2 Identifying GPU Health Incidents

Explain the benefits of the task, the purpose of the task, who should perform the task, and when to perform the task in 50 words or fewer.

NVSM uses NVIDIA Data Center GPU Manager (DCGM) to continuously monitor GPU health, and reports GPU health issues as “GPU health incidents”. Whenever GPU health incidents are present, NVSM indicates this state in the “Status_HealthRollup” property of the /systems/localhost/gpus CLI target.

“Status_HealthRollup” captures the overall health of all GPUs in the system in a single value. Check the “Status_HealthRollup” property before checking other properties when checking for GPU health incidents.

To check for GPU health incidents, do the following,

1. Display the “Properties” section of GPU health

```
~$ sudo nvsm
nvsm-> cd /systems/localhost/gpus
nvsm(/systems/localhost/gpus)-> show -display properties
```

A system with a GPU-related issue might report the following.

```
Properties:
  Status_HealthRollup = Critical
  Status_Health = OK
```

The “Status_Health = OK” property in this example indicates that NVSM did not find any system-level problems, such as missing drivers or incorrect device file permissions.

The “Status_HealthRollup = Critical” property indicates that at least one GPU in this system is exhibiting a “Critical” health incident.

2. To find this GPU, issue the following command to list the health status for each GPU..

```
~$ sudo nvsm
nvsm-> show -display properties=*health /systems/localhost/gpus/*
```

The GPU with the health incidents will be reported as in the following example for GPU14.

```
/systems/localhost/gpus/GPU14
Properties:
  Status_Health = Critica
```

3. Issue the following command to show the detailed health information for a particular GPU (GPU14 in this example).

```
nvsm-> cd /systems/localhost/gpus
nvsm(/systems/localhost/gpus)-> show -level all GPU14/health
```

The output shows all the incidents involving that particular GPU.

```
/systems/localhost/gpus/GPU14/health
Properties:
  Health = Critical
Targets:
  incident0
Verbs:
```

(continues on next page)

(continued from previous page)

```

    cd
    show/systems/localhost/gpus/GPU2/health/incident0
Properties:
    Message = GPU 14's NVLink link 2 is currently down.
    Health = Critical
    System = NVLink
Verbs:
    cd
    show

```

The output in this example narrows down the scope to a specific incident (or incidents) on a specific GPU. DCGM will monitor for a variety of GPU conditions, so check “Status_HealthRollup ” using NVSM CLI to understand each incident.

3.4.6. Show Processors

Information for all CPUs installed on the system can be viewed using the “show processors” command.

```
user@dgx-2$ sudo nvsm show processors
```

From within an NVSM CLI interactive session, the same information is available under the /systems/localhost/processors target.

```

user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/processors
nvsm(/systems/localhost/processors)-> show

```

Example output:

```

/systems/localhost/processors
Targets:
    CPU0
    CPU1
    alerts
    policy
Verbs:
    cd
    show

```

Details for any particular CPU can be viewed using the “show” command.

For example:

```

nvsm(/systems/localhost/processors)-> show CPU0/systems/localhost/processors/CPU0
Properties:
    Id = CPU0
    InstructionSet = x86-64
    Manufacturer = Intel(R) Corporation
    MaxSpeedMHz = 3600
    Model = Intel(R) Xeon(R) Platinum 8168 CPU @ 2.70GHz
    Name = Central Processor
    ProcessorArchitecture = x86
    ProcessorId_EffectiveFamily = 6
    ProcessorId_EffectiveModel = 85

```

(continues on next page)

(continued from previous page)

```

ProcessorId_IdentificationRegisters = 0xBFEBFBFF00050654
ProcessorId_Step = 4
ProcessorId_VendorId = GenuineIntel
ProcessorType = CPU
Socket = CPU 0
Status_Health = OK
Status_State = Enabled
TotalCores = 24
TotalThreads = 48
Verbs:
  cd
  show

```

3.4.6.1 Show Processor Alerts

Processor alerts are generated when the DSHM monitoring daemon detects a CPU Internal Error (IERR) or Thermal Trip and attempts to alert the user (via email or otherwise). Past processor alerts can be viewed within an NVSM CLI interactive session under the `/systems/localhost/processors/alerts` target.

```

user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/processors/alerts
nvsm(/systems/localhost/processors/alerts)-> show

```

Example output:

```

/systems/localhost/processors/alerts
Targets:
  alert0
  alert1
  alert2
Verbs:
  cd
  show

```

This example listing appears to show three processor alerts associated with this system. The contents of these alerts can be viewed with the “show” command.

For example:

```

nvsm(/systems/localhost/processors/alerts)-> show alert2

```

```

/systems/localhost/processors/alerts/alert2
Properties:
  system_name = xpl-bu-06
  component_id = CPU0
  description = CPU is reporting an error.
  event_time = 2018-07-18T16:42:20.580050
  recommended_action =
    1. Please run nvsysinfo
    2. Please open a case with NVIDIA Enterprise Support at this address https://
    ↪nvid.nvidia.com/enterpriselogin
    3. Attach this notification and the nvsysinfo log file from /tmp/nvsysinfo-XYZ*
  severity = Critical
  alert_id = NV-CPU-02
  system_serial = To be filled by O.E.M.

```

(continues on next page)

(continued from previous page)

```

message = System entered degraded mode, CPU0 is reporting an error.
message_details = CPU Thermtrip has occurred, processor socket temperature
↳exceeded the thermal specifications of the component.
Verbs:
  cd
  show

```

Possible categories for processor alerts are given in the table below.

Alert ID	Severity	Details
NV-CPU-01	Critical	An unrecoverable CPU Internal error has occurred.
NV-CPU-02	Critical	CPU Thermtrip has occurred, processor socket temperature exceeded the thermal specifications of the component.

3.4.7. Show Memory

Information for all system memory (i.e. all DIMMs installed near the CPU, not including GPU memory) can be viewed using the “show memory” command.

```
user@dgx-2:~$ sudo nvsm show memory
```

From within an NVSM CLI interactive session, system memory information is accessible under the /systems/localhost/memory target.

```

lab@xpl-dvt-42:~$ sudo nvsm
nvsm-> cd /systems/localhost/memory
nvsm(/systems/localhost/memory)-> show

```

Example output:

```

/systems/localhost/memory
Targets:
  CPU0_DIMM_A1
  CPU0_DIMM_A2
  CPU0_DIMM_B1
  CPU0_DIMM_B2
  CPU0_DIMM_C1
  CPU0_DIMM_C2
  CPU0_DIMM_D1
  CPU0_DIMM_D2
  CPU0_DIMM_E1
  CPU0_DIMM_E2
  CPU0_DIMM_F1
  CPU0_DIMM_F2
  CPU1_DIMM_G1
  CPU1_DIMM_G2
  CPU1_DIMM_H1
  CPU1_DIMM_H2
  CPU1_DIMM_I1
  CPU1_DIMM_I2

```

(continues on next page)

(continued from previous page)

```

CPU1_DIMM_J1
CPU1_DIMM_J2
CPU1_DIMM_K1
CPU1_DIMM_K2
CPU1_DIMM_L1
CPU1_DIMM_L2
alerts      policy
Verbs:
  cd
  show

```

Details for any particular memory DIMM can be viewed using the “show” command.

For example:

```
nvsm(/systems/localhost/memory)-> show CPU2_DIMM_B1
```

```

/systems/localhost/memory/CPU2_DIMM_B1
Properties:
  CapacityMiB = 65536
  DataWidthBits = 64
  Description = DIMM DDR4 Synchronous
  Id = CPU2_DIMM_B1
  Name = Memory Instance
  OperatingSpeedMhz = 2666
  PartNumber = 72ASS8G72LZ-2G6B2
  SerialNumber = 1CD83000
  Status_Health = OK
  Status_State = Enabled
  VendorId = Micron
Verbs:
  cd
  show

```

3.4.7.1 Show Memory Alerts

On DGX systems with a Baseboard Management Controller (BMC), the BMC will monitor DIMMs for correctable and uncorrectable errors. Whenever memory error counts cross a certain threshold (as determined by SBIOS), a memory alert is generated by the DSHM daemon in an attempt to notify the user (via email or otherwise).

Past memory alerts are accessible from an NVSM CLI interactive session under the `/systems/localhost/memory/alerts` target.

```

user@dgx-2:~$ sudo nvsm
nvsm-> cd /systems/localhost/memory/alerts
nvsm(/systems/localhost/memory/alerts)-> show

```

Example output:

```

/systems/localhost/memory/alerts
Targets:
  alert0
Verbs:

```

(continues on next page)

(continued from previous page)

```
cd
show
```

This example listing appears to show one memory alert associated with this system. The contents of this alert can be viewed with the “show” command.

For example:

```
nvsm(/systems/localhost/memory/alerts)-> show alert0
```

```
/systems/localhost/memory/alerts/alert0
Properties:
  system_name = xpl-bu-06
  component_id = CPU1_DIMM_A2
  description = DIMM is reporting an error.
  event_time = 2018-07-18T16:48:09.906572
  recommended_action =
    1. Please run nvsysinfo
    2. Please open a case with NVIDIA Enterprise Support at this address https://
    ↪ nvid.nvidia.com/enterpriselogin
    3. Attach this notification and the nvsysinfo log file from /tmp/nvsysinfo-XYZ*
  severity = Critical
  alert_id = NV-DIMM-01
  system_serial = To be filled by O.E.M.
  message = System entered degraded mode, CPU1_DIMM_A2 is reporting an error.
  message_details = Uncorrectable error is reported.
Verbs:
  cd
  show
```

Possible categories for memory alerts are given in the table below.

Alert Type	Severity	Details
NV-DIMM-01	Critical	Uncorrectable error is reported.

3.4.8. Show Fans and Temperature

NVSM CLI provides a “show fans” command to display information for each fan on the system.

```
~$ sudo nvsm show fans
```

Likewise, NVSM CLI provides a “show temperatures” command to display temperature information for each temperature sensor known to NVSM.

```
~$ sudo nvsm show temperatures
```

Within an NVSM CLI interactive session, targets related to fans and temperature are located under the /chassis/localhost/thermal target.

```
~$ sudo nvsm
nvsm-> cd /chassis/localhost/thermal
nvsm(/chassis/localhost/thermal)-> show
```

Example output:

```
/chassis/localhost/thermal
Targets:
  alerts
  fans
  policy
  temperatures
Verbs:
  cd
  show
```

3.4.8.1 Show Thermal Alerts

The DSHM daemon monitors fan speed and temperature sensors. When the values of these sensors violate certain threshold criteria, DSHM generates a thermal alert in an attempt to notify the user (via email or otherwise).

Past thermal alerts can be viewed in an NVSM CLI interactive session under the `/chassis/localhost/thermal/alerts` target.

```
user@dgx-2:~$ sudo nvsm
nvsm-> cd /chassis/localhost/thermal/alerts
nvsm(/chassis/localhost/thermal/alerts)-> show
```

Example output:

```
/chassis/localhost/thermal/alerts
Targets:
  alert0
Verbs:
  cd
  show
```

This example listing appears to show one thermal alert associated with this system. The contents of this alert can be viewed with the “show” command.

For example:

```
nvsm(/chassis/localhost/thermal/alerts)-> show alert0
/chassis/localhost/thermal/alerts/alert0
Properties:
  system_name = system-name
  component_id = FAN1_R
  description = Fan Module is reporting an error.
  event_time = 2018-07-12T15:12:22.076814
  recommended_action =
    1. Please run nvsysinfo
    2. Please open a case with NVIDIA Enterprise Support at this address https://
    ↪nvid.nvidia.com/enterpriselogin      3. Attach this notification and the
    ↪nvsysinfo log file from /tmp/nvsysinfo-XYZ*
  severity = Critical
  alert_id = NV-FAN-01
  system_serial = To be filled by O.E.M.
  message = System entered degraded mode, FAN1_R is reporting an error.
  message_details = Fan speed reading has fallen below the expected speed setting.
Verbs:    cd    show
```

From the message in this alert, it appears that one of the rear fans is broken in this system. This is the exact message that the user would have received at the time this alert was generated, assuming alert notifications were enabled.

Possible categories for thermal-related (fan and temperature) alerts are given in the table below.

Alert ID	Severity	Details
NV-FAN-01	Critical	Fan speed reading has fallen below the expected speed setting.
NV-FAN-02	Critical	Fan readings are inaccessible.
NV-PDB-01	Critical	Operating temperature exceeds the thermal specifications of the component.

3.4.8.2 Show Fans

Within an NVSM CLI interactive session, each fan on the system is represented by a target under the `/chassis/localhost/thermal/fans` target. The “show” command can be used to obtain a listing of fans on the system.

```
user@dgx-2:~$ sudo nvsm
```

```
nvsm-> cd /chassis/localhost/thermal/fans
```

```
nvsm(/chassis/localhost/thermal/fans)-> show
```

Example output:

```
/chassis/localhost/thermal/fans
```

Targets:

```
FAN10_F
FAN10_R
FAN1_F
FAN1_R
FAN2_F
FAN2_R
FAN3_F
FAN3_R
FAN4_F
FAN4_R
FAN5_F
FAN5_R
FAN6_F
FAN6_R
FAN7_F
FAN7_R
FAN8_F
FAN8_R
FAN9_F
FAN9_R
PDB_FAN1
PDB_FAN2
PDB_FAN3
```

(continues on next page)

(continued from previous page)

```
PDB_FAN4
Verbs:
  cd
  show
```

Again using the “show” command, the details for any given fan can be obtained as follows.

For example:

```
nvsm(/chassis/localhost/thermal/fans)-> show PDB_FAN2
/chassis/localhost/thermal/fans/PDB_FAN2
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = PDB_FAN2
  MemberId = 21
  ReadingUnits = RPM
  LowerThresholdNonCritical = 11900.000
  Reading = 13804 RPM
  LowerThresholdCritical = 10744.000
Verbs:
  cd
  show
```

3.4.8.3 Show Temperatures

Each temperature sensor known to NVSM is represented as a target under the /chassis/localhost/thermal/temperatures target. A listing of temperature sensors on the system can be obtained using the following commands.

```
nvsm(/chassis/localhost/thermal/temperatures)-> show
```

Example output:

```
/chassis/localhost/thermal/temperatures
Targets:
  PDB1
  PDB2
Verbs:
  cd
  show
```

As with fans, the details for any temperature sensor can be viewed with the “show” command.

For example:

```
nvsm(/chassis/localhost/thermal/temperatures)-> show PDB2
/chassis/localhost/thermal/temperatures/PDB2
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = PDB2
  PhysicalContext = PDB
  MemberId = 1
  ReadingCelsius = 20 degrees C
```

(continues on next page)

(continued from previous page)

```

UpperThresholdNonCritical = 127.000
SensorNumber = 66h
UpperThresholdCritical = 127.000
Verbs:
  cd
  show

```

3.4.9. Show Power Supplies

NVSM CLI provides a “show power” command to display information for all power supplies present on the system.

```
user@dgx-2:~$ sudo nvsm show power
```

From an NVSM CLI interactive session, power supply information can be found under the /chassis/localhost/power target.

```

user@dgx-2:~$ sudo nvsm
nvsm-> cd /chassis/localhost/power
nvsm(/chassis/localhost/power)-> show

```

Example output:

```

/chassis/localhost/power
Targets:
  PSU1
  PSU2
  PSU3
  PSU4
  PSU5
  PSU6
  alerts      policyVerbs:  cd    show

```

Details for any particular power supply can be viewed using the “show” command as follows.

For example:

```
nvsm(/chassis/localhost/power)-> show PSU4
```

```

/chassis/localhost/power/PSU4
Properties:
  Status_State = Present
  Status_Health = OK
  LastPowerOutputWatts = 442
  Name = PSU4
  SerialNumber = DTHTCD18240
  MemberId = 3
  PowerSupplyType = AC
  Model = ECD16010081
  Manufacturer = Delta
Verbs:
  cd
  show

```

3.4.9.1 Show Power Alerts

The DSHM daemon monitors PSU status. When the PSU status is not Ok, DSHM generates a power alert in an attempt to notify the user (via email or otherwise).

Prior power alerts can be viewed under the `/chassis/localhost/power/alerts` target of an NVSM CLI interactive session.

```
user@dgx-2:~$ sudo nvsm
nvsm-> cd /chassis/localhost/power/alerts
nvsm(/chassis/localhost/power/alerts)-> show
```

Example output:

```
/chassis/localhost/power/alerts
Targets:
  alert0
  alert1
  alert2
  alert3
  alert4
Verbs:
  cd
  show
```

This example listing shows a system with five prior power alerts. The details for any one of these alerts can be viewed using the “show” command.

For example:

```
nvsm(/chassis/localhost/power/alerts)-> show alert4
/chassis/localhost/power/alerts/alert4
Properties:
  system_name = system-name
  component_id = PSU4
  description = PSU is reporting an error.
  event_time = 2018-07-18T16:01:27.462005
  recommended_action =
    1. Please run nvsysinfo
    2. Please open a case with NVIDIA Enterprise Support at this address https://
    ↪nvid.nvidia.com/enterpriselogin
    3. Attach this notification and the nvsysinfo log file from /tmp/nvsysinfo-XYZ*
  severity = Warning
  alert_id = NV-PSU-05
  system_serial = To be filled by O.E.M.
  message = System entered degraded mode, PSU4 is reporting an error.
  message_details = PSU is missing
Verbs:
  cd
  show
```

Possible categories for power alerts are given in the table below.

Alert ID	Severity	Details
NV-PSU-01	Critical	Power supply module has failed.
NV-PSU-02	Warning	Detected predictive failure of the Power supply module.
NV-PSU-03	Critical	Input to the Power supply module is missing.
NV-PSU-04	Critical	Input voltage is out of range for the Power Supply Module.
NV-PSU-05	Warning	PSU is missing

3.5. Show Network Adapters

NVSM CLI provides a “show networkadapters” command to display information for each physical network adapter in the chassis.

```
~$ sudo nvsm show networkadapters
```

Within an NVSM CLI interactive session, targets related to network adapters are located under the /chassis/localhost/NetworkAdapters target.

```
~$ sudo nvsm
nvsm-> cd /chassis/localhost/NetworkAdapters
nvsm(/chassis/localhost/NetworkAdapters)-> show
```

3.5.1. Show Network Ports

NVSM CLI provides a “show networkports” command to display information for each physical network port in the chassis.

```
~$ sudo nvsm show networkports
```

Within an NVSM CLI interactive session, targets related to network adapters are located under the /chassis/localhost/NetworkAdapter/<id>/NetworkPort target, where <id> is one of the network adapter IDs displayed from the `nvsm show networkadapters` command.

```
~$ sudo nvsm
nvsm-> cd /chassis/localhost/NetworkAdapters/<id>/NetworkPorts
nvsm(/chassis/localhost/NetworkAdapters/<id>/NetworkPorts)-> show
```

3.5.2. Show Network Device Functions

NVSM CLI provides a “show networkdevicefunctions” command to display information for each network adapter-centric PCIe function in the chassis.

```
~$ sudo nvsm show networkdevicefunctions
```

Within an NVSM CLI interactive session, targets related to network device functions are located under the /chassis/localhost/NetworkAdapter/<id>/NetworkDeviceFunctions target, where <id> is one of the network adapter IDs displayed from the nvsm show networkadapters command..

```
~$ sudo nvsm
nvsm-> cd /chassis/localhost/NetworkAdapters/<id>/NetworkDeviceFunctions
nvsm(/chassis/localhost/NetworkAdapters/<id>/NetworkDeviceFunctions)-> show
```

3.5.3. Show Network Interfaces

NVSM CLI provides a “show networkinterfaces” command to display information for each logical network adapter on the system.

```
~$ sudo nvsm show networkinterfaces
```

Within an NVSM CLI interactive session, targets related to network adapters are located under the /system/localhost/networkinterfaces target.

```
~$ sudo nvsm
nvsm-> cd /system/localhost/NetworkInterfaces
nvsm(/system/localhost/NetworkInterfaces)-> show
```

3.6. System Monitoring Configuration

NVSM provides a DSHM service that monitors the state of the DGX system.

NVSM CLI can be used to interact with the DSHM system monitoring service via the NVSM API server.

3.6.1. Configuring Email Alerts

In order to receive the Alerts generated by DSHM through email, configure the Email settings in the global policy using NVSM CLI. User shall receive email whenever a new alert gets generated. The sender address, recipient address(es), SMTP server IP address and SMTP server Port number must be configured according to the SMTP server settings hosted by the user.

Email configuration properties

Property	Description
email_sender	Sender email address Must be a valid email address, otherwise no emails will be sent. [sender@domain.com]
email_recipients	List of recipients to which the email shall be sent [user1@domain.com ,user2@domain.com]
email_smtp_server_name	SMTP server name that the user wants to use for re-laying email [smtp.domain.com]
email_smtp_server_port	Port Number used by the SMTP server for providing SMTP relay service. Numeric value

The following examples illustrate how to configure email settings in global policy using NVSM CLI.

```
user@dgx-2:~$sudo nvsm set /policy email_sender=dgx-admin@nvidia.com
```

```
user@dgx-2:~$sudo nvsm set /policy email_smtp_server_name=smtpserver.nvidia.com
```

```
user@dgx-2:~$sudo nvsm set /policy email_recipients=jdoe@nvidia.com,jdeer@nvidia.com
```

```
user@dgx-2:~$sudo nvsm set /policy email_smtp_server_port=465
```

3.6.2. Understanding System Monitoring Policies

From within an NVSM CLI interactive session, system monitor policy settings are accessible under the following targets.

CLI Target		Description
/policy		Global NVSM monitoring policy, such as email settings for alert notifications.
/systems/localhost/gpus/policy		
/systems/localhost/memory/policy		NVSM policy for monitoring DIMM correctable and uncorrectable errors.
/systems/localhost/processors/policy		NVSM policy for monitoring CPU machine-check exceptions (MCE)
/systems/localhost/storage/policy		NVSM policy for monitoring storage drives and volumes
/chassis/policy		
/chassis/localhost/thermal/policy		NVSM policy for monitoring fan speed and temperature as reported by the baseboard management controller (BMC)
/chassis/localhost/power/policy		NVSM policy for monitoring power supply voltages as reported by the BMC
/chassis/localhost/NetworkAdaptors/policy		NVSM policy for monitoring the physical network adaptors
/chassis/localhost/NetworkAdaptors/<ETH>/NetworkPorts/policy	x	NVSM policy for monitoring the network ports for the specified Ethernet network adapter
/chassis/localhost/NetworkAdaptors/<IB>/NetworkPorts/policy	y	NVSM policy for monitoring the network ports for the specified InfiniBand network adapter
/chassis/localhost/NetworkAdaptors/<ETH>/NetworkDevicesFunctions/policy	x	NVSM policy for monitoring the PCIe functions for the specified Ethernet network adapter
/chassis/localhost/NetworkAdaptors/<IB>/NetworkDevicesFunctions/policy	y	NVSM policy for monitoring the PCIe functions for the specified InfiniBand network adapter

3.6.2.1 Global Monitoring Policy

Global monitoring policy is represented by the /policy target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /policy
```

Example output:

```
/policy
Properties:
  email_sender = NVIDIA DSHM Service
  email_smtp_server_name = smtp.example.com
  email_recipients = jdoe@nvidia.com,jdeer@nvidia.com
  email_smtp_server_port = 465
Verbs:
  cd
  set
  show
```

The properties for global monitoring policy are described in the table below.

Property	Description
email_sender	Sender email address [sender@domain.com]
email_recipients	List of recipients to which the email shall be sent [user1@domain.com ,user2@domain.com]
email_smtp_server_name	SMTP server name that the user wants to use for re-laying email [smtp.domain.com]
email_smtp_server_port	Port Number used by the SMTP server for providing SMTP relay service. Numeric value

3.6.2.2 Memory Monitoring Policy

Memory monitoring policy is represented by the /systems/localhost/memory/policy target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /systems/localhost/memory/policy
```

Example output:

```
/systems/localhost/memory/policy
Properties:
  mute_notification =
  mute_monitoring =

Verbs:
  cd
  set
  show
```

The properties for memory monitoring policy are described in the table below.

Property	Syntax	Description
mute_notification	List of comma separated DIMM IDs Example: CPU1_DIMM_A1,CPU2_DIMM_F2	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated DIMM IDs Example: CPU1_DIMM_A1,CPU2_DIMM_F2	Health monitoring is suppressed for devices in the list.

3.6.2.3 Processor Monitoring Policy

Processor monitoring policy is represented by the `/systems/localhost/processors/policy` target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /systems/localhost/processors/policy
```

Example output:

```
/systems/localhost/processors/policy
Properties:
    mute_notification =
    mute_monitoring =

Verbs:
    cd
    set
    show
```

The properties for processor monitoring policy are described in the table below.

Property	Syntax	Description
<code>mute_notification</code>	List of comma separated CPU IDs. Example: CPU0,CPU1	Email alert notification is suppressed for devices in the list.
<code>mute_monitoring</code>	List of comma separated CPU IDs. Example: CPU0,CPU1	Health monitoring is suppressed for devices in the list.

3.6.2.4 Storage Monitoring Policy

Storage monitoring policy is represented by the `/systems/localhost/storage/1/policy` target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /systems/localhost/storage/policy
```

Example output:

```
/systems/localhost/storage/policy
Properties:
    volume_mute_monitoring =
    volume_poll_interval = 10
    drive_mute_monitoring =
    drive_mute_notification =
    drive_poll_interval = 10
    volume_mute_notification =

Verbs:
    cd
    set
    show
```

The properties for storage monitoring policy are described in the table below.

Property	Syntax	Description
drive_mute_notification	List of comma separated drive slots Example: 0, 1 etc	Email alert notification is suppressed for drives in the list.
drive_mute_monitoring	List of comma separated drive slots Example: 0, 1 etc	Health monitoring is suppressed for drives in the list.
drive_poll_interval	Positive integer	DSHM checks the health of the drives periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.
volume_mute_notification	List of comma separated volume identifier Example: md0, md1 etc	Email alert notification is suppressed for volumes in the list
volume_mute_monitoring	List of comma separated volume identifier Example: md0, md1 etc	Health monitoring is suppressed for volumes in the list
volume_poll_interval	Positive integer	DSHM checks the health of the volumes periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.

3.6.2.5 Thermal Monitoring Policy

Thermal monitoring policy (for fan speed and temperature) is represented by the `/chassis/localhost/thermal/policy` target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /chassis/localhost/thermal/policy
```

Example output:

```
/chassis/localhost/thermal/policy
Properties:
  fan_mute_notification =
  pdb_mute_monitoring =
  fan_mute_monitoring =
  pdb_mute_notification =
Verbs:
  cd
  set
  show
```

The properties for thermal monitoring policy are described in the table below.

Property	Syntax	Description
fan_mute_notification	List of comma separated FAN IDs. Example: FAN2_R,FAN1_L,PDB_FAN2	Email alert notification is suppressed for devices in the list.
fan_mute_monitoring	List of comma separated FAN IDs Example: FAN6_F,PDB_FAN1	Health monitoring is suppressed for devices in the list.
pdb_mute_notification	List of comma separated PDB IDs. Example: PDB1,PDB2	Email alert notification is suppressed for devices in the list.
pdb_mute_monitoring	List of comma separated PDB IDs Example: PDB1	Health monitoring is suppressed for devices in the list.

3.6.2.6 Power Monitoring Policy

Power monitoring policy is represented by the /chassis/localhost/power/policy target of NVSM CLI.

```
user@dgx-2:~$ sudo nvsm show /chassis/localhost/power/policy
```

Example output:

```
/chassis/localhost/power/policy
Properties:
  mute_notification =
  mute_monitoring =

Verbs:
  cd
  set
  show
```

The properties for power monitoring policy are described in the table below.

Property	Syntax	Description
mute_notification	List of comma separated PSU IDs. Example: PSU4,PSU2	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated FAN IDs Example: PSU1,PSU4	Health monitoring is suppressed for devices in the list.

3.6.2.7 PCIe Monitoring Policy

Memory monitoring policy is represented by the /systems/localhost/pcie/policy target of NVSM CLI.

```
:~$ sudo nvsm show /systems/localhost/pcie/policy
```

Example output:

```
/systems/localhost/pcie/policy
Properties:
```

(continues on next page)

(continued from previous page)

```

    mute_notification =
    mute_monitoring =

Verbs:
    cd
    set
    show

```

The properties for memory monitoring policy are described in the table below.

Property	Syntax	Description
mute_notification	List of comma separated PCIe IDs	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated PCIe IDs	Health monitoring is suppressed for devices in the list.

3.6.2.8 GPU Monitoring Policy

Memory monitoring policy is represented by the `/systems/localhost/gpus/policy` target of NVSM CLI.

```
:~$ sudo nvsm show /systems/localhost/gpus/policy
```

Example output:

```

/systems/localhost/gpus/policy
Properties:
    mute_notification =
    mute_monitoring =

Verbs:
    cd
    set
    show

```

The properties for memory monitoring policy are described in the table below.

Property	Syntax	Description
mute_notification	List of comma separated GPU IDs	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated GPU IDs	Health monitoring is suppressed for devices in the list.

3.6.2.9 Network Adaptor Monitoring Policies

3.6.2.9.1 Network Adapter Policy

The physical network adapter monitoring policy is represented by the `/chassis/localhost/NetworkAdaptors/policy` target of NVSM CLI.

```
:~$ sudo nvsm show /chassis/localhost/NetworkAdaptors/policy
```

Example output:

```
/chassis/localhost/NetworkAdaptors/policy
Properties:
  mute_notification =
  mute_monitoring =
Verbs:
  cd
  set
  show
```

The properties are described in the following table.

Property	Syntax	Description
<code>mute_notification</code>	List of comma separated physical network adapter IDs.	Email alert notification is suppressed for devices in the list.
<code>mute_monitoring</code>	List of comma separated physical network adapter IDs.	Health monitoring is suppressed for devices in the list.

3.6.2.9.2 Network Ports Policy

The physical network port monitoring policy is represented by the `/chassis/localhost/NetworkAdaptors/<network-id>/NetworkPorts/policy` target of NVSM CLI.

The following uses the network port IB0 to demonstrate this command.

```
:~$ sudo nvsm show /chassis/localhost/NetworkAdaptors/IB0/NetworkPorts/policy
```

Example output:

```
/chassis/localhost/NetworkAdaptors/IB0/NetworkPorts/policy
Properties:
  mute_notification =
  mute_monitoring =
Verbs:
  cd
  set
  show
```

The properties are described in the following table.

Property	Syntax	Description
mute_notification	List of comma separated physical network port IDs.	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated physical network port IDs	Health monitoring is suppressed for devices in the list.

3.6.2.9.3 Network Devices Functions Policy

The network devices functions monitoring policy is represented by the /chassis/localhost/NetworkAdaptors/<network-id>/NetworkDevicesFunctions/policy target of NVSM CLI.

The following uses the network port IB0 to demonstrate this command.

```
:~$ sudo nvsm show /chassis/localhost/NetworkAdaptors/IB0/NetworkDevicesFunctions/
↪policy
```

Example output:

```
/chassis/localhost/NetworkAdaptors/IB0/NetworkDevicesFunctions/policy
Properties:
  mute_monitoring =
  mute_notification =
  rx_collision_threshold = 5
  rx_crc_threshold = 5
  tx_collision_threshold = 5
Verbs:
  cd
  set
  show
```

The properties are described in the following table.

Property	Syntax	Description
mute_notification	List of comma separated network-centric PCIe function IDs. Example: PSU4,PSU2	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated network-centric PCIe function IDs. Example: PSU1,PSU4	Health monitoring is suppressed for devices in the list.
rx_collision_threshold	Positive integer	
rx_crc_threshold	Positive integer	
tx_collision_threshold	Positive integer	

3.7. Performing System Management Tasks

This section describes commands for accomplishing some system management tasks.

3.7.1. Rebuilding a RAID 1 Array

For DGX systems with two NVMe OS drives configure as a RAID 1 array, the operating system is installed on volume md0. You can use NVSM CLI to view the health of the RAID volume and then rebuild the RAID array on two healthy drives.

3.7.1.1 Viewing a Healthy RAID Volume

On a healthy system, this volume appears with two drives and “Status_Health = OK”. For example:

```
nvsm-> cd /systems/localhost/storage
nvsm(/systems/localhost/storage)-> show volumes/md0
/systems/localhost/storage/volumes/md0
Properties:
  Status_State = Enabled
  Status_Health = OK
  Name = md0
  Encrypted = False
  VolumeType = RAID-1
  Drives = [ nvme0n1, nvme1n1 ]
  CapacityBytes = 893.6G
  Id = md0
Targets:
  rebuild
Verbs:
  cd
  show
```

3.7.1.2 Viewing a Degraded RAID Volume

On a system with degraded OS volume, the md0 volume will appear with only one drive, with messages “Status_Health = Warning”, and “Status_State = Degraded” reported as follows.

```
nvsm-> cd /systems/localhost/storage
nvsm(/systems/localhost/storage)-> show volumes/md0
/systems/localhost/storage/volumes/md0
Properties:
  Status_State = Degraded
  Status_Health = Warning
  Name = md0
  Encrypted = False
  VolumeType = RAID-1
  Drives = [ nvme1n1 ]
  CapacityBytes = 893.6G
  Id = md0Targets:
  rebuild
```

(continues on next page)

(continued from previous page)

```
Verbs:
  cd
  show
```

In this situation, the OS volume is missing its parity drive.

3.7.1.3 Rebuilding the RAID 1 Volume

To rebuild the RAID array, make sure that you have installed a known good NVMe drive for the parity drive.

The RAID rebuilding process should begin automatically upon turning on the system. If it does not start automatically, use NVSM CLI to manually rebuild the array as follows.

1. Start an NVSM CLI interactive session and switch to the storage target.

```
$ sudo nvsm
nvsm-> cd /systems/localhost/storage
```

2. Start the rebuilding process and be ready to enter the device name of the replaced drive.

```
nvsm(/systems/localhost/storage)-> start volumes/md0/rebuild
PROMPT: In order to rebuild this volume, a spare drive
        is required. Please specify the spare drive to use
        to rebuild md0.
Name of spare drive for md0 rebuild (CTRL-C to cancel): nvmeXn1
WARNING: Once the volume rebuild process is started, the
        process cannot be stopped.
Start RAID-1 rebuild on md0? [y/n] y
```

3. After entering **y** at the prompt to start the RAID 1 rebuild, the “Initiating rebuild ...” message appears.

```
/systems/localhost/storage/volumes/md0/rebuild started at 2018-10-12 15:27:26.
↪525187
Initiating RAID-1 rebuild on volume md0...
0.0% [ \ ]
```

After about 30 seconds, the “Rebuilding RAID-1 ...” message should appear.

```
/systems/localhost/storage/volumes/md0/rebuild started at 2018-10-12 15:27:26.
↪525187
Rebuilding RAID-1 rebuild on volume md0...
31.0% [===== / ]
```

If this message remains at “Initiating RAID-1 rebuild” for more than 30 seconds, then there is a problem with the rebuild process. In this case, make sure the name of the replacement drive is correct and try again.

The RAID 1 rebuild process should take about 1 hour to complete.

For more detailed information on replacing a failed NVMe OS drive, see the [NVIDIA DGX-2 Service Manual](#).

3.7.2. Setting MaxQ/MaxP on DGX-2 Systems

Beginning with DGX OS 4.0.5, you can set two GPU performance modes – MaxQ or MaxP.

Note: Support on DGX-2 systems requires BMC firmware version 1.04.03 or later. MaxQ/MaxP is not supported on DGX-2H systems.

3.7.2.1 MaxQ

- ▶ Maximum efficiency mode
- ▶ Allows two DGX-2 systems to be installed in racks that have a power budget of 18 kW.
- ▶ Switch to MaxQ mode as follows.

```
$ sudo nvsm set powermode=maxq
```

The settings are preserved across reboots.

3.7.2.2 MaxP

- ▶ Default mode for maximum performance
- ▶ GPUs operate unconstrained up to the thermal design power (TDP) level.
In this setting, the maximum DGX-2 power consumption is 10 kW.
- ▶ Provides reduced but better performance than MaxQ when only 3 or 4 PSUs are working.
- ▶ If you switch to MaxQ mode, you can switch back to MaxP mode as follows:

```
$ sudo nvsm set powermode=maxp
```

The settings are preserved across reboots.

3.7.3. Configuring Support for Custom Drive Partitioning

DGX systems incorporate data drives configured as either RAID 0 or RAID 5 arrays, depending on the product. You can alter the default configuration by adding or removing drives, or by switching between a RAID 0 configuration and a RAID 5 configuration. If you alter the default configuration, you must let NVSM know so that the utility does not flag the configuration as an error, and so that NVSM can continue to monitor the health of the drives.

To configure NVSM to support a custom drive partitioning, perform the following.

1. Stop NVSM services.

```
$ systemctl stop nvsm
```

2. Edit `/etc/nvsm/nvsm.config` and set the `"use_standard_config_storage"` parameter to `false`

```
"use_standard_config_storage":false
```

3. Remove the NVSM database.

```
$ sudo rm /var/lib/nvsm/sqlite/nvsm.db
```

4. Restart NVSM.

```
$ systemctl restart nvsm
```

Remember to set the parameter back to `true` if you restore the drive partition back to the default configuration.

Chapter 4. Configuring NVSM Security

This chapter explains shows how to secure the NVSM API installation.

4.1. Overview of NVSM Security

NVSM APIs are served using the HTTPS protocol. HTTPS requires the NVSM API server to possess a public-private key pair as well as a certificate that it presents to connecting clients. The certificate also needs to be signed by a certificate authority (CA) using the private key of that CA.

For proper security, this certificate+key should be provided by users. It cannot be provided by NVIDIA because

- ▶ The private key should be known only to the user, and should not be known to NVIDIA, and
- ▶ NVIDIA is not a [Certificate Authority](#)

To allow the NVSM software stack to work right out of the box, the installation process creates some sample key pairs and certificates. These certificates are created with dummy values for country, organization, organization unit, etc. because the installation does not include these details. Also, the generated CA certificate is self signed. These sample certificates must NOT be used in a production environment.

NVSM allows you to provide your own key-pairs and certificates with correct values that are properly signed by a trusted CA. Details of key generation and certificate chains is beyond the scope of this document. However, an example setup is shown below to show how NVSM can be configured with customer provided/generated keys and certificates.

4.2. What You Need to Configure NVSM Security

To configure NVSM security, you need the following, either copied from a CA provider or generated locally and copied to a location on the system.

- ▶ X.509 certificate for the NVSM REST server
Example path and filename: `/pki/node1.crt`
- ▶ Private key file corresponding to the above certificate
Example path and filename: `/pki/node1.key`

- The certificate of the CA who issued the above certificate

Example path and filename: `/pki/ca.crt`

An explanation of how to generate or obtain these certificates and keys is beyond the scope of this document since these have to be in compliance with the overall security architecture of the data center. In the most simplest form, users might use commands such as OpenSSL to generate their own certificate chain and keys. You may prefer to use free services such as <https://letsencrypt.org/> to acquire them.

4.3. How to Configure NVSM Security

1. Edit the NVSM configuration file to use the paths and filenames of your certificate files and key file.

Edit the `ca_cert`, `https_cert`, and `https_priv_key` configuration parameters to specify the path and filenames that NVSM shall use. The following use the example path and filenames.

```
"ca_cert":"/pki/ca.crt",
```

```
"https_cert":"/pki/node1.crt",
```

```
"https_priv_key":"/pki/node1.key",
```

2. Restart the NVSM service.

```
$ sudo systemctl restart nvsm
```

Chapter 5. NVSM Call Home

The NVIDIA System Manager (NVSM) Call Home, when enabled and with an internet connection, provides additional automation to NVSM health monitoring functionality. Instead of having to contact NVIDIA Enterprise Support to report critical alerts from NVSM, submit system logs, nvsm dump health files, and DGX serial numbers to create a support ticket, NVSM Call Home automates those tasks. This reduces overall turnaround time for resolving issues.


5.1. NVSM Call Home Overview

When NVSM raises a critical alert, NVSM Call Home performs the following actions:

- ▶ Proactively pushes **Critical** level alerts to NVIDIA Enterprise Services.
- ▶ Collects the system and nvsm dump, and system serial number, and uploads them to NVIDIA Enterprise Services..
- ▶ Creates a Support Case on the Enterprise Support portal.

NVSM Call Home also clears resolved alerts and pushes the updated status to NVIDIA Enterprise Services.

The figure below illustrates the end-to-end flow of the NVSM Call Home feature.



When NVSM Call Home is enabled on the DGX system and a critical alert is raised, the NVSM daemon on the DGX system initiates an HTTPS connection to the secured NVIDIA Enterprise Services backend and communicates the alert details and logs using the RESTful interface. The information is validated

and then a new Support Ticket is created on the NVIDIA Enterprise Support Portal. Communication of all alerts, including status changes for the alerts, is through REST calls.

NVSM Call Home operates in three different modes. It is not enabled by default, so to use NVSM Call Home you must enable one of the following modes:

5.1.1. Policy-enabled “automatic” Mode

- ▶ This mode batches alert submissions at regular intervals and then pushes them to the NVIDIA Enterprise Support portal.
- ▶ An internet connection is required.
- ▶ To enable, issue

```
$ sudo nvsm set /policy callhome_enable=true
```

See the section *Using NVSM Call Home in Automatic Mode* for details as well as configuration options for automatic mode.

5.1.2. Policy-enabled Offline Mode

- ▶ This mode is useful for air-gapped or highly-secured environments where access to the internet is limited.
- ▶ Instead of sending batched submissions to NVIDIA Enterprise Support, the alert and system information are stored on the local system. Users need to manually provide the gathered information to NVIDIA Enterprise Services to create a support case.
- ▶ To enable, issue

```
$ sudo nvsm set /policy offline_callhome_enable=true
```

See the section *Using NVSM Offline Call Home* for details as well as configuration options for offline mode.

5.1.3. On-demand Mode

- ▶ This mode gathers alert information and creates a submission on-the-fly to the Enterprise Support Portal.
- ▶ An internet connection is required.
- ▶ To initiate a Call Home submission on-demand, issue

```
$ sudo nvsm set /callhome trigger=true
```

See the section *Using NVSM On-Demand Mode* for details as well as configuration options for on-demand mode.

5.2. Using NVSM Call Home

NVSM Call Home can operate in two modes:

- ▶ **Automatic Mode** - NVSM Call Home operates automatically at regular intervals.
- ▶ **On-demand Mode** - NVSM Call Home sequence is initiated manually.

You can also set up NVSM Call Home to run offline; for example, on air-gapped systems.

5.2.1. Prerequisites for Using NVSM Call Home in Automatic or On-Demand Mode

5.2.1.1 Enabling Ports

Since NVSM Call Home communicates with the external NVIDIA server, port 443 must be enabled prior to operating NVSM Call Home.

5.2.1.2 Enabling Access

You need to register your system for NVSM Call Home so that the NVIDIA Services Cloud recognizes the system. Contact [NVIDIA Enterprise Services](#) to set up NVIDIA Call Home for your DGX system.

5.2.1.3 Validating NVSM Call Home Readiness

Before using NVSM Call Home, make sure the server is ready to support NVSM Call Home by performing a diagnostic test. The test does not create a ticket with NVIDIA Enterprise Services, but does test that the system is able to communicate with the NVIDIA Enterprise Services infrastructure.

To run the diagnostic, issue the following:

```
$ sudo nvsm set /callhome trigger=true diagtest=true
```

Note: This uses the on-demand mode of NVIDIA Call Home, explained in more detail in the section [Using NVSM On-demand Mode](#).

To see the result of the last diagnostic test run, issue the following:

```
$ sudo nvsm show /callhome
```

Example output confirming the setup is ready for Call Home operation. Lines of interest are identified in bold.

```
/callhome
Properties:
  Trigger = False
  Op_Description = User initiated call home operation.
  Op_DiagTest = True
```

(continues on next page)

(continued from previous page)

```
Op_CaseId = none
Op_State = Succeeded
Op_StartTime = 2019-06-24T06:10:17Z
Op_Message = Call Home operation Succeeded
Op_Email =
```

If the output reports errors or failures, contact NVIDIA technical support for assistance.

5.2.2. Using NVSM Call Home in Automatic Mode

When automatic mode is enabled, NVSM monitors the server continuously and pushes critical or cleared alerts to NVESC and creates a support case on behalf of the registered user.

Automatic Mode Syntax

To enable automatic mode, issue the following.

```
$ sudo nvsm set /policy callhome_enable=true [callhome_email_contact=<email>]
↳[callhome_batch_interval=<time-in-seconds>]
```

Automatic Mode Configuration Options

You can configure NVSM Call Home using the following parameters:

► **callhome_email_contact**

An option to set an email-id. This should be a registered user of the NVIDIA Enterprise Support Portal. The email gets embedded in the case/ticket created in NVIDIA Enterprise Support Portal.

► **callhome_batch_interval**

Enabling automatic mode batches alert submissions at regular intervals and then pushes them to the NVIDIA Enterprise Support portal. Any raised alerts within that time frame will be sent (as individual Support Cases). By default, the interval is 600 seconds (10 minutes), but you can use this option to specify other intervals (in seconds).

Automatic Mode Example

The following example illustrates how to use these parameters.

```
$ sudo nvsm set /policy callhome_enable=true callhome_email_contact=123@example.com
↳callhome_batch_interval=610
```

Verifying Automatic Mode Status

To verify the status of the current setup, issue the following.

```
$ sudo nvsm show /policy
```

Example output:

```
/policy
Properties:
  callhome_batch_interval = 610
  callhome_email_contact = 123@example.com
  callhome_enable = True
  email_recipients =
```

(continues on next page)

(continued from previous page)

```
email_sender =
email_smtp_server_name =
email_smtp_server_port = 0
```

`callhome_enable = True` indicates that Call Home automatic mode is enabled.

5.2.3. Using NVSM On-Demand Mode

NVSM Call Home On-Demand mode is a user-triggered call-home action. Triggering Call Home on-demand creates a Support Case with NVIDIA Enterprise Support Portal that includes a captured system dump ('nvsm dump health'). NVSM Call Home On-Demand can be used whether or not automatic mode is enabled.

On-Demand Mode Syntax

To trigger an NVSM Call Home sequence on-demand, issue the following.

```
# sudo nvsm set /callhome trigger=true [description="<description>"] [email=<email>
```

To cancel an on-demand Call Home in progress, issue the following.

```
# sudo nvsm set /callhome trigger=false
```

See the next section for an explanation of the optional parameters.

On-Demand Mode Configuration Options

You can configure NVSM Call Home triggered on-demand using the following parameters:

► **email**

This option sets an email-id. The email gets embedded in the case/ticket created in NVIDIA Enterprise Support Portal.

► **description**

This option lets you describe the purpose or the details for triggering On-Demand Call Home.

Examples of descriptive strings:

"Testing"

"System running low in performance, takes several minutes to perform an nvidia-smi command."

On-Demand Example

The following example illustrates how to use these optional parameters.

```
# sudo nvsm set /callhome trigger=true description="testing" email=123@example.com
```

Verifying On-Demand Status

To check the status of a Call Home sequence initiated on-demand, issue the following.

```
# sudo nvsm show /callhome
```

The following example output shows the progress of the Call Home sequence.

```
/callhome
Properties:
  Trigger = True
  Op_Description = testing
  Op_CaseId = none
  Op_State = Running
  Op_StartTime = 2019-06-12T08:28:45Z
  Op_Message = Collecting logs
  Op_Email = 123@example.com
```

The following example output shows that a case ID was created in the NVIDIA Enterprise Support portal.

```
/callhome
Properties:
  Trigger = False
  Op_Description = testing
  Op_CaseId = 0001XXX
  Op_State = Succeeded
  Op_StartTime = 2019-06-12T08:28:45Z
  Op_Message = Call Home operation Succeeded
  Op_Email = 123@example.com
```

`Trigger = False` indicates that the on-demand sequence is not running - in this case because it has completed.

5.2.4. Using NVSM Offline Call Home

To support DGX systems installed in air-gapped or highly-secured environments where access to the internet is limited, NVSM Call Home can be operated in offline mode (Offline Call Home). Like standard Call Home, Offline Call Home software proactively monitors the health of the DGX system and automatically

- ▶ Collects system dump and logs, and
- ▶ Collects alerts and system information.

However, instead of sending the information to NVIDIA Enterprise Services, NVSM Offline Call Home stores the information in a user-specified directory on the DGX system. Also, unlike standard Call Home, Offline Call Home operates in automatic mode only; there is no on-demand mode in Offline Call Home.

Prerequisites

Offline Call Home and standard Call Home cannot be enabled at the same time. To ensure that standard automatic-mode Call Home is not enabled, issue the following before enabling Offline Call Home.

```
$ sudo nvsm set /policy callhome_enable=false
```

Enabling Offline Call Home

Like standard Call Home, use the `nvsm set /policy` command to enable Offline Call Home.

```
$ sudo nvsm set /policy offline_callhome_enable=true \
offline_callhome_dump_destination_location=<path/to/location> \
offline_callhome_batch_interval=<batch-interval> \
offline_callhome_no_of_dumps_allowed=<number>
```

Offline Call Home Configuration Options

► offline_callhome_dump_destination_location

By default, Offline Call Home stores the system logs at `/var/log/nvsm_offline_callhome`. You can set a different location using this option.

► offline_callhome_batch_interval

Enabling Offline Call Home creates a batch of alerts at regular intervals and then pushes them to local storage. By default, the interval is 600 seconds (10 minutes), but you can use this option to specify other intervals (in seconds).

► offline_callhome_no_of_dumps_allowed

By default, NVSM Offline Call Home will store 9999999 different log files, but you can specify a smaller number for which to allocate space as needed.

Example of Enabling Offline Call Home

The following example illustrates how to use these parameters.

```
$ sudo nvsm set /policy \
offline_callhome_enable=true \
offline_callhome_dump_destination_location=/tmp/offline_callhome_dump \
callhome_batch_interval=610 \
offline_callhome_no_of_dumps_allowed=10
```

Verifying the Offline Call Home Configuration

To verify the status of the current setup, issue the following.

```
$ sudo nvsm show /policy
```

Example output showing the offline Call Home policy details.

```
/policy
Properties:
  offline_callhome_batch_interval = 610
  offline_callhome_enable = True
  offline_callhome_dump_destination_location = /tmp/offline_callhome_dump
  offline_callhome_no_of_dumps_allowed = 10
```

Verifying Contents of the Dump File

The contents of each batch is stored in a tar file.

► Tar file naming format:

`offline_callhome-nvsm-health-<timestamp>_<serial-number>.tar.xz`

► The tar file contents -

► System dump

► JSON metadata file that lists the system and critical alerts.

JSON file naming format:

`offline_callhome_notifications-<timestamp>_<hostname>_<serial-number>.json`

► JSON file format showing type of data included :

```
{
  "system_serial": "<serial number>",
  "system_name": "<hostname>",
  "notifications": [
    {
      "alert_id": "",
      "clear_time": "-",
      "component_id": "",
      "description": "",
      "event_time": "",
      "message": "",
      "message_details": ".",
      "recommended_action": "",
      "severity": "",
      "system_name": "",
      "system_serial": "",
      "type": ""
    }
  ]
}
```

Chapter 6. Third-Party Licenses

This NVIDIA product contains third party software that is being made available to you under their respective open source software licenses. Some of those licenses also require specific legal information to be included in the product. This section provides such information.

6.1. `mattn/go-sqlite3`

The `mattn/go-sqlite3` software (<https://github.com/mattn/go-sqlite3>) is provided under the following terms: The MIT License (MIT)

Copyright (c) 2014 Yasuhiro Matsumoto

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Chapter 7. Notices

7.1. Notice

THE INFORMATION IN THIS GUIDE AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS GUIDE IS PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the product described in this guide shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS GUIDE IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this guide will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this guide. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this guide, or (ii) customer product designs.

Other than the right for customer to use the information in this guide with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this guide. Reproduction of information in this guide is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

7.2. Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

7.3. Copyright

© 2018 NVIDIA Corporation. All rights reserved.

Copyright

©2019-2023, NVIDIA