



# NVIDIA SYSTEM MANAGEMENT

DU-09242-001 \_v02 | November 2018

## User Guide



# TABLE OF CONTENTS

<b>Chapter 1. Introduction.....</b>	<b>1</b>
1.1. DGX System Health Monitors.....	2
1.2. Configurable DSHM Features.....	3
1.2.1. Health Monitor Alerts.....	3
1.2.2. DSHM Alert List.....	3
1.2.3. Health Monitor Policies.....	6
1.3. Verifying the Installation.....	6
1.3.1. Verifying DSHM Services.....	6
1.3.2. Verifying NVSM APIs Services.....	7
<b>Chapter 2. Using the NVSM CLI.....</b>	<b>8</b>
2.1. Using the NVSM CLI Interactively.....	8
2.2. Using the NVSM CLI Non-Interactively.....	9
2.3. Getting Help.....	9
2.3.1. nvsm "man" Page.....	10
2.3.2. nvsm --help Flag.....	10
2.3.3. Help for NVSM CLI Commands.....	10
2.4. Examining System Health.....	10
2.4.1. List of Basic Commands.....	10
2.4.2. Show Health.....	12
2.4.3. Dump Health.....	12
2.4.4. Show Storage.....	12
2.4.4.1. Show Storage Alerts.....	13
2.4.4.2. Show Storage Drives.....	14
2.4.4.3. Show Storage Volumes.....	15
2.4.5. Show GPUs.....	16
2.4.5.1. Showing Individual GPUs.....	17
2.4.5.2. Identifying GPU Health Incidents.....	17
2.4.6. Show Processors.....	19
2.4.6.1. Show Processor Alerts.....	20
2.4.7. Show Memory.....	21
2.4.7.1. Show Memory Alerts.....	22
2.4.8. Show Fans and Temperature.....	23
2.4.8.1. Show Thermal Alerts.....	23
2.4.8.2. Show Fans.....	24
2.4.8.3. Show Temperatures.....	25
2.4.9. Show Power Supplies.....	26
2.4.9.1. Show Power Alerts.....	27
2.5. System Monitoring Configuration.....	28
2.5.1. Configuring Email Alerts.....	28
2.5.2. Understanding System Monitoring Policies.....	29

2.5.2.1. Global Monitoring Policy.....	29
2.5.2.2. Memory Monitoring Policy.....	30
2.5.2.3. Processor Monitoring Policy.....	30
2.5.2.4. Storage Monitoring Policy.....	31
2.5.2.5. Thermal Monitoring Policy.....	32
2.5.2.6. Power Monitoring Policy.....	33
2.6. Performing System Management Tasks.....	33
2.6.1. Rebuilding a RAID 1 Array.....	34

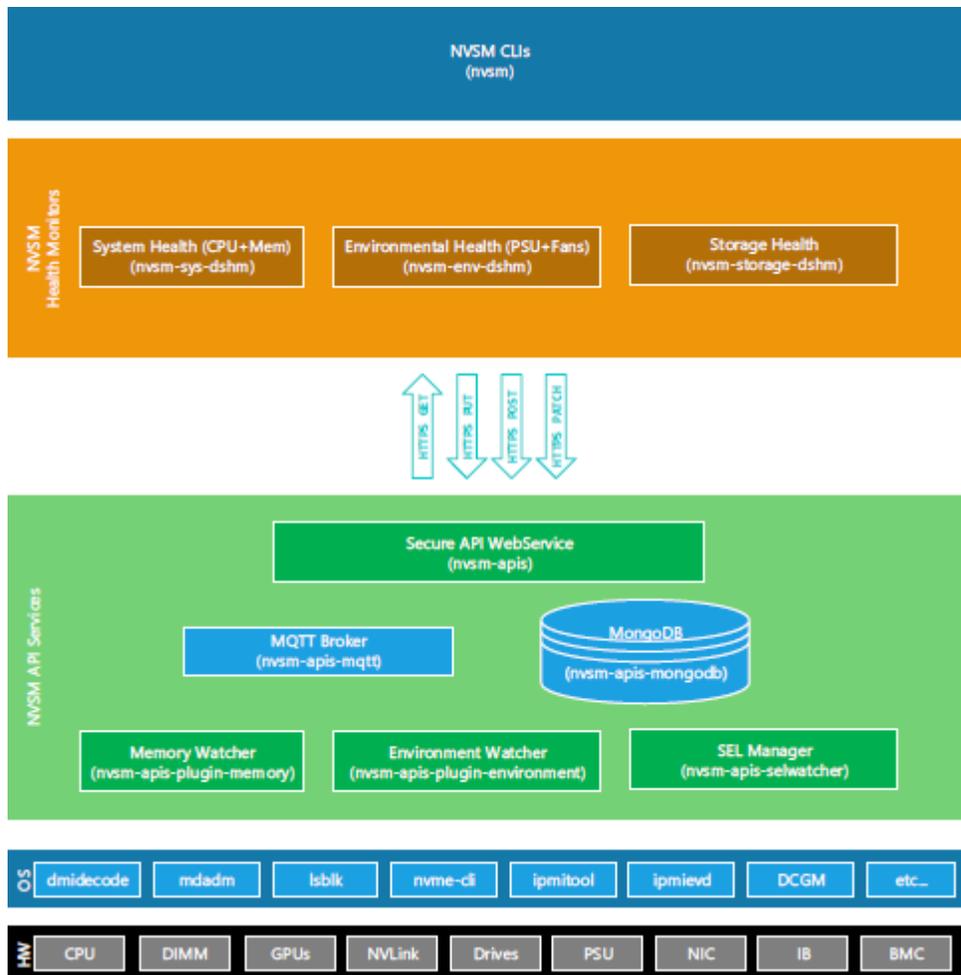


# Chapter 1.

## INTRODUCTION

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. It 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.



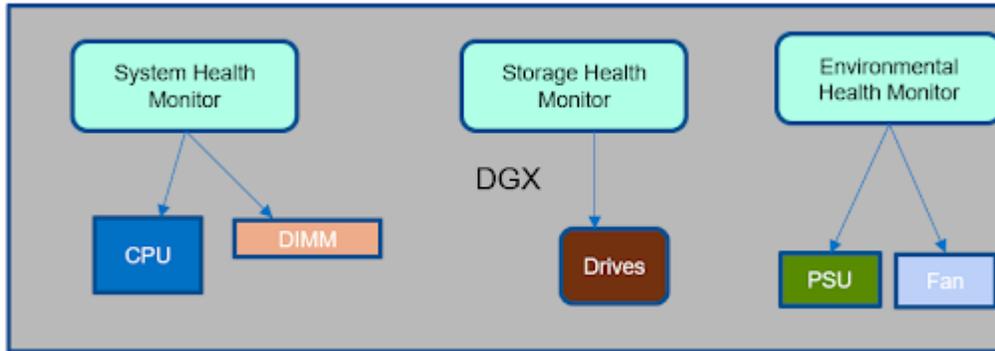
## 1.1. DGX System Health Monitors

The NVSM software incorporates the DGX System Health Monitor (DSHM), which probes critical hardware components in a DGX system and provides notification of fluctuations in system health, faults, and potential failures.

Health monitors are responsible for monitoring the health of critical DGX system components and informing users when an event of significance is detected. Below are the list of health monitors.

- ▶ System Health Monitors
  - ▶ CPU
  - ▶ DIMM
- ▶ Storage Health Monitor
- ▶ Environment Health Monitors
  - ▶ PSU
  - ▶ Fan

The following diagram illustrates the individual health monitors within DSHM.



Each health monitor is launched as a systemd service and leverages NVSM APIs to perform health management responsibilities. Periodic polling of critical system events are performed by each monitor and on identifying an event of significance, the monitor raises an alert. The alert is recorded in persistent storage (on the OS drive) and a notification is sent to configured users.

## 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.

Event	Alert ID	Component ID	Message	Severity
Drive missing	NV-DRIVE-01	<drive slot>	Drive missing in slot <slot number>	Critical

Event	Alert ID	Component ID	Message	Severity
Media errors in drive	NV-DRIVE-02	<drive slot>	Media errors detected in drive <slot number>	Warning
IO errors in drive	NV-DRIVE-03	<drive slot>	IO errors detected in drive <slot number>	Warning
NVMe controller failure in drive	NV-DRIVE-04	<drive slot>	NVMe controller failure detected in drive <slot number>	Critical
Drive available capacity below 10 percent	NV-DRIVE-05	<drive slot>	Available capacity percentage below critical threshold for drive <slot number>	Critical
Drive used percentage above 90	NV-DRIVE-06	<drive slot>	Drive used percentage above critical threshold for drive <slot number>	Critical
Unsupported drive inserted	NV-DRIVE-07	<drive slot>	System has unsupported drive <slot number>	Warning
RAID-0 corrupted	NV-VOL-01	NA	RAID-0 corrupted	Critical
RAID-1 corrupted	NV-VOL-02	NA	RAID-1 corrupted	Critical
ESP-1 corrupted	NV-VOL-03	NA	EFI System Partition 1 is corrupted	Warning
ESP-2 corrupted	NV-VOL-04	NA	EFI System Partition 2 is corrupted	Warning
Power supply failure detected	NV-PSU-01	<PSU#> where # is the PSU number.	Power supply module has failed.	Critical
PSU Predictive failure	NV-PSU-02	<PSU#> where # is the PSU number.	Detected predictive failure of the Power supply module.	Warning
PSU Input lost (AC/DC)	NV-PSU-03	<PSU#> where # is the PSU number.	Input to the Power supply module is missing	Critical
PSU input lost or out of range	NV-PSU-04	<PSU#> where # is the PSU number.	Input voltage is out of range for the Power Supply Module.	Critical
PSU Absent	NV-PSU-05	<PSU#> where # is the PSU number.	PSU is missing.	Warning
PDB Thermal exceeded	NV-PDB-01	<PDB#> where # is the PDB number	Operating temperature exceeds the thermal specifications of the component.	Critical
Fan speed exceeded	NV-FAN-01	<FAN#_F> or <FAN#_R>	Fan speed reading has exceeded the expected speed setting	Critical

Event	Alert ID	Component ID	Message	Severity
		where # is the fan module number. F is for front fan. R is for rear fan.		
Fan speed readings unavailable	NV-FAN-02	<FAN#_F> or <FAN#_R> where # is the fan module number. F is for front fan. R is for rear fan.	Fan readings are inaccessible.	Critical
CPU Internal error	NV-CPU-01	<CPU#> where # is the CPU socket number (CPU0 or CPU1)	An unrecoverable CPU Internal error has occurred.	Critical
CPU Thermtrip	NV-CPU-02	<CPU#> where # is the CPU socket number (CPU0 or CPU1)	CPU Thermtrip has occurred, processor socket temperature exceeded the thermal specifications of the component.	Critical
DIMM Uncorrectable ECC	NV-DIMM-01	<CPU#_DIMM_@ \$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Uncorrectable error is reported.	Critical
DIMM Correctable ECC	NV-DIMM-02	<CPU#_DIMM_@ \$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Correctable errors reported exceeds the configured threshold.	Warning
DIMM Critical	NV-DIMM-03	<CPU#_DIMM_@ \$> where # = (1, 2) @ = (A, B, C, D, E, F) \$ = (1, 2)	Unrecoverable error is observed on the DIMM, specific details of the error are unavailable.	Critical
GPU Critical	NV-GPU-01		System entered degraded mode, GPU is reporting an error	Critical

Event	Alert ID	Component ID	Message	Severity
PCI Sub-system Link Speed Warning	NV-PCI-01		System entered degraded mode, PCI is reporting an error on the GPU endpoint	Warning
PCI Sub-system Link Width Warning	NV-PCI-02		System entered degraded mode, PCI is reporting an error on the GPU endpoint	Warning

### 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 BaseOS 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.

```
$ systemctl status nvsm-env-dshm

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

Other modules can be verified using similar commands:

To verify the storage module:

```
$ sudo systemctl status nvsm-storage-dshm
```

To verify the system module:

```
$ sudo systemctl status nvsm-sys-dshm
```

To verify the environment module

```
$ sudo systemctl status nvsm-env-dshm
```

## 1.3.2. Verifying NVSM APIs Services

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

**nvsm-apis-plugin-environment**

**nvsm-apis-mqtt**

**nvsm-apis-plugin-memory**

**nvsm-apis-mongodb**

**nvsm-apis**

**nvsm-apis-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-apis-plugin-memory
```

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

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

# Chapter 2.

## 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.

### 2.1. 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
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:~$

```

## 2.2. 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.

## 2.3. 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.

### 2.3.1. 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
```

### 2.3.2. 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'.
```

### 2.3.3. 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
```

## 2.4. 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.

### 2.4.1. 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.

Global Commands	Descriptions
\$ sudo nvsm show alerts	
\$ sudo nvsm show policy	
Health Commands	
\$ sudo nvsm show health	Displays overall system health
\$ sudo nvsm dump health	Generates a health report file
Storage Commands	
\$ sudo nvsm show storage	Displays all storage-related information
\$ sudo nvsm show drives	Displays the storage drives
\$ sudo nvsm show volumes	Displays the storage volumes
GPU Commands	
\$ sudo nvsm show gpus	
Processor Commands	
\$ sudo nvsm show processors	Displays information for all CPUs in the system
\$ sudo nvsm show cpus	Alias for "show processors"
Memory Commands	
\$ sudo nvsm show memory	Displays information for all installed DIMMs
\$ sudo nvsm show dimms	Alias for "show memory"
Thermal Commands	
\$ sudo nvsm show fans	
\$ sudo nvsm show temperatures	
\$ sudo nvsm show temps	Alias for "show temperatures"
Power Commands	
\$ sudo nvsm show power	
\$ sudo nvsm show psus	Alias for "show power"

## 2.4.2. 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.

## 2.4.3. 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             vdisplay
```

## 2.4.4. 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/1
nvsm(/systems/localhost/storage/1)-> show
```

Example output:

```
/systems/localhost/storage/1
Properties:
  DriveCount = 10
  Volumes = [ md0, md1, nvme0n1p1, nvme1n1p1 ]
Targets:
  alerts
  drives
  policy
  volumes
Verbs:
  cd
  show
```

#### 2.4.4.1. 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/1/alerts
nvsm(/systems/localhost/storage/1/alerts)-> show
```

Example output:

```
/systems/localhost/storage/1/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/1/alerts)-> show alert1
```

```
/systems/localhost/storage/1/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

#### 2.4.4.2. 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/1/drives
nvsm(/systems/localhost/storage/1/drives)-> show

```

**Example output:**

```

/systems/localhost/storage/1/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/1/drives)-> show nvme2n1
```

```

/systems/localhost/storage/1/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

```

**2.4.4.3. 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 nvsmnvsm-> cd /systems/localhost/storage/1/volumes
nvsm(/systems/localhost/storage/1/volumes)-> show

```

**Example output:**

```

/systems/localhost/storage/1/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/1/volumes)-> show md0
```

```
/systems/localhost/storage/1/volumes/md0P
properties:
  Status_State = Enabled
  Status_Health = OK
  Name = md0
  Encrypted = False
  VolumeType = RAID-1
  Drives = [ nvme0n1, nvme1n1 ]
  CapacityBytes = 893.6G
  Id = md0
Verbs:
  cd
  show
```

## 2.4.5. 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.



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:
  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.

## 2.4.6. 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
    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
```

### 2.4.6.1. 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.
  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.

Alert ID	Severity	Details
NV-CPU-02	Critical	CPU Thertmtrip has occurred, processor socket temperature exceeded the thermal specifications of the component.

## 2.4.7. 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
  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

```

### 2.4.7.1. 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:
  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.

## 2.4.8. 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
```

### 2.4.8.1. 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:

```
nvsml(/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.

### 2.4.8.2. 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 nvsml
nvsml-> cd /chassis/localhost/thermal/fans
nvsml(/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
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

```

### 2.4.8.3. 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
  UpperThresholdNonCritical = 127.000
  SensorNumber = 66h
  UpperThresholdCritical = 127.000
Verbs:
  cd
  show
```

## 2.4.9. 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

```

### 2.4.9.1. 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

## 2.5. 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.

### 2.5.1. 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 relaying 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
```

## 2.5.2. 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/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/1/policy	NVSM policy for monitoring storage drives and volumes
/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

### 2.5.2.1. 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 ]

Property	Description
email_smtp_server_name	SMTP server name that the user wants to use for relaying email [ smtp.domain.com ]
email_smtp_server_port	Port Number used by the SMTP server for providing SMTP relay service. Numeric value

### 2.5.2.2. 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 =
  poll_interval = 10
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.
poll_interval	Positive integer	DSHM checks the health of the devices periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.

### 2.5.2.3. 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 =
poll_interval = 30
Verbs:
  cd
  set
  show

```

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

Property	Syntax	Description
mute_notification	List of comma separated CPU IDs. Example: CPU0,CPU1	Email alert notification is suppressed for devices in the list.
mute_monitoring	List of comma separated CPU IDs Example: CPU0,CPU1	Health monitoring is suppressed for devices in the list.
poll_interval	Positive integer	DSHM checks the health of the devices periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.

#### 2.5.2.4. 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/1/policy
```

Example output:

```

/systems/localhost/storage/1/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.

Property	Syntax	Description
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.

### 2.5.2.5. 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 =
  fan_poll_interval = 20
  pdb_poll_interval = 10
  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.
fan_poll_interval	Positive integer	DSHM checks the health of the devices periodically. By default, this polling occurs every 10

Property	Syntax	Description
		seconds. The poll interval can be configured through this property.
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.
pdb_poll_interval	Positive integer	DSHM checks the health of the devices periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.

### 2.5.2.6. 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 =
  poll_interval = 10
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.
poll_interval	Positive integer	DSHM checks the health of the devices periodically. By default, this polling occurs every 10 seconds. The poll interval can be configured through this property.

## 2.6. Performing System Management Tasks

This section describes commands for accomplishing some system management tasks.

## 2.6.1. 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.

### 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
```

### 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
Verbs:
  cd
  show
```

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

## 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](#).

## 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.

## Trademarks

NVIDIA, the NVIDIA logo, DGX, DGX-1, DGX-2, and DGX Station 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.

## Copyright

© 2018 NVIDIA Corporation. All rights reserved.