ES3000 V5 NVMe PCIe SSD

User Guide

Issue 07

Date 2019-03-19

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://www.huawei.com

Email: support@huawei.com

About This Document

Purpose

This document provides the product information about the Huawei ES3000 V5 PCIe SSD (ES3000 V5 for short) and describes how to install, configure, operate, and maintain the ES3000 V5.

NOTE

The document does not contain the commands for production, equipment, and return detection. These commands are commonly used during engineering implementation and fault locating. Incorrect use of the commands results in device faults or service interruption. To obtain the commands, contact Huawei technical support.

Intended Audience

This document is intended for:

- Server installation engineers
- Server maintenance engineers

Symbol Conventions

The symbols that may be found in this document are defined as follows:

Symbol	Description
DANGER	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
CAUTION	Indicates a potentially hazardous situation which, if not avoided, may result in minor or moderate injury.
NOTICE	Indicates a potentially hazardous situation which, if not avoided, could result in equipment damage, data loss, performance deterioration, or unanticipated results. NOTICE is used to address practices not related to personal injury.

Symbol	Description
NOTE	Calls attention to important information, best practices and tips. NOTE is used to address information not related to personal injury, equipment damage, and environment deterioration.
TIP	Provides a tip that may help you solve a problem or save time.

Change History

Changes between document issues are cumulative. The latest document issue contains all the changes made in earlier issues.

Issue	Date	Description
07	2019-03-19	Modified notes on the hot swap function.
06	2019-03-11	Modified command description.
05	2018-12-22	Modified the Information Collection chapter.
04	2018-11-13	Added information about SSD cards.
03	2018-10-23	Modified descriptions of the command line and updated output of some commands.
02	2018-08-06	Modified some commands and added operations for querying formatting parameter support information.
01	2018-07-20	This issue is the first official release.

Contents

1 Safety

About This Chapter

This section describes the safety precautions to be observed when you install and maintain devices.

1.1 General Safety Notice

This section describes the safety precautions you must take before installing or maintaining Huawei equipment.

Overview

- To ensure safety of humans and the equipment, pay attention to the safety symbols on the equipment and all the safety instructions in this document.
- The "NOTE", "CAUTION", and "WARNING" marks in other documents do not represent all the safety instructions. They are only supplements to the safety instructions.

Local Safety Regulations

When operating Huawei equipment, you must follow the local laws and regulations. The safety instructions in this document are only supplements to the local laws and regulations.

Basic Requirements

- Installation and maintenance personnel must understand basic safety precautions to avoid hazards.
- When operating Huawei equipment, in addition to following the general precautions in this document, follow the specific safety instructions given by Huawei.
- Only trained and qualified personnel are allowed to install, operate, and maintain Huawei equipment.

Ground

- Do not damage the ground conductor or operate the device in the absence of a properly installed ground conductor. Conduct the electrical inspection carefully.
- The device must be connected permanently to the protection ground before an operation.

Human Safety

- Do not operate the device or cables during lightning strikes.
- Move or lift the chassis by holding its lower edge. Do not hold the handles on certain modules such as power supply, fans, and boards because they cannot support the weight of the device.
- At least two persons are required to lift the chassis. When lifting it, keep your back straight and move stably.
- Do not look into the optical port without eye protection.
- Do not wear jewelry or watches when you operate the device.

Installation

- The device (or system) must be installed in an access-controlled location.
- The device must be fixed securely on the floor or to other immovable objects such as walls and mounting racks before operation.
- When installing the unit, always make the ground connection first and disconnect it at the end.
- Do not block the ventilation while the device is operating. Keep a minimum distance of 5 cm between the device and the wall or other objects that may block the ventilation.
- Tighten the thumbscrews by using a tool after initial installation and subsequent access to the panel.

1.2 Labels on Devices

This section describes the warning labels, grounding labels, and electrostatic discharge (ESD) labels on devices.

Table 1-1 describes the safety labels on devices.

Table 1-1 Safety label

Label	Meaning	Description
	Warning	This label indicates that wrong operations may cause device damage or human injury.

Label	Meaning	Description
	External grounding	This label indicates grounding of external devices. The terminals of the ground cable are connected to different devices. This ensures normal running of the devices and the safety of the operator.
	Internal grounding	This label indicates grounding of internal devices. The terminals of the ground cable are connected to different components of the same device. This ensures normal running of the devices and the safety of the operator.
	ESD-Preventive Label	This label indicates a static sensitive area. Do not touch the device with your hands. When operating the device within this area, take electrostatic discharge (ESD)-preventive measures. For example, wear an ESD-preventive wrist strap.

1.3 Electrical Safety

This section describes safety notice about high voltage, thunderstorm, high electrical leakage, power cable, fuse, electrostatic discharge.

High Voltage

⚠ CAUTION

The high voltage power supply provides power for the device operation. Direct or indirect contact (through damp objects) with high voltage and AC mains supply may result in fatal danger.

Non-standard and improper high voltage operations may result in fire and electric shock.

Thunderstorm

⚠ CAUTION

Do not perform any operation, including high voltage and AC operations, on a steel tower or mast during a thunderstorm.

High Electrical Leakage

⚠ CAUTION

Ground the device before powering it on. Otherwise, personal injury or device damage may be caused by high leakage current.

If a "high electrical leakage" tag is present on the power terminal of the device, you must ground the device before powering it on.

Power Cable

⚠ CAUTION

Do not install or remove power cables when the device is on. Transient contact between the core of the power cable and the conductor may generate electric arcs or sparks, which may cause fire or hurt human eyes.

- Before installing or removing the power cable, turn off the power switch.
- Before connecting a power cable, check that the label on the power cable is correct.

Fuse

⚠ CAUTION

If a fuse is to be replaced, the new fuse must be of the same type and specifications.

Electrostatic Discharge

⚠ CAUTION

The static electricity generated by human bodies may damage the electrostatic-sensitive components on boards, for example, the large-scale integrated (LSI) circuits.

Human body movement, friction between human bodies and clothes, friction between shoes and floors, or handling of plastic articles causes static electromagnetic fields on human bodies. These static electromagnetic fields cannot be eliminated until the static is discharged.

To prevent electrostatic-sensitive components from being damaged by the static on human bodies, you must wear a well-grounded ESD wrist strap when touching the device or handling boards or application-specific integrated circuits (ASICs).

Figure 1-1 shows how to wear an ESD wrist strap.

Figure 1-1 Wearing an ESD wrist strap

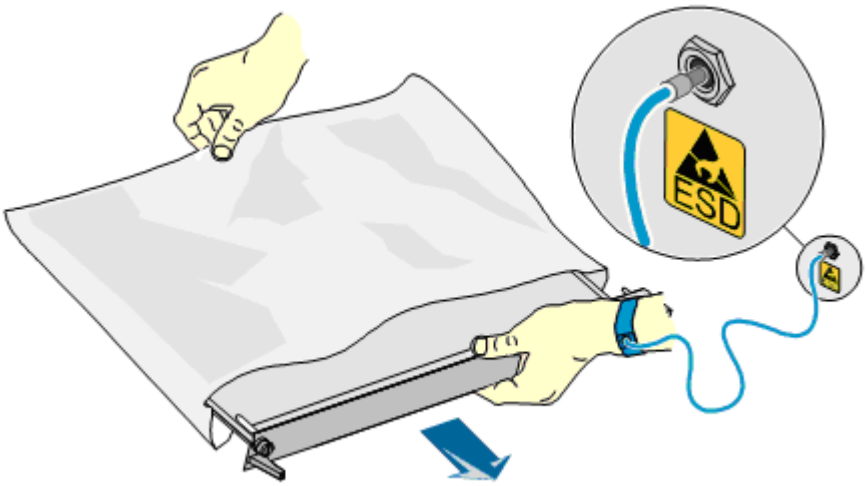

1.4 Certification

ES3000 V5 SSD Disk

RoHS(China), RoHS(Europe), REACH, WEEE, CE, FCC, IC, NRTL-UL, VCCI, RCM, KSA, KUCAS, SONCAP, UNBS, ACAP, CB, KCC, EAC, PCI-SIG, and NVMe protocol consistency

この装置は、クラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。 VCCI－B

ES3000 V5 SSD Card

RoHS(China), RoHS(Europe), REACH, WEEE, CE, FCC, IC, NRTL-UL, VCCI, RCM, KSA, KUCAS, SONCAP, UNBS, ACAP, CB

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI－A

Warning

This is a class A product, operation in a residential environment this product may cause radio interference in which case the user may be required to take adequate measures.

2 Getting to Know the ES3000 V5

About This Chapter

This document describes the features of ES3000 V5.

2.1 Security

This section describes the information security features of the ES3000 V5.

Huawei is committed to ensuring stability and security of customer networks and services.

As a storage device within a server or storage host, the ES3000 V5 provides the following security features based on the security hardening measures taken according to the minimum permission security rules:

- Rights control: The ES3000 V5 allows only system administrators to use its device maintenance management tools.
- Operation logs: The tool logs all the set operations on the ES3000 V5.
- Security assurance: A complete data clearance tool is provided to destroy data at the NAND flash memory chip level, preventing data leakage.
- Firmware upgrade: Firmware upgrade through in-band management channels ensures upgrade reliability and network security. The upgrade tool verifies the integrity of the ES3000 V5 firmware upgrade file to ensure that the upgrade is successful. Firmware can be loaded only after it is authenticated by the Hi1812E controller, which ensures firmware validity and data security.
- Device management: The ES3000 V5 uses out-of-band management based on I2C channels and in-band management based on SNMP. Both management functions support only query, preventing malicious and illegal write operations.

NOTE

To use the management tool, obtain the administrator rights. For example, the Windows administrator group must run as administrator.

2.2 Overview

This section describes the ES3000 V5 appearance and standards.

Huawei ES3000 V5 series NVMe PCIe SSD is an enterprise-class high-performance storage and acceleration component. It features high performance, fast response, and high reliability, greatly improving storage I/O performance. Two products are included: ES3000 V5 NVMe PCIe SSD disk and ES3000 V5 NVMe PCIe SSD card. The disk consists of the ES3500P V5 and ES3600P V5 series. The card consists of the ES3600C V5 series. The ES3000 V5 NVMe PCIe SSD consists of the ES3500P V5 and ES3600P V5 series.

The ES3000 V5 NVMe PCIe SSD disk is a 2.5-inch PCIe 3.0 x 4 SSD disk that uses the HiSilicon Hi1812E chip and complies with the SFF-8639 standard. This SSD disk can be installed in universal servers, such as rack servers.

Figure 2-1 shows the appearance of the ES3000 V5 NVMe PCIe SSD disk.

Figure 2-1 ES3000 V5 NVMe PCIe SSD disk

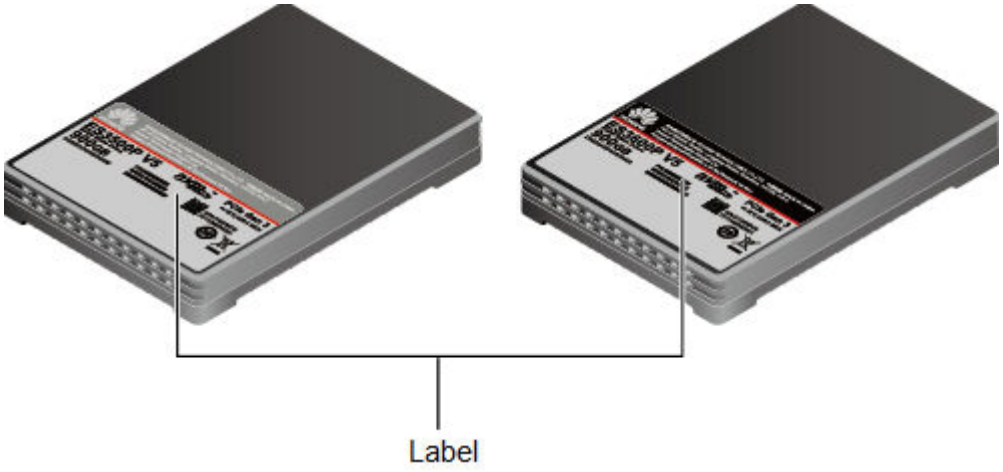

Label

The label on an ES3000 V5 NVMe PCIe SSD disk provides the World Wide Name (WWN) and serial number (S/N).

A WWN uniquely identifies an ES3000 V5 NVMe PCIe SSD disk and is a string generated from the MAC address. A WWN is required when you apply for technical support from Huawei.

An example S/N is 032VUE10G3000001. **Figure 2-2** describes the S/N format.

Figure 2-2 Example S/N

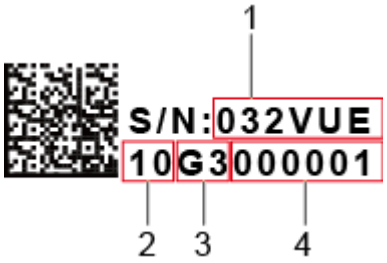

No.	Description
1	Indicates the material identification code (6 characters).
2	Indicates the vendor code (two characters).
3	Indicates the year and month (two characters).
4	Indicates the serial number (six characters).

The ES3000 V5 NVMe PCIe SSD card is a standard half-height half-length (HHHL) PCIe 3.0 x 4 card. It can be installed on universal servers, such as rack servers.

Figure 2-3 shows the appearance of the ES3000 V5 NVMe PCIe SSD.

Figure 2-3 ES3000 V5 NVMe PCIe SSD card

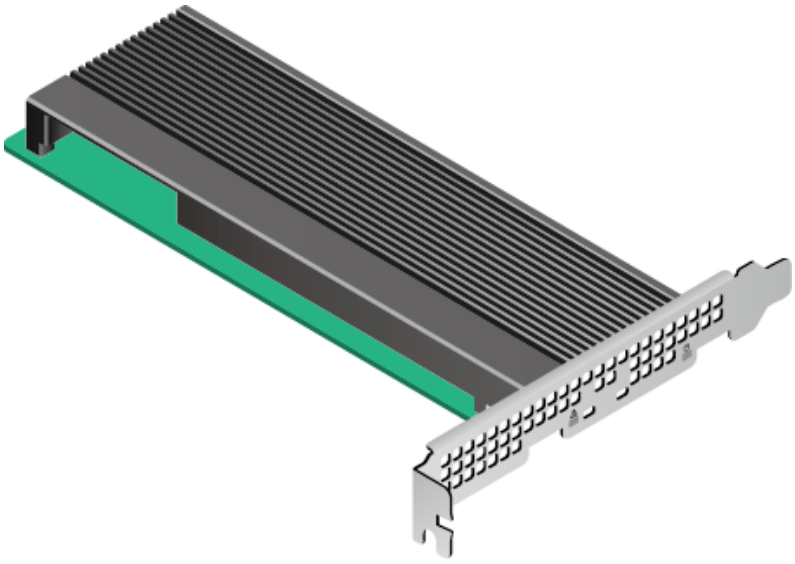

2.3 Features

The ES3000 V5 NVMe PCIe SSD provides the following NVMe features:

- Huawei ES3000 V5 SSD supports the NVMe 1.3 specification.

- Besides the basic feature and the mandatory commands, it also supports Name Space feature, a maximum of 64 Name Spaces are supported.

The ES3000 V5 NVMe PCIe SSD provides the following features:

- Uses the PCIe 3.0 x4 interface, complies with PCIe 3.0 electrical specifications, and is backward compatible with PCIe 2.0 and 1.0.
- Provides various single-drive capacities to meet different application requirements.
  - ES3500P V5 series single-drive capacities: 1000 GB, 2000 GB, 4000 GB, and 8000 GB
  - ES3600P V5 series single-drive capacities: 800 GB, 1600 GB, 3200 GB, and 6400 GB
  - ES3600C V5 series single-card capacities: 800 GB, 1600 GB, 3200 GB, 6400 GB
- Supports standard storage device operations.
  - Block device operations in Linux
  - Drive device operations in Windows
  - Drive device operations in ESXi
- Supports device diagnosis and detecting, and SMART information reporting.
  - Checks the power supply, temperature, voltage, capacitance, and number of bad blocks.
  - Checks whether the flash component is normal.
  - Monitors, queries, and warns about the SSD service life.
  - Checks the SSD power consumption.
  - Performs I/O statistics.
- ES3500P V5 and ES3600P V5 series support orderly hot swap, and support surprise hot swap on some V5 servers only. ES3600C V5 series support orderly hot swap on some V5 servers.
- Uses multiple technologies to ensure device reliability.
  - Enhanced Low-Density Parity-Check (LDPC) algorithm: provides higher error correction capability than that required by flash chips and ensures device reliability.
  - Embedded RAID-like algorithm: implements channel-based error correcting. Data can be restored if an error occurs in a channel.
  - Flexible RAID algorithm: automatically restores data on the flash memory when the flash memory has failed and continues to perform RAID data protection.
  - Intelligent wear leveling algorithm: intelligently levels the flash chip wear pressure and prolongs the device service life.
  - Advanced flash access technology: combines the read retry and adaptive read technologies of flash chips and ensures data validity.
  - Data inspection technology: periodically inspects data and prevents errors.
  - Power-off protection: prevents drive data loss when a power failure occurs on the server.
  - Provides end-to-end data protection through the DIF domain, and data check for every data transmission between modules within the SSD.
- Provides multiple maintenance and management tools with comprehensive functions and high performance.

- Provides in-band online upgrades to facilitate routine maintenance.
- Provides out-of-band management functions that comply with NVMe standards to implement real-time detecting of model, capacity, temperature, life, and health status on GUIs, facilitating management.
- Provides comprehensive command line management functions for users to perform routine device management by using various commands.
- Provides the manufacture dates and serial numbers to facilitate asset management.

2.4 Specifications

This section describes the specifications of the ES3000 V5.

The ES3000 V5 comes in two form factors: disk and card. The disk consists of the ES3500P V5, ES3600P V5. The card consists of the ES3600C V5 series.The ES3000 V5 NVMe PCIe SSD consists of the ES3500P V5 and ES3600P V5 series.The default power consumption level 0 ensures optimal performance.

A non-default power consumption level is not recommended.

NOTE

Logical Block Addressing (LBA) is a general mechanism used to describe the block of data on a storage device. It is generally used on drives. An LBA format indicates an LBA size of the LV as which an NVMe SSD is formatted. An LBA block has a variety of sizes, which may be 512 bytes, 4096 bytes, or the like.

The ES3000 V5 can be formatted into multiple LBA formats based on service requirements, such as 512 bytes and 4096 bytes. ES3000 V5 I/O performance varies according to LBA formats. The 4096 byte format delivers the best performance.

2.4.1 ES3500P V5 Specifications

Nameplates

Figure 2-4 shows the nameplate for the ES3500P V5 series.

Figure 2-4 ES3500P V5 series nameplate

Power Consumption

Table 2-1 lists the power consumption specifications for the ES3500P V5 series.

Table 2-1 Power consumption specifications for the ES3500P V5 series

Parameter	Specifications
Power consumption	● 1000 GB: 13 W maximum power consumption[a] and 5 W standby power consumption. ● 2000 GB: 13.8 W maximum power consumption and 5 W standby power consumption. ● 4000 GB: 18.8 W maximum power consumption and 5.5 W standby power consumption. ● 8000 GB: 21 W maximum power consumption and 6.5 W standby power consumption.
Form factor	Standard 2.5-inch drive
a: The maximum power consumption is the test result with the maximum sequential write bandwidth.	

Power consumption specifications in the preceding table are test results based on CentOS 7.0 and the standard storage test tool fio 2.6, with the ES3500P V5 formatted into a 512-byte LBA block device.

Reliability Specifications

Table 2-2 lists the reliability specifications for the ES3500P V5 series.

Table 2-2 Reliability specifications for the ES3500P V5 series

Parameter	Specifications
Mean time between failures (MTBF)	2.5×10^6 hours
Annual failure rate (AFR)	≤ 0.35%
Bit error rate (BER)	10^{-18}
Data protection	≥ 3 months (40°C, powered off)

Environmental Specifications

Table 2-3 lists the environmental specifications for the ES3500P V5 series.

Table 2-3 Environmental specifications for the ES3500P V5 series

Parameter	Specifications
Temperature	● Storage temperature: –40°C to +85°C (–40°F to +185°F) ● Drive case temperature: 0°C to 70°C (32°F to 158°F)
Altitude	● Operating: –305 m to +5486 m (–1000.64 ft to +17998.47 ft) ● Non-operating: –305 m to 12192 m (–1000.64 ft to 39999.51 ft) The highest operating temperature decreases by 1°C (1.8°F) for every increase of 220 m (721.78 ft) in altitude.
Air speed	1000 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	2000 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	4000 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	8000 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
Humidity	● Storage humidity: 5% RH to 95% RH ● Operating humidity: 5% RH to 95% RH
Vibration	● Operating vibration: 2.17 G_{RMS} (5 Hz to 700 Hz) ● Non-operating vibration: 3.13 G_{RMS} (5 Hz to 800 Hz)
Shock	● Operating shock: 1000 G/0.5 ms ● Non-operating shock: 1000 G/0.5 ms

2.4.2 ES3600P V5 Specifications

Nameplates

Figure 2-5 shows the nameplate for the ES3600P V5 series.

Figure 2-5 ES3600P V5 series nameplate

Power Consumption

Table 2-4 lists the power consumption specifications for the ES3600P V5 series.

Table 2-4 Power consumption specifications for the ES3600P V5 series

Parameter	Specifications
Power consumption	● 800 GB: 13 W maximum power consumption[a] and 5 W standby power consumption. ● 1600 GB: 13.8 W maximum power consumption and 5 W standby power consumption. ● 3200 GB: 18.8 W maximum power consumption and 5.5 W standby power consumption. ● 6400 GB: 21 W maximum power consumption and 5.7 W standby power consumption.
Form factor	Standard 2.5-inch drive
a: The maximum power consumption is the test result with the maximum sequential write bandwidth.	

Power consumption specifications in the preceding table are test results based on CentOS 7.0 and the standard storage test tool fio 2.6, with the ES3600P V5 formatted into a 512-byte LBA block device.

Reliability Specifications

Table 2-5 lists the reliability specifications for the ES3600P V5 series.

Table 2-5 Reliability specifications for the ES3600P V5 series

Parameter	Specifications
Mean time between failures (MTBF)	2.5 x 10^6 hours
Annual failure rate (AFR)	≤ 0.35%
Bit error rate (BER)	10^{-18}
Data protection	≥ 3 months (40°C, powered off)

Environmental Specifications

Table 2-6 lists the environmental specifications for the ES3600P V5 series.

Table 2-6 Environmental specifications for the ES3600P V5 series

Parameter	Specifications
Temperature	● Storage temperature: – 40°C to +85°C (– 40°F to +185°F) ● Drive case temperature: 0°C to 70°C (32°F to 158°F)
Altitude	● Operating: – 305 m to +5486 m (– 1000.64 ft to +17998.47 ft) ● Non-operating: – 305 m to +12192 m (– 1000.64 ft to +39999.51 ft) The highest operating temperature decreases by 1°C (1.8°F) for every increase of 220 m (721.78 ft) in altitude.
Air speed	800 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	1600 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at a ambient temperature of 55°C (131°F)
	3200 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	6400 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
Humidity	● Storage humidity: 5% RH to 95% RH ● Operating humidity: 5% RH to 95% RH
Vibration	● Operating vibration: 2.17 G_{RMS} (5 Hz to 700 Hz) ● Non-operating vibration: 3.13 G_{RMS} (5 Hz to 800 Hz)

Parameter	Specifications
Shock	● Operating shock: 1000 G/0.5 ms ● Non-operating shock: 1000 G/0.5 ms

2.4.3 ES3600C V5 Specifications

Nameplates

Figure 2-6 shows the nameplate for the ES3600C V5 series.

Figure 2-6 ES3600C V5 series nameplate

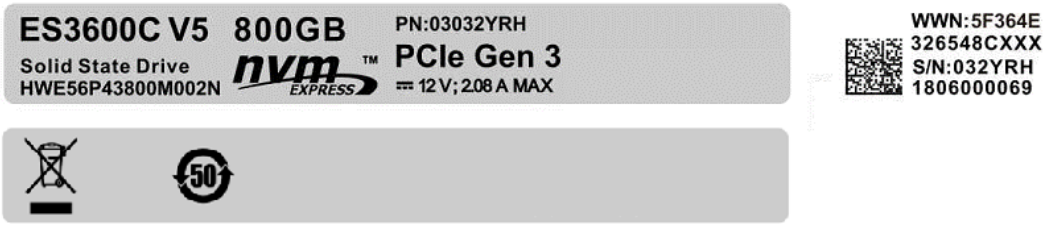

Power Consumption

Table 2-4 lists the power consumption specifications for the ES3600C V5 series.

Table 2-7 Power consumption specifications for the ES3600C V5 series

Parameter	Specifications
Power consumption	● 800 GB: maximum power consumption[a]: 13 W; standby power consumption: 5 W ● 1600 GB: maximum power consumption: 13.8 W; standby power consumption: 5 W ● 3200 GB: maximum power consumption: 18.8 W; standby power consumption:5.5 W ● 6400 GB: maximum power consumption: 21 W; standby power consumption:5.7 W
Form factor	Standard HHHL PCIe card. Dimensions: 167.5 mm × 68.9 mm × 18.6 mm (6.6 in. x 2.7 in. x 0.7 in.)
a: The maximum power consumption is the test result with the maximum sequential write bandwidth.	

Power consumption specifications in the preceding table are test results based on CentOS 7.0 and the standard storage test tool fio 2.6, with the ES3600P V5 formatted into a 512-byte LBA block device.

Reliability Specifications

Table 2-5 lists the reliability specifications for the ES3600C V5 series.

Table 2-8 Reliability specifications for the ES3600C V5 series

Parameter	Specifications
Mean time between failures (MTBF)	2.5 x 10^6 hours
Annual failure rate (AFR)	≤ 0.35%
Bit error rate (BER)	10^{-18}
Data protection	≥ 3 months (40°C, powered off)

Environmental Specifications

Table 2-6 lists the environmental specifications for the ES3600C V5 series.

Table 2-9 Environmental specifications for the ES3600C V5 series

Parameter	Specifications
Temperature	● Storage temperature: – 40°C to +85°C (– 40°F to +185°F) ● Drive case temperature: 0°C to 70°C (32°F to 158°F)
Altitude	● Operating: – 305 m to +5486 m (– 1000.64 ft to +17998.47 ft) ● Non-operating: – 305 m to +12192 m (– 1000.64 ft to +39999.51 ft) The highest operating temperature decreases by 1°C (1.8°F) for every increase of 220 m (721.78 ft) in altitude.
Air speed	800 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	1600 GB ● > 155 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at a ambient temperature of 55°C (131°F)
	3200 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)
	6400 GB ● > 237 LFM at an ambient temperature of 40°C (104°F) ● > 650 LFM at an ambient temperature of 55°C (131°F)

Parameter	Specifications
Humidity	● Storage humidity: 5% RH to 95% RH ● Operating humidity: 5% RH to 95% RH
Vibration	● Operating vibration: 2.17 G_{RMS} (5 Hz to 700 Hz) ● Non-operating vibration: 3.13 G_{RMS} (5 Hz to 800 Hz)
Shock	● Operating shock: 1000 G/0.5 ms ● Non-operating shock: 1000 G/0.5 ms

2.5 System Requirements

This section describes the ES3000 V5 requirements for server software and hardware.

Hardware Requirements

The minimum requirements for server hardware are as follows:

- 2 GB memory
- To use the ES3000 V5 NVMe PCIe SSD disk, the server must provide at least one 2.5-inch SSD disk slot that supports the PCIe SFF-8639 connector.
- To use the ES3000 V5 NVMe PCIe SSD card, the server must provide at least one HHHL (167.5 mm x 70 mm x 18.6 mm) PCIe 3.0 x4 card slot.

Software Requirements

- Use the **Huawei Server Compatibility Checker** to check the OS types and versions supported by the ES3000 V5.
- OSs listed in this section support common I/O services. To obtain the list of OSs that support hot swap, see **4.1 OSs Supporting Hot Swap**. To obtain the list of OSs that can be installed on the ES3000, see **G Installing an OS on the ES3000**.

3 Installation and Configuration

About This Chapter

This section describes the process for installing and configuring the ES3000 V5.

CAUTION

If the ES3000 V5 exceeds the overtemperature threshold, handle the problem promptly; otherwise, management operations on the SSD may fail. For details about the command for querying temperature information, see 6.8.1 Querying the Temperature Information About a Device.

3.1 Installation and Configuration Process

The installation and configuration process includes the overall procedures for installing and configuring the ES3000 V5. You can learn about ES3000 V5 installation and configuration logic.

The flowchart for installing and configuring ES3000 V5 is shown in **Figure 3-1**.

Figure 3-1 Installation and configuration flowchart

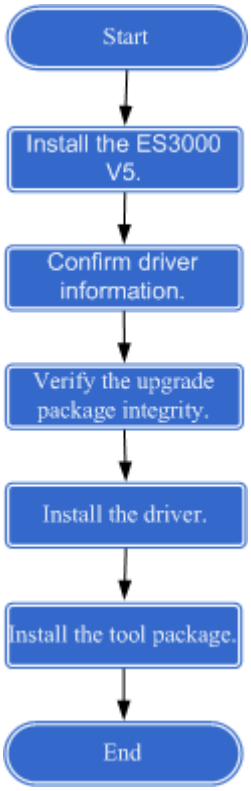

3.2 Installing the ES3000 V5

Operation Scenario

The method of installing the ES3000 V5 NVMe PCIe SSD disk is the same as the method of installing a common hard disk. For details, see **3.2.1 Installing the ES3000 V5 NVMe PCIe SSD Disk**.

If the PCIe card slots are on a riser card, the ES3000 V5 NVMe PCIe SSD card needs to be installed on the riser card. For details, see **3.2.2 Installing the ES3000 V5 NVMe PCIe SSD Card on a Riser Card**.

If the PCIe card slots are on the mainboard, the ES3000 V5 NVMe PCIe SSD card needs to be installed on the mainboard. For details, see **3.2.3 Installing the ES3000 V5 NVMe PCIe SSD Card on the Mainboard**.

Preparations

- Stop all services on the server.
- Unpack and inspect the devices.

  Check that the items shipped are complete and intact based on the following list:
  - ES3000 V5 NVMe PCIe SSD
  - Quick Start Guide
  - Warranty card
- Tools

- ESD gloves

  Used to prevent ESD damage when you insert, remove, or hold the ES3000 V5 or a precision device.

CAUTION

Pay attention to the following in Linux:

- Do not perform direct read or write operations on a raw SSD to avoid processing the wrong SSD due to a changed drive letter. Restarting the server, installing or uninstalling the SSD driver, inserting and removing the SSD may change the SSD drive letter.
- If you have to perform direct read or write operations on a raw SSD, use the **hioadm info** command to check the SSD information and confirm the drive letter through the SSD SN.
- It is recommended that you use the **mount** command to mount the SSD partitions to a specified path before performing read and write operations. To ensure normal release of system resources, use the unmount command to cancel the mounted SSD partitions before installing or uninstalling the driver, and inserting or removing the SSD.
- You can also add information to the **/etc/fstab** file so that SSD partitions are automatically mounted during system startup, driver installation and uninstallation, and SSD insertion and removal. For details, run the **man fstab** command to view help information. After enabling the automatic mounting of SSD partitions, you do not need to run the **umount** command before installing or uninstalling the driver, and inserting or removing the SSD.

Follow-up Procedure

After the ES3000 V5 is connected to power, the startup process takes about 30 seconds.

If the activity indicator associated with the ES3000 V5 NVMe PCIe SSD disk is on and the fault indicator is off, the SSD hardware is installed successfully. If an installation error occurs, see the **Huawei Servers Troubleshooting**.

If the green indicator (active indicator) on the ES3000 V5 NVMe PCIe SSD card is on and the orange indicator (fault indicator) is off, the SSD card hardware is installed successfully.

3.2.1 Installing the ES3000 V5 NVMe PCIe SSD Disk

Step 1 Wear ESD gloves.

Step 2 Take the ES3000 V5 out of the ESD bag. Check whether a 2.5-inch or 3.5-inch drive tray is used.

- For a 2.5-inch drive tray, go to **Step 3**.
- For a 3.5-inch drive tray, go to **Step 5**.

Step 3 Place the ES3000 V5 into the drive tray. See (1) in **Figure 3-2**.

Step 4 Use four screws to secure the ES3000 V5 to the drive tray. See (2) in **Figure 3-2**. Then go to **Step 8**.

Figure 3-2 Placing the ES3000 V5 into the drive tray

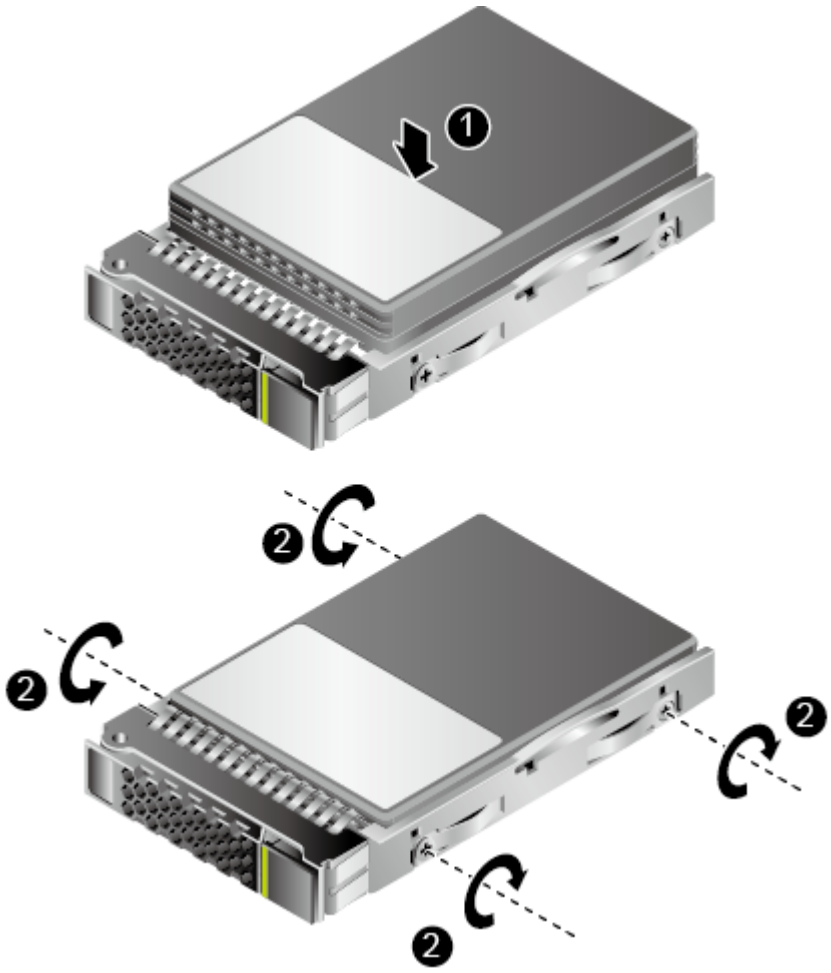

Step 5 Place the ES3000 V5 to the converter. See (1) in **Figure 3-3**.

Step 6 Use two screws to secure the ES3000 V5 to the converter. See (2) in **Figure 3-3**.

Figure 3-3 Securing the ES3000 V5 to the converter

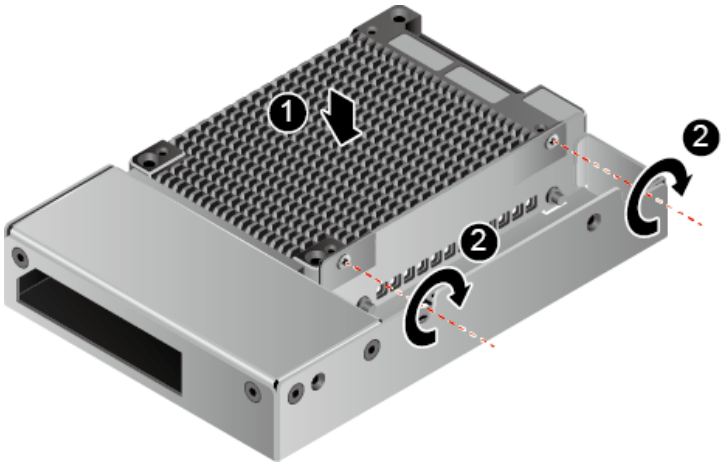

Step 7 Use four screws to secure the ES3000 V5 to the drive tray. See (3) in **Figure 3-4**.

Figure 3-4 Securing the ES3000 V5 to the drive tray

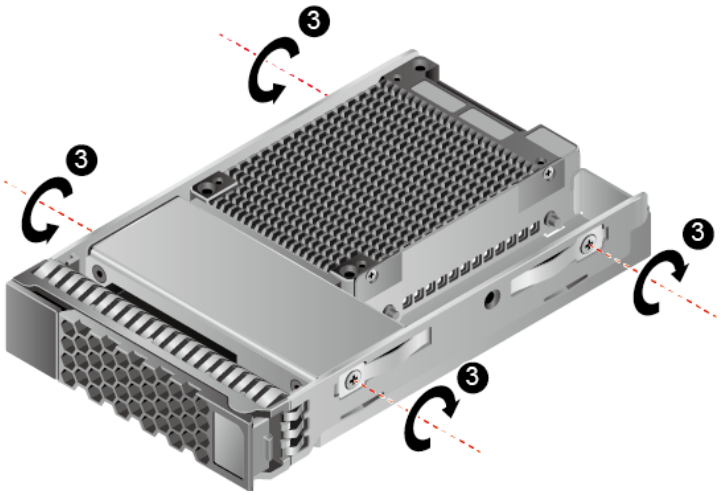

Step 8 Open the ejector lever and push the ES3000 V5 into the server until it does not move. See (1) in **Figure 3-5**.

Step 9 Close the ejector lever to push the ES3000 V5 into place. See (2) in **Figure 3-5**.

Figure 3-5 Installing the ES3000 V5

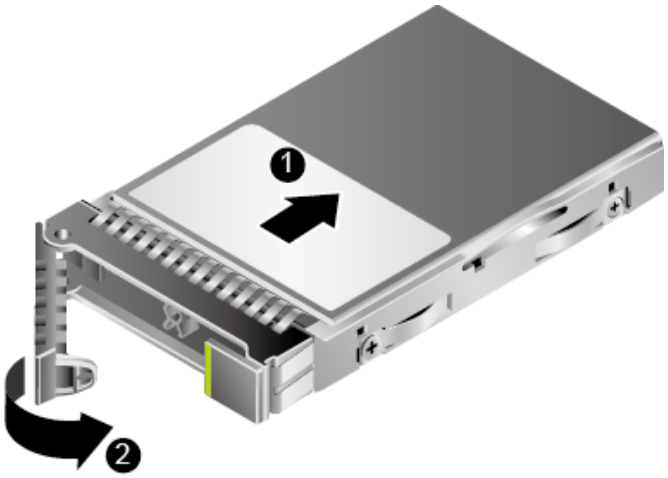

----End

3.2.2 Installing the ES3000 V5 NVMe PCIe SSD Card on a Riser Card

Step 1 Wear ESD gloves.

Step 2 Turn off the power to the server and remove the chassis cover.

Step 3 Take the new ES3000 V5 out of the ESD bag.

Step 4 Hold the upper edge of the ES3000 V5, align its connecting part with a PCIe slot on the riser card, and insert the ES3000 V5 slowly into the PCIe slot. See **Figure 3-6**.

Figure 3-6 Installing the ES3000 V5 on a riser card

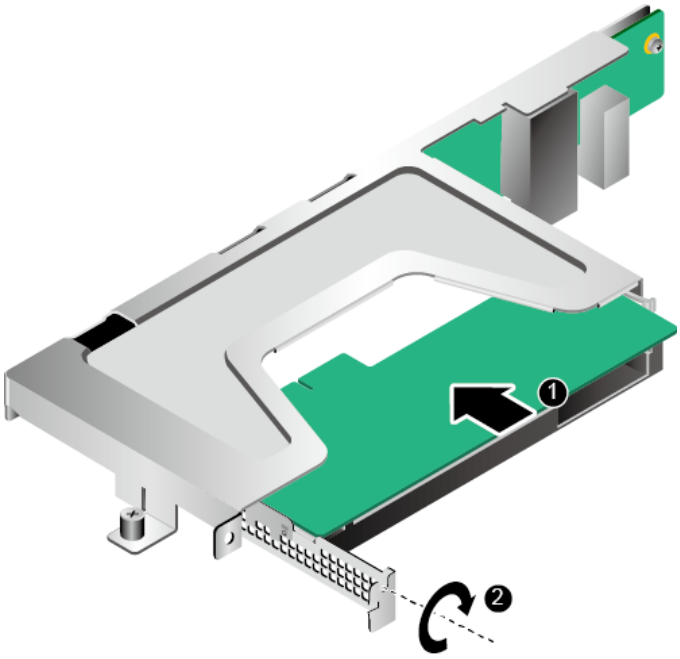

Step 5 Install the riser card to the server. For details, see related server manuals.

Step 6 Turn on the power to the server.

----End

3.2.3 Installing the ES3000 V5 NVMe PCIe SSD Card on the Mainboard

Step 1 Wear ESD gloves.

Step 2 Turn off the power to the server and remove the chassis cover.

Step 3 Hold the upper edge of the ES3000 V5, align its connecting part with the PCIe slot, and vertically insert the ES3000 V5 downwards into the PCIe slot slowly. See **Figure 3-7**.

Figure 3-7 Installing the ES3000 V5 on the mainboard

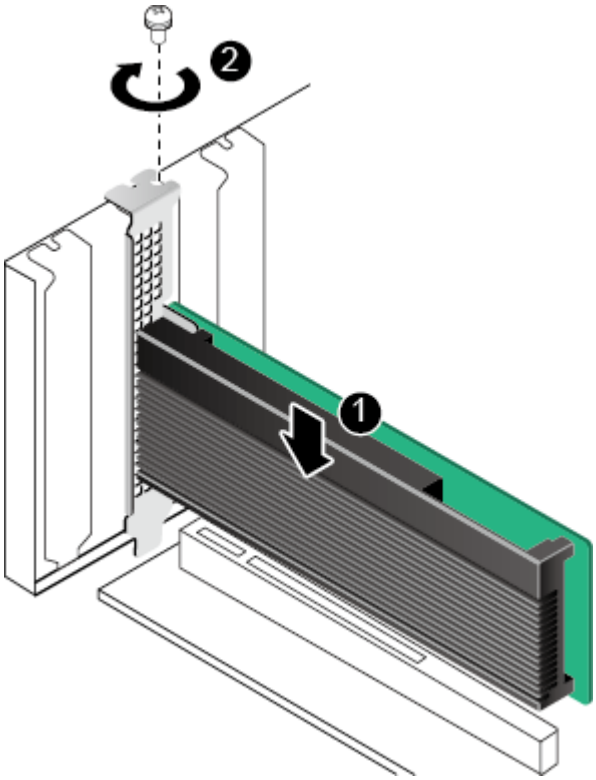

Step 4 Turn on the power to the server.

----End

3.3 Confirming Driver Information

This section describes how to obtain the ES3000 V5 software package.

Confirm the scenario before downloading the driver.

- If the server uses only ES3000 V5 NVMe PCIe SSD and does not use SSD provided by other vendors. You can use either the NVMe driver integrated in the OS or that provided by Huawei. For details about how to install Huawei's NVMe driver, see **3.6 Installing the Driver**.
- If the server running Linux uses both ES3000 V5 NVMe PCIe SSD and SSD provided by other vendors, use the NVMe driver integrated in the OS and you do not need to install Huawei's NVMe driver.
- If Huawei's NVMe driver has been installed in Linux and you need to use SSD provided by other vendors, uninstall Huawei's NVMe driver.
- For Windows Server 2008 R2 or Windows Server 2012, if the **KB2990941** patch has been installed, the OS is integrated with the NVMe driver. To install Huawei's NVMe driver, uninstall the **KB2990941** patch first. For details, see **3.6.2 Installing the ES3000 V5 Driver in Windows**.
- The following OSs are also integrated with the NVMe driver:
  - Windows Server 2012 R2 or later
  - ESXi 6.0 or later

- RHEL 6.5 or later
- CentOS 6.5 or later
- OEL 6.5 or later
- Ubuntu 14.04 or later
- SUSE Linux Enterprise Server (SLES) 11 SP4 or later

You can download drivers as follows: log in to Huawei Enterprise support website **http://support.huawei.com/enterprise**, choose **Products** > **Cloud Data Center** > **Servers** > **Intelligent Accelerator Components** > **ES3000 V5 SAS SSD Disk** > **Technical Support**", click **Download Software**, and download the required software package. The software package contains host tools (including hioadm) and Huawei's NVMe driver.

3.4 Verifying the Upgrade Package Integrity

Verify that the obtained upgrade package is consistent with the original one at the website http://support.huawei.com/enterprise.

Click  and  , load the digital certificate and software package respectively.

For details about the verification tool and method, visit the following link:

Software digital signature (Open Verify) validation tool

3.5 Transferring Files by Using WinSCP

This section describes how to transfer files by using WinSCP. You can also transfer files by using other software.

Scenarios

Use WinSCP to transfer files from a PC to a remote host.

Prerequisites

Conditions

The File Transfer Protocol (FTP) service has been enabled on the remote host.

Data

You have obtained the following data:

- IP address of the remote host to be connected
- User name and password for logging in to the remote host to be connected

Software

WinSCP.exe: free software available on the Internet.

Procedure

Step 1 Open the **WinSCP** folder, and double-click **WinSCP.exe**.

The **WinSCP Login** dialog box is displayed, as shown in **Figure 3-8**.

NOTE

To change the UI language, click **Languages**.

Figure 3-8 WinSCP login

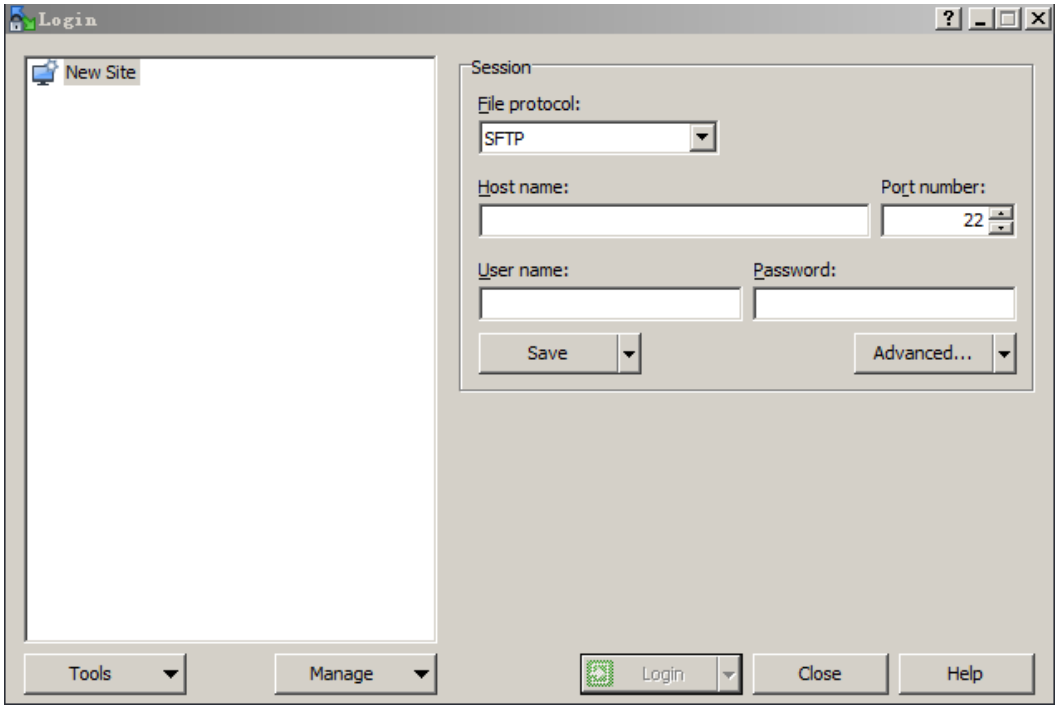

Step 2 Set the login parameters.

The parameters are described as follows:

- **Host name**: Enter the IP address of the remote host to be connected. For example, **191.100.34.32**.
- **Port number**: The default value is **22**.
- **User name**: Enter the user name. For example, **admin**.
- **Password**: Enter the password.

Step 3 Click **Login**.

The **WinSCP** file transfer window is displayed.

NOTE

- If a key file was not selected during the first login, a warning **Continue connecting and add host key to cache** is displayed. Click **Yes**. The **WinSCP** file transfer page is displayed.
- By default, the **C:\Documents and Settings\Administrator\My Documents** directory of the local PC is opened in the left pane, and the **/root** directory of the remote host is opened in the right pane.

Step 4 In the right pane, choose a directory of the remote host for storing files, for example, **/root/firmware**.

Step 5 In the left pane, select a directory of the local PC for storing files, for example **E:\Software**, and then select the file to be transferred.

Step 6 Choose **File** > **Copy**.

The **Copy** dialog box is displayed.

Step 7 Confirm the file to be copied and the directory of the remote host, and then click **Copy**.

The system starts to copy the file from the local PC to the remote host.

----End

3.6 Installing the Driver

This section describes how to install the ES3000 V5 driver in different operating systems (OSs). The ES3500P V5 and the ES3600P V5 and the ES3600C V5 series use the same drivers and firmware.

3.6.1 Installing the ES3000 V5 Driver in Linux

Operation Scenario

If you decide to use Huawei's NVMe driver after reading **3.3 Confirming Driver Information**, install it by referring to this section before using the ES3000 V5 for the first time.

If you install multiple ES3000 V5s on a server, you need to install the driver only once.

NOTE

The Linux driver package naming rules are as follows:

- RHEL, CentOS, or OEL driver package name: **kmod-hiodriver-Driver version-Supported kernel version.System distributor.Hardware platform.rpm**, such as **kmod-hiodriver-5.0.0.6-3.10.0_123.el7.centos.x86_64.rpm**, **kmod-hiodriver-5.0.0.6-3.10.0_123.el7.redhat.x86_64.rpm**, or **kmod-hiodriver-5.0.0.6-3.10.0_123.el7.oracle.x86_64.rpm**.
- SLES driver package name: **hiodriver-kmp-Kernel feature-Driver version-Supported kernel version.System distributor.Hardware platform.rpm**, such as **hiodriver-kmp-default-5.0.0.6-3.12.28_4.suse12.x86_64.rpm**.
- Ubuntu or Debian driver package name: **hiodriver -Driver version-Supported kernel version.System distributor.Hardware platform.deb**, such as **hiodriver-5.0.0.6-3.19.0_15_generic.ubuntu.amd64.deb**.
- UVP driver package name: **hiodriver-kmp-xen-Driver version-Supported kernel version.System distributor.Hardware platform.rpm**.
- Citrix driver software package: **hiodriver-citrix-**<*system version*>**-**<*version number*>**.iso**, for example, **hiodriver-citrix-6u2-5.0.0.8.iso**.

Preparations

Before installing the ES3000 V5 driver, check that:

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the ES3000 V5 driver package to the server, and verified the package integrity.

NOTE

If the system version is Citrix 6.2, use the 32-bit driver package, for example, **i386**. If the system version is Citrix 6.5 or 7.1, use a 64-bit driver package, for example, **x86_64**.

- You have uploaded the driver package to a directory (for example, **/home**) on the server by using WinSCP.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Go to the directory that stores the ES3000 V5 driver package, for example, **/home**.

Step 3 Run the installation command based on the OS type.

- For RHEL, CentOS, OEL, SLES, UVP, or RHEL, run **rpm -ivh** *driver package name*.

  An example package name for CentOS is **rpm -ivh kmod-hiodriver-5.0.0.6-2.6.32_431.el6.centos.x86_64.rpm**.
- For Ubuntu or Debian, run **dpkg -i** *driver package name*.

  An example package name for Ubuntu is **dpkg -i hiodriver-5.0.0.6-3.19.0_15_generic.ubuntu.amd64.deb**.
- For Citrix:

  a. Check the version of Citrix
     - Citrix 7.1 and below, go to **b**
     - Citrix 7.1, go to **f**

  b. Run the **mount -o loop <*Driver package name*>** command to mount the ISO package to a directory, for example, **/tmp/citrix**.

     The following uses Citrix 6.2 as an example: **mount -o loop hiodriver-3.3.9.1-6u2.iso /tmp/citrix.**

  c. Access the mount directory where the ISO package is located, for example, **/tmp/citrix**.

     cd /tmp/citrix

  d. Run **install.sh** in **/tmp/citrix** to install the driver.

     ./install.sh

  e. Run the following command to query driver information:

     rpm -qa | grep hiodriver

  f. Run **xe update-upload file-name=<*Driver package name*>** to upload the update.

     xe update-upload file-name=hiodriver-citrix7u1-5.0.0.8.iso

     NOTE

     The UUID of the pack is returned when the upload completes.

  g. Apply the pack by **xe update-apply uuid=<pack_uuid>**

     xe update-apply uuid=320232df-7adb-4cbe-a7a3-8515240879e1

  h. Run the following command to query driver information.

     xe update-list

NOTE

- If the CPU hardware capabilities are insufficient for processing high-speed I/O services of NVMe devices during interruption, it is recommended that you configure the NVMe driver to work in threaded_irq mode before installing it. The configuration method is as follows:

  1. Run the **vim /etc/modprobe.d/nvme.conf** command to create or open the configuration file for editing.
  2. Enter **i**.
  3. In the configuration file, change the value of **options nvme use_threaded_interrupts** to **1**:

     options nvme use_threaded_interrupts=1
  4. Press **Esc**, enter **:wq**, and press **Enter** to save the edits.
  5. Run the **chmod +x /etc/modprobe.d/nvme.conf** command to change the permission on **nvme.conf**.

- If the NVMe driver works in threaded_irq mode, the I/O latency increases: Both the 4 KB read latency and write latency increase by 5 μs.

Step 4 Check whether the following message indicating an incorrect installation package is displayed.

```
Error: The installing package does not match the current OS. Please execute
command: rpm -e kmod-hiodriver to uninstall the installing rpm.
```

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Uninstall the installed package as prompted. Upload the correct driver package for the OS to the server and perform **Step 2** again to install the driver.

Step 6 Check whether the following message is displayed. The message indicates that the original NVMe driver fails to be uninstalled during the installation.

```
Warning: fail to uninstall the old nvme kernel module. Please uninstall the old
nvme kernel module manually or reboot the system.
```

- If yes, go to **Step 9**.
- If no, go to **Step 7**.

Step 7 Check whether the following message is displayed. The message indicates that the new NVMe driver fails to be installed during the installation.

```
Warning: fail to install the new nvme kernel module. Please install the new nvme
kernel module manually or reboot the system.
```

- If yes, go to **Step 8**.
- If no, go to **Step 12**.

Step 8 Check whether restarting the OS will affect services.

- If yes, restart the OS when no service is running and go to **Step 12**.
- If no, restart the OS and go to **Step 12**.

Step 9 Check whether restarting the OS will affect services.

- If yes, go to **Step 10**.
- If no, restart the OS and go to **Step 12**.

Step 10 Run the **modprobe -r nvme** command to uninstall the original driver. During the uninstallation, check whether a message is displayed indicating that the driver is in use. If yes, the uninstallation fails.

- If yes, restart the OS when no service is running and go to **Step 12**.
- If no, go to **Step 11**.

Step 11 Run the **modprobe nvme** command to install the new driver. Check whether a warning or error is displayed during the installation.

- If yes, restart the OS and go to **Step 12**.
- If no, go to **Step 12**.

Step 12 Run the **modinfo nvme** command to check the driver author information to determine whether the driver is successfully installed.

If the author in the command output is as follows, the driver is successfully installed.

```
author:         Huawei Technologies Co., Ltd.
```

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 If a PCIe SSD device has been installed on the server, check whether the driver has initialized the device.

Check whether the NVMe device exists in the **/dev** directory to determine whether the device initialization is complete.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 if the OS uses the irqbalance service to balance CPU interrupts, restart the service by using one of the following commands after re-loading the driver module:

- For RHEL 6 or CentOs 6, run **service irqbalance restart**.
- For SLES 11 or UVP, run **service irq_balancer restart**.
- For SLES 12 or RHEL 7, run **systemctl restart irqbalance.service**.

NOTE

- For RHEL 6, CentOS 6, and SLES 12, you can add IRQBALANCE_ARGS="--hintpolicy=exact" to the end of the **/etc/sysconfig/irqbalance** file content to better balance interrupts.
- To configure the irqbalance service in other OSs, refer to documents related to the OSs.

If the OS cannot find the service, it indicates that the OS does not use the irqbalance service for balance interrupts. If the command can be run, check whether errors are reported after running the command for restarting the irqbalance service.

- If yes, go to **Step 15**.
- If no, no further action is required.

Step 15 Contact Huawei technical support for help. For details, see **H Getting Help**.

NOTE

If the ES3000 V5 hiodriver is used, certain operating systems such as SUSE Linux will record the **module verification failed** error in dmesg logs. This error message is a notification indicating that an operating system detects a third-party driver with an unmatched signature. It does not affect services. For detailed reasons and solutions, see **"Why the module verification failed" Message Is Displayed When ES3000 V3 PCIe SSDs Use Huawei-developed SUSE NVMe Drivers**.

----End

3.6.2 Installing the ES3000 V5 Driver in Windows

Operation Scenario

If you decide to use Huawei's NVMe driver after reading **3.3 Confirming Driver Information**, install it by referring to this section before using the ES3000 V5 for the first time.

NOTE

- If you install multiple ES3000 V5s on a server, you need to install the driver only once.
- The ES3000 V5 tool is incompatible with the NVMe driver built in the Windows OS. Ensure that you have installed Huawei's NVMe driver before using the tool.
- To install Windows on the ES3000 V5, see ***G Installing an OS on the ES3000***.

Preparations

Before installing the ES3000 V5 driver, check that:

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the driver package and verified its integrity. For details, see **3.4 Verifying the Upgrade Package Integrity**.
- You have uploaded the driver package to the server by using WinSCP or other software.

Procedure

Select an operation based on the OS type.

- For Windows Server 2012 R2, go to **1** in **Install the driver**.
- For Windows Server 2008 R2 or Windows Server 2012, go to **1** in **Check whether the SSD driver is provided by the OS**.

Check whether the SSD driver is provided by the OS.

1. Log in to the OS as an administrator.
2. Right-click **Computer** and choose **Properties** from the shortcut menu. The **System** window is displayed, as shown in **Figure1 System**.

Figure 3-9 System

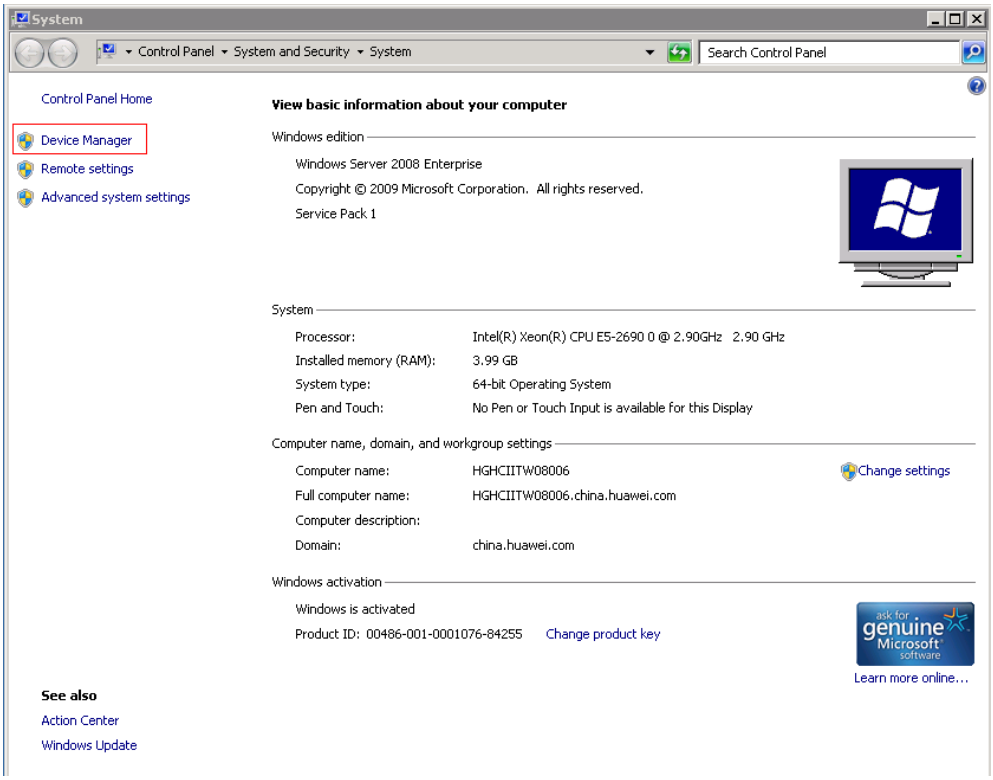

3. Click **Device Manager**. The **Device Manager** is displayed, as shown in **Figure 3-10**.

Figure 3-10 Device Manager

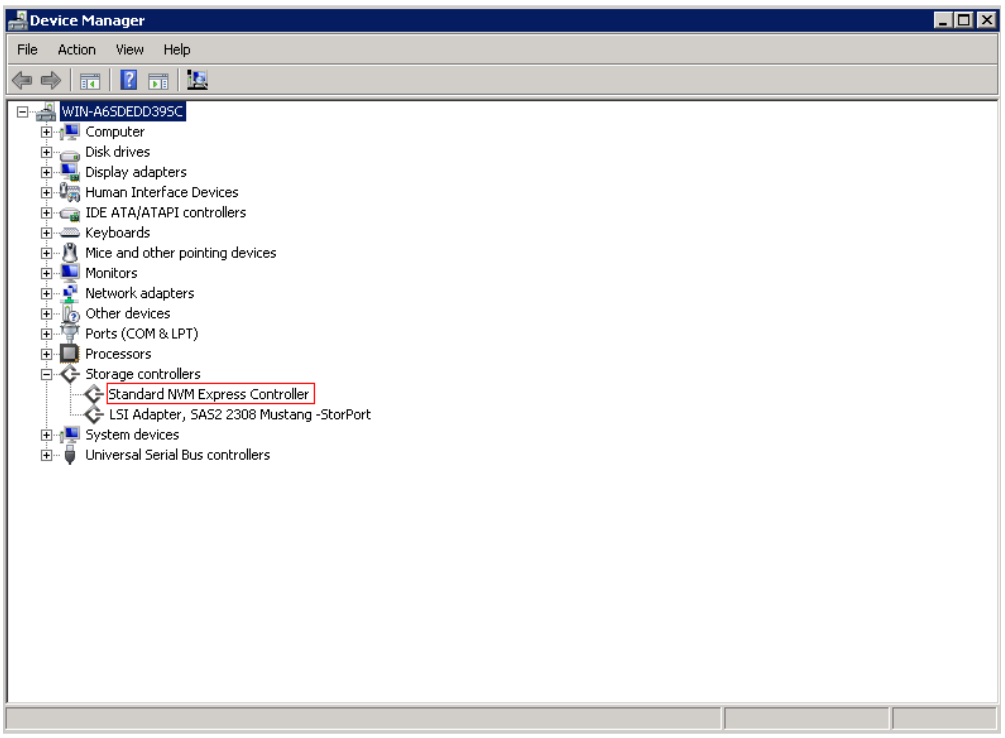

The following are three possible cases:

- If the device name is **Standard NVM Express Controller**, the device uses the NVMe driver provided by the OS.
- If the device name is **Huawei ES3000V5 storage controller**, the device uses Huawei's SSD driver.
- The device does not have a driver.

4. Determine the next step based on the driver type.
   - If the driver is provided by the OS, go to **5**.
   - If the driver is provided by Huawei, uninstall the driver by referring to **5.1.2 Uninstalling the Driver in Windows** and go to **1** in **Install the driver**.
   - If the device does not have a driver, go to **1** in **Install the driver**.
5. Choose **Start** > **Control Panel** > **Programs and Features**.
6. Click **View installed updates**. The **Installed Updates** window is displayed, as shown in **Figure 3-11**.

Figure 3-11 Installed Updates

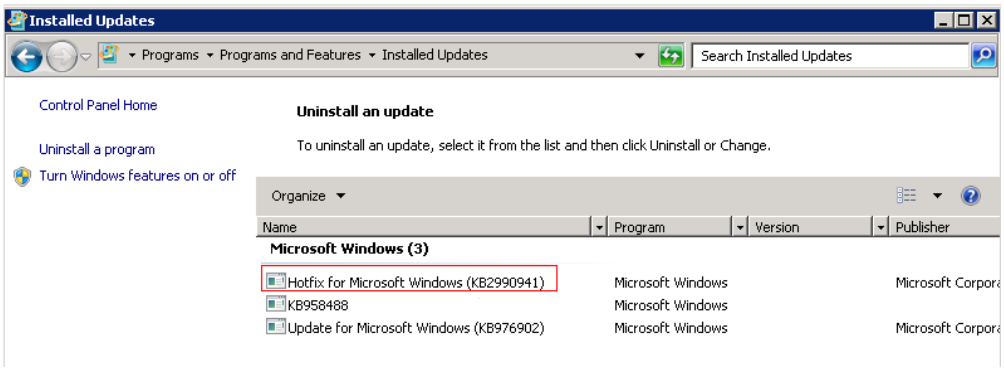

7. Right-click the **KB2990941** patch in the **Microsoft Windows** column and choose **Uninstall** from the shortcut menu.

Install the driver.

1. Double-click the driver installation file **hiodriver_V5.3.5.2_2008R2_x84_64.msi**.

   NOTE

   There are four driver installation file types for four Windows versions. ***.*.*.*** indicates the version number. This procedure uses hiodriver_V5.3.5.2_2008R2_x84_64.msi as an example.

   - hiodriver_v*.*.*.*_2008R2_x84_64.msi
   - hiodriver_v*.*.*.*_2012_x84_64.msi
   - hiodriver_v*.*.*.*_2012R2_x84_64.msi
   - hiodriver_v*.*.*.*_2016R2_x84_64.msi

   The installation interface is displayed, as shown in **Figure 3-12**.

Figure 3-12 Installation interface

2. Click **Next**.

   The **End-User License Agreement** window is displayed, as shown in **Figure 3-13**.

Figure 3-13 End-User License Agreement window

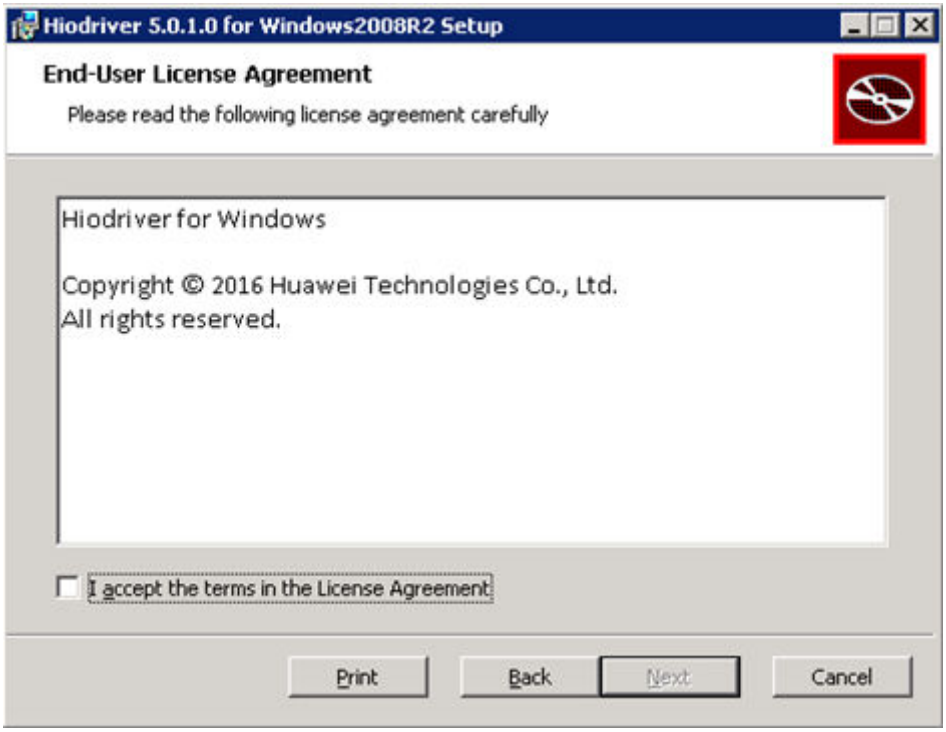

3. Select **I accept the terms in the License Agreement** and click **Next**.

The **Choose Setup Type** window is displayed, as shown in **Figure 3-14**.

Figure 3-14 Choose Setup Type window

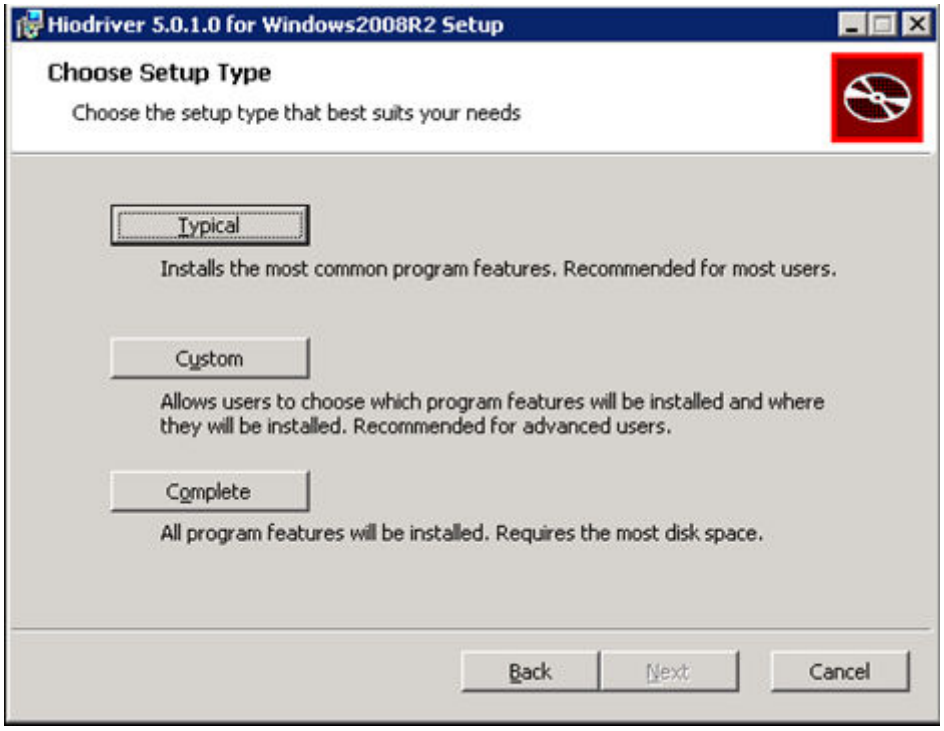

4. Click **Complete**.

   The **Ready to install** window is displayed, as shown in **Figure 3-15**.

Figure 3-15 Ready to install window

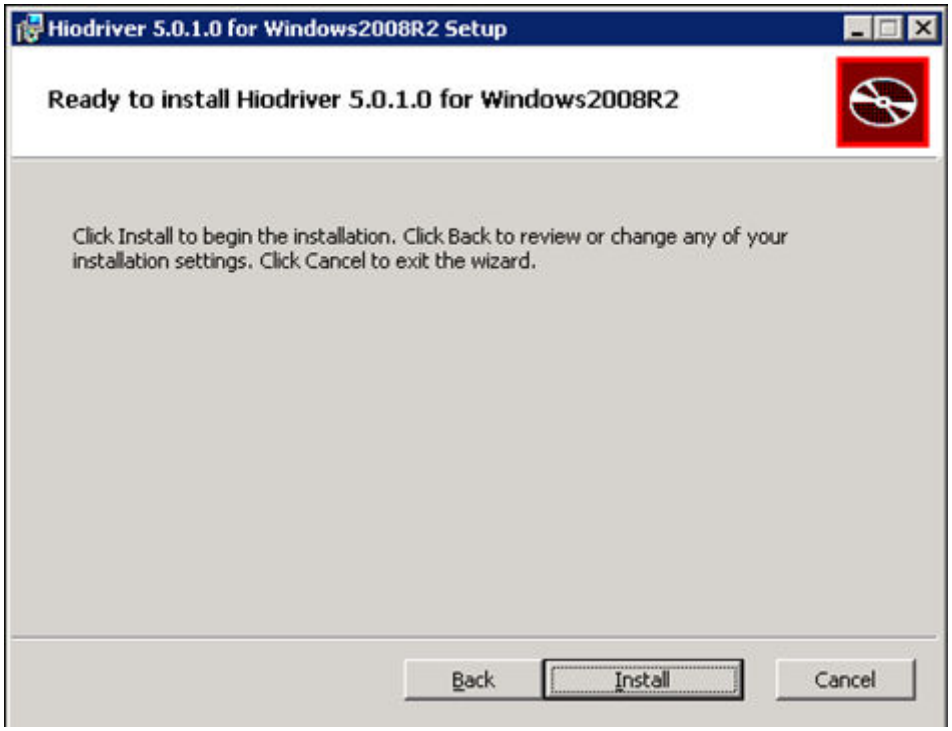

5. Click **Install**.

   When the installation is complete, the **Completing Setup Wizard** window is displayed, as shown in **Figure 3-16**.

Figure 3-16 Completing Setup Wizard window

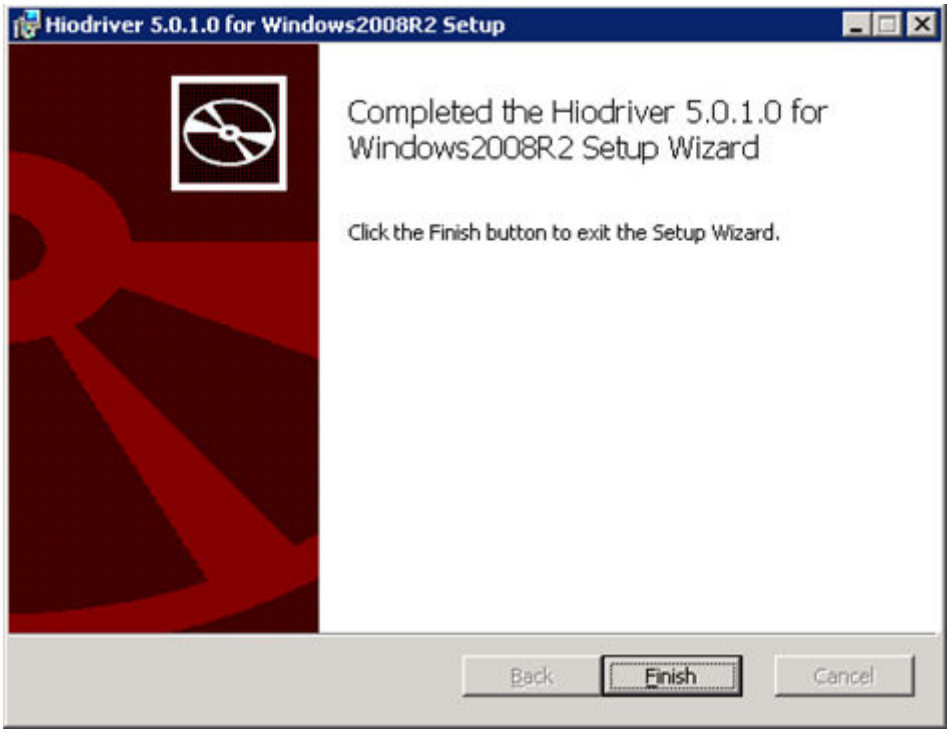

6. Click **Finish**.
7. Open Device Manager and expand **Storage controllers**.

   Check whether **Huawei ES3000V5 storage controller** exists under **Storage controllers**. See **Figure 3-17**.

   - If yes, go to **1** in **Install the driver**.
   - If no, contact Huawei technical support. For details, see **H Getting Help**.

Figure 3-17 Expanding Storage controllers

8. Expand **Disk drives**.

   Check whether **NVMe HWE32P43032M000N SCSI Disk Device** exists under **Disk drives**. See **Figure 3-18**.

   - If yes, the driver has been successfully installed.
   - If no, contact Huawei technical support for help. For details, see **H Getting Help**.

Figure 3-18 Expanding Disk drives

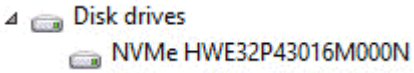

3.6.3 Installing the ES3000 V5 Driver in ESXi

Operation Scenario

You need to install the ES3000 V5 driver on the server before using the ES3000 V5 for the first time.

NOTE

- If you install multiple ES3000 V5s on a server, you need to install the driver only once.
- To install ESXi on the ES3000 V5, see G Installing an OS on the ES3000.

Preparations

Before installing the ES3000 V5 driver, check that:

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the ES3000 V5 driver package to the server, and verified the package integrity.
- You have uploaded the driver package to the server by using WinSCP or other software.

Procedure

Step 1 Log in to the OS as an administrator or a member of the administrator group.

Step 2 Install the ES3000 V5 driver.

The following procedure assumes that the driver package is in the **/tmp** directory, and the driver package name is **hiodriver-5.0.0.6-1OEM.550.0.0.2768847.x86_64.vib**.

Run the following command to install the ES3000 V5 driver package:

esxcli software vib install -v *driver package in a directory*

The following is an example command:

esxcli software vib install -v */tmp/hiodriver-5.0.0.6-1OEM.550.0.0.2768847.x86_64.vib*

NOTE

- In the driver file name, **5.0.0.6** is the driver version. The actual driver version may differ.
- Enter a complete directory after **-v**.
- If the "Could not find a trusted signer" error occurs indicating that the driver version is not certified by ESXi IOVP, add **--no-sig-check** at the end of the command to skip signature check.

Step 3 After the driver is installed, run the **reboot** command to restart the OS.

Step 4 After the restart, run the **esxcli software vib list | grep hiodriver** command. If the following command output is displayed, the driver is installed properly.

```
[root@esxi113:~] esxcli software vib list | grep hiodriverhiodriver 5.0.0.6-1OEM.
550.0.0.1391871 Huawei ESXiCertified 2016-01-04
```

Step 5 If the SSD device has been inserted into the server, run the **lspci | grep "Huawei ES3000"** command. If the device is displayed, it is detected by the OS.

```
[root@esxi113:~] lspci | grep "Huawei ES3000"
0000:06:00.0 Mass storage controller: Huawei Technologies Co Ltd. Huawei ES3000
storage Controller [vmhba3]
```


Step 6 Run the **vmkload_mod -l | grep hiodriver** command to check whether the driver is loaded.

If the command output is similar to the following, the device is found, which indicates that the driver is loaded.

```
[root@esxi113:~] vmkload_mod -l | grep hiodriver
hiodriver          2    112
```

- If yes, no further action is required.
- If no, go to **Step 7**.

Step 7 Contact Huawei technical support for help. For details, see **H Getting Help**.

----End

3.7 Installing the Tool Package

This section describes how to install the tool package in Linux, Windows, and ESXi operating systems (OSs).

3.7.1 Installing the Tool Package (Linux)

Operation Scenario

Install the tool package for managing the ES3000 V5.

NOTE

The tool package naming rules are as follows:

- For RHEL, CentOS, OEL, or SLES, the package name is hioadm-*version number-release number.hardware architecture*.rpm, for example, **hioadm-5.0.0.6-1.x86_64.rpm**.
- For Ubuntu or Debian, the package name is hioadm-*version number-release number.hardware architecture*.deb, for example, **hioadm-5.0.0.6-1.amd64.deb**.
- Citrix driver software package: **hioadm-**<*Version number*>**-**<*System version*>**.iso**, for example, **hioadm-citrix-6u2-5.0.0.8.iso**.

Preparations

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the tool package and verified its integrity.
- You have used WinSCP or other software to upload the tool package to a directory (for example, **/home**) on the server.

  NOTE

  If the system version is Citrix 6.2, use the 32-bit driver package, for example, **i386**. If the system version is Citrix 6.5 or 7.1, use a 64-bit driver package, for example, **x86_64**.

Procedure

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Go to the directory that stores the tool package, for example, **/home**.

Step 3 Run the installation command based on the OS type.

- For RHEL, CentOS, OEL, or SLES, run **rpm -ivh** *tool package name*.

Example: **rpm -ivh hioadm-5.0.0.6-1.x86_64.rpm**

- For Ubuntu or Debian, run **dpkg -i** *tool package name*.

  Example: **dpkg -i hioadm-5.0.0.6-1.amd64.deb**

- For Citrix:

  a. Check the version of Citrix

     - If Citrix is earlier than 7.1, perform **b**, **c**, and **d**.
     - If Citrix is later than 7.1, perform **e** and **f**.

  b. Run the **mount -o loop <*Tool package name*>** command to mount the ISO package to a directory, for example, **/tmp/citrix**.

     The following uses Citrix 6.2 as an example: **mount -o loop hioadm-5.0.0.6-6u2.iso /tmp/citrix.**

  c. Access the mount directory where the ISO package is located, for example, **/tmp/citrix**.

     cd /tmp/citrix

  d. Run **install.sh** in **/tmp/citrix** to install the tool.

  e. Run **xe update-upload file-name=hioadm-citrix7u1-5.0.0.8.iso** to upload the ISO package.

     NOTE

     The UUID of the pack is returned when the upload completes.

  f. Run **xe update-apply uuid=320232df-7adb-4cbe-a7a3-8515240879e1** to apply the package.

Step 4 Run the following commands to obtain help information about the tool or query tool information:

hioadm -h or **hioadm --help**

rpm -qa | grep hioadm

For Citrix 7.1, run **xe update-list** to query tool information

----End

3.7.2 Installing the Tool Package (Windows)

Operation Scenario

Install the tool package for managing the ES3000 V5. The ES3000 V5 tool is incompatible with the NVMe driver built in the Windows OS. Ensure that you have installed Huawei's NVMe driver before using the tool.

Preparations

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the tool package and verified its integrity.
- Upload the tool package through a shared directory, or log in to the remote desktop and directly upload the tool package to the server.

Procedure

Step 1 Log in to the operating system (OS) as an administrator or administrator group member.

Step 2 Double-click the tool package file **hioadm_V5.3.5.4_x86_64.msi**.

NOTE

V5.3.5.4 indicates a driver version number. This procedure uses **hioadm_V5.3.5.4_x86_64.msi** as an example.

The installation interface is displayed, as shown in **Figure 3-19**.

Figure 3-19 Installation interface

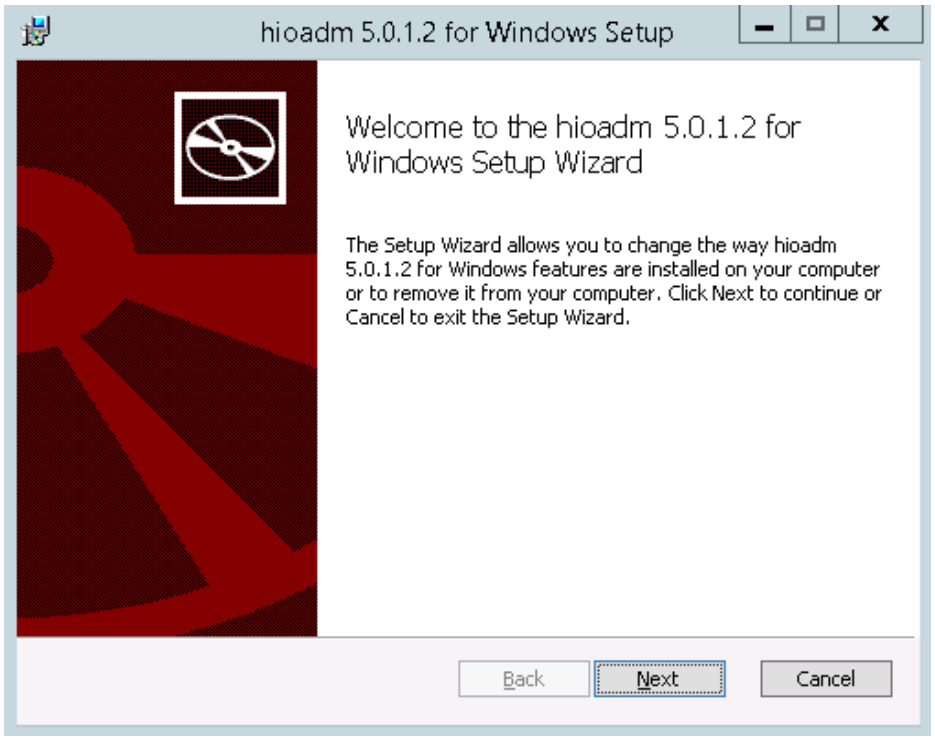

Step 3 Click **Next**.

The **End-User License Agreement** window is displayed, as shown in **Figure 3-20**.

Figure 3-20 End-User License Agreement window

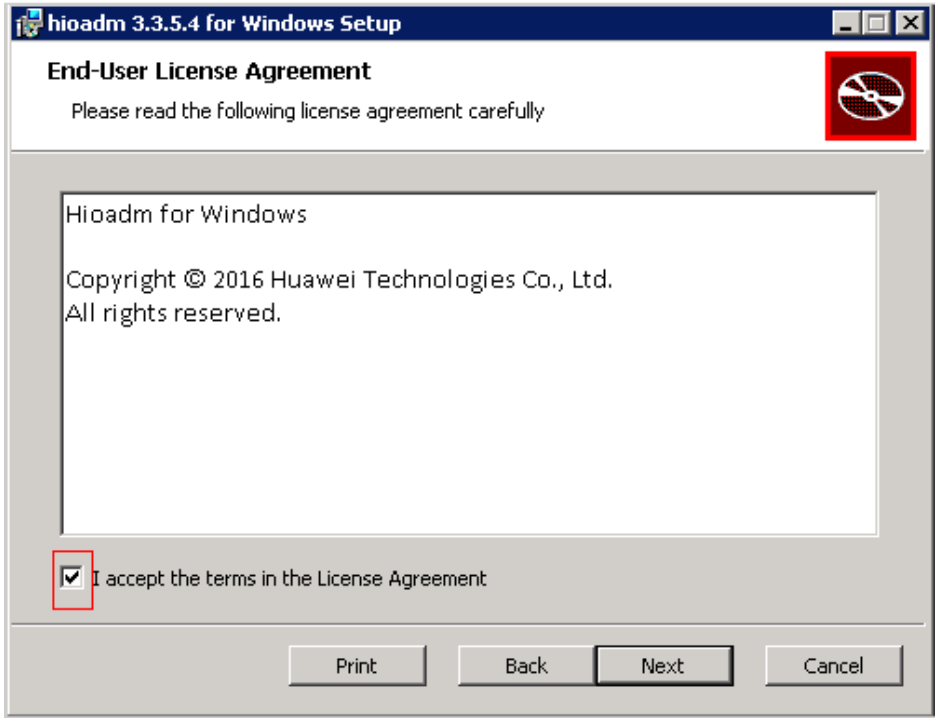

Step 4 Select **I accept the terms in the License Agreement** and click **Next**.

The **Destination Folder** window is displayed, as shown in **Figure 3-21**.

Figure 3-21 Destination Folder window

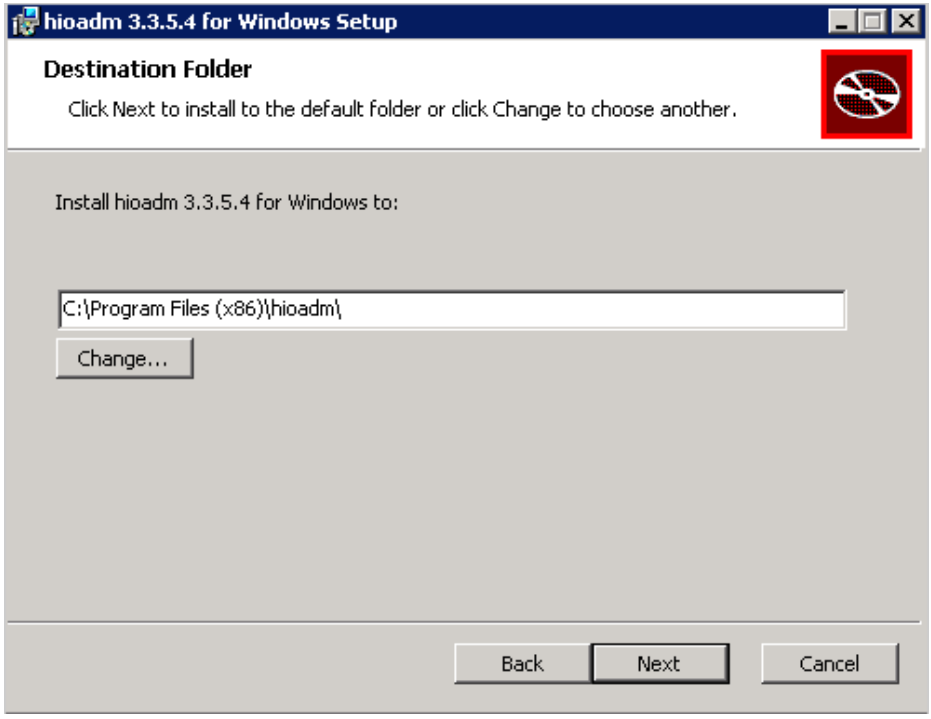

Step 5 Select a folder for installing the tool and click **Next**.

The **Ready to install** window is displayed, as shown in **Figure 3-22**.

Figure 3-22 Ready to install window

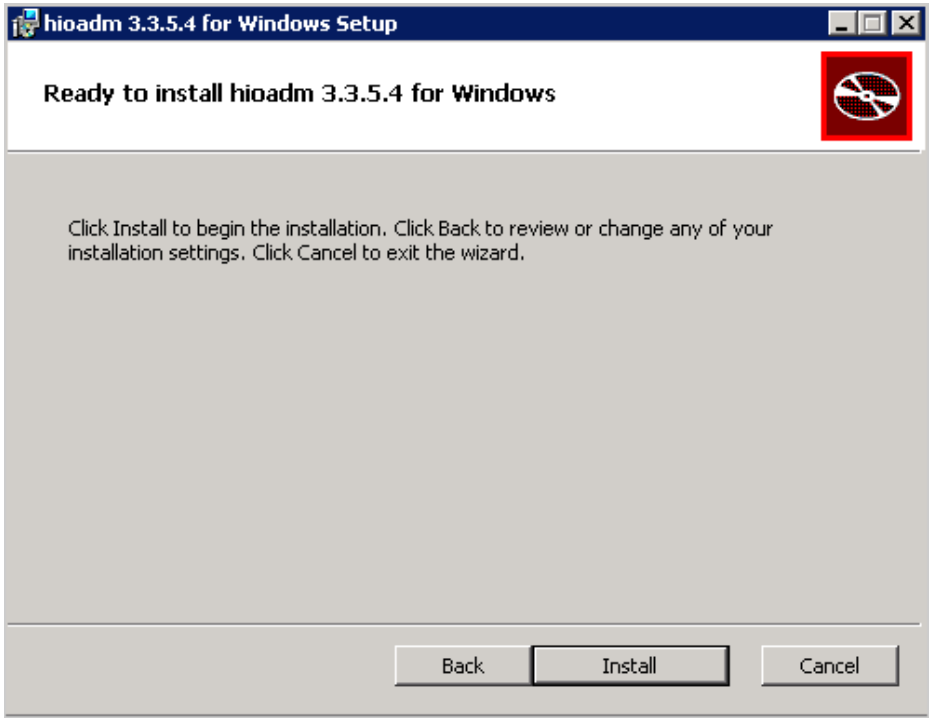

Step 6 Click **Install** to start installation.

When the installation is complete, the **Completing Setup Wizard** window is displayed, as shown in **Figure 3-23**.

Figure 3-23 Completing Setup Wizard window

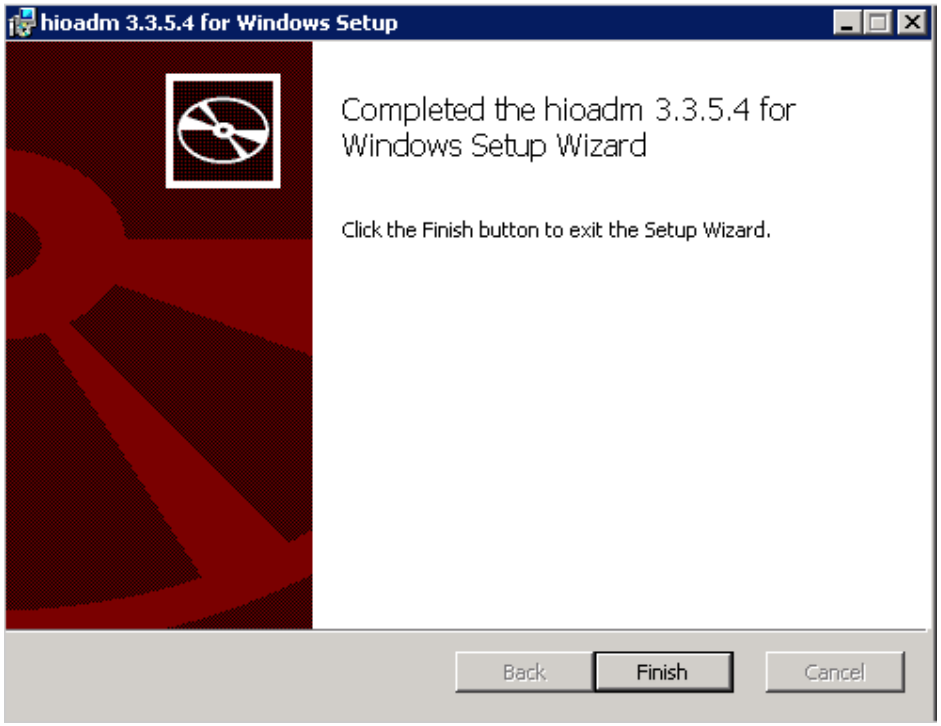

Step 7 Click **Finish**.

----End

3.7.3 Installing the Tool (ESXi)

Operation Scenario

Install the tool for managing the ES3000 V5.

Procedure

Step 1 Log in to the OS as an administrator or a member of the administrator group.

Step 2 Install the ES3000 V5 tool package.

The following procedure assumes that the tool package is in the **/tmp** directory, and the tool package name is **hioadm-5.0.1.0-1OEM.vib**.

Run the following command to install the ES3000 V5 tool package:

esxcli software vib install -v *tool package in a directory*

The following is an example command:

esxcli software vib install -v */tmp/hioadm-5.0.1.0-1OEM.vib*

NOTE

- In the tool package name, **5.0.1.0** is the software version. The actual tool version may differ.
- Enter a complete directory after **-v**.
- If the "Could not find a trusted signer" error occurs indicating that the tool version is not certified by ESXi IOVP, add **--no-sig-check** at the end of the command to skip signature check.

Step 3 After the tool is installed, run the following command to obtain help information about the tool.

hioadm -h

----End

4 Hot Swap

About This Chapter

This section describes how to hot-swap a PCIe SSD.

Hot swap includes the following two ways:

- Orderly hot swap: Users can directly insert a PCIe SSD when the OS is running, but needs to notify the OS before removing a PCIe SSD.
- Surprise hot swap: Users can directly insert or remove a PCIe SSD.

The ES3000 V5 NVMe PCIe SSD disk supports orderly hot swap and surprise hot swap on V5 servers, but does not support surprise hot swap on V3 servers. The ES3000 V5 NVMe PCIe SSD card does not support orderly or surprise hot swap.

CAUTION

- For details about operating systems that support orderly hot swap, see **4.1 OSs Supporting Hot Swap**.
- Swap only one ES3000 V5 SSD at a time. After one ES3000 V5 SSD is removed or inserted, wait for at least 3 seconds before swapping or inserting another SSD.
- SSD hot swap may affect services running on the SSD. Before performing a hot swap, notify system maintenance personnel to ensure service security.
- If you insert the ES3000 V5 SSD when it is not installed on a drive tray, it may fail to be detected.
- After the ES3000 V5 SSD is removed, wait for at least 3 seconds before inserting it again. If you insert the ES3000 V5 SSD again when it is halfway removed, it may fail to be detected.
- Do not pause when inserting the ES3000 V5 SSD, otherwise, it may fail to be detected and leads to system exceptions. If the SSD cannot be detected or system exceptions occur, power cycle the server after the SSD is properly installed.
- If the ES3000 V5 fails to be detected due to improper insertion, remove and insert it again.

4.1 OSs Supporting Hot Swap

NOTE

Ensure that the server and BIOS support hot swap.

The following OSs support orderly hot swap of the ES3000 V5:

- Mainstream Linux OSs:
  - RHEL 6.6, 6.7, 7.0, 7.1, 7.2, 7.3 and 7.4
  - CentOS 6.6, 6.7, 6.8, 7.0, 7.1, and 7.2
  - SLES 11SP4, 12, 12SP1
  - Oracle OEL 7.0, OEL 7.1, and OEL 7.2
  - Ubuntu 14.04 LTS, 15.04
- Mainstream Windows OSs:
  - Windows Server 2008 R2 64-bit
  - Windows Server 2012 64-Bit
  - Windows Server 2012 R2 64-bit
- Mainstream ESXi OSs:
  - ESXi 6.5 and later

The following Oss support surprise hot swap of the ES3000 V5:

- Mainstream Windows OSs:
  - Windows Server 2012 R2
  - Windows Server 2016
- Mainstream Linux OSs: RHEL 7.3 and 7.4

4.2 Orderly Hot Swap

You can insert the ES3000 V5 SSDs to running servers. For Linux and Windows, notify the OS before hot removing the ES3000 V5. For ESXi OSs, notify the OS before hot-swapping the ES3000 V5.

4.2.1 Orderly Hot Swap of a Drive in Linux

Operation Scenario

In Linux, perform the operations described in this section before performing an orderly hot swap of an ES3000 V5 SSD.

Preparations

You have correctly installed the ES3000 V5 SSD on the server.

You have installed the Linux tool package. For details, see **3.7.1 Installing the Tool Package (Linux)**.

Configuring Linux Kernel Parameters

In Linux, such as Red Hat Enterprise Linux (RHEL) 7.0, hot swap of the ES3000 V5 NVMe PCIe SSD disk will cause unmatched maximum payload size, I/O write errors, and other issues. This is because Linux does not fully support hot swap. To avoid this issue, make changes to the Linux kernel file. The changes vary with the Linux OS type.

The Linux kernel parameter configurations also vary depending on OS types:

- **Kernel Parameter Configuration on RHEL 7.3/7.4**
- **Kernel Parameter Configuration on RHEL 7.0/7.1/7.2 and CentOS 7.0/7.1/7.2**
- **Kernel Parameter Configuration on RHEL 6.6/6.7 and CentOS 6.6/6.7/6.8**
- **Kernel Parameter Configuration on Oracle Enterprise Linux 7.3/7.4**
- **Kernel Parameter Configuration on Oracle Enterprise Linux 7.0/7.1/7.2**
- **Kernel Parameter Configuration on SLES 12.2/12.3**
- **Kernel Parameter Configuration on SLES 12/12SP1**
- **Kernel Parameter Configuration on SLES 11 SP4**
- **Kernel Parameter Configuration on Ubuntu 16.04/16.04.2**
- **Kernel Parameter Configuration on Ubuntu 14.04/15.04**

Kernel Parameter Configuration on RHEL 7.3/7.4

The following procedure uses RHEL 7.3 as an example.

Step 1 Log in to the OS as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'Red Hat Enterprise Linux Server (3.10.0-514.el7.x86_64) 7.3 (Maipo)' --
class red --class gnu-linux --class gnu --class os --unrestricted
$menuentry_id_option 'gnulinux-3.10.0-514.el7.x86_64-
advanced-13434f8c-9fd6-4b12-81bc-91f85ad32a89' {
 load_video
 set gfxpayload=keep
 insmod gzio
 insmod part_msdos
 insmod xfs
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
8f04e328-185b-40d9-8f8a-19d3561dd769
```



```
 else
     search --no-floppy --fs-uuid --set=root 8f04e328-185b-40d9-8f8a-19d3561dd769
 fi
 linux16 /vmlinuz-3.10.0-514.el7.x86_64 root=/dev/mapper/rhel-root ro
crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet
LANG=en_US.UTF-8 pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd16 /initramfs-3.10.0-514.el7.x86_64.img
}
menuentry 'Red Hat Enterprise Linux Server (0-
rescue-144afff501d4492e8c752cca7ff44971) 7.3 (Maipo)' --class red --class gnu-
linux --class gnu --class os --unrestricted $menuentry_id_option 'gnulinux-0-
rescue-144afff501d4492e8c752cca7ff44971-
advanced-13434f8c-9fd6-4b12-81bc-91f85ad32a89' {
 load_video
 insmod gzio
 insmod part_msdos
 insmod xfs
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
8f04e328-185b-40d9-8f8a-19d3561dd769
 else
     search --no-floppy --fs-uuid --set=root 8f04e328-185b-40d9-8f8a-19d3561dd769
 fi
 linux16 /vmlinuz-0-rescue-144afff501d4492e8c752cca7ff44971 root=/dev/mapper/rhel-
root ro crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb quiet
pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd16 /initramfs-0-rescue-144afff501d4492e8c752cca7ff44971.img
}
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on RHEL 7.0/7.1/7.2 and CentOS 7.0/7.1/7.2

The following procedure uses CentOS 7.0 as an example.

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'CentOS Linux, with Linux 3.10.0-123.el7.x86_64' --class centos --class
```



```
gnu-linux --class gnu --class os --unrestricted $menuentry_id_option
'gnulinux-3.10.0-123.el7.x86_64-advanced-b054ccd0-c3d3-47eb-a449-f897f63f1b76' {
        load_video
        set gfxpayload=keep
        insmod gzio
        insmod part_msdos
        insmod xfs
        set root='hd0,msdos1'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'  b64fedb9-
b1e9-465d-994f-9cee287c1d57
        else
          search --no-floppy --fs-uuid --set=root b64fedb9-
b1e9-465d-994f-9cee287c1d57
        fi
        linux16 /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=b054ccd0-c3d3-47eb-a449-
f897f63f1b76 ro rd.lvm.lv=centos/swap vconsole.font=latarcyrheb-sun16
rd.lvm.lv=centos/root crashkernel=auto  vconsole.keymap=us rhgb quiet
LANG=en_US.UTF-8 pciehp.pciehp_force=1 pci=pcie_bus_perf
        initrd16 /initramfs-3.10.0-123.el7.x86_64.img
}
menuentry 'CentOS Linux, with Linux 0-rescue-17e12da386e9435aa7f2b6b718079678' --
class centos --class gnu-linux --class gnu --class os --unrestricted
$menuentry_id_option 'gnulinux-0-rescue-17e12da386e9435aa7f2b6b718079678-advanced-
b054ccd0-c3d3-47eb-a449-f897f63f1b76' {
        load_video
        insmod gzio
        insmod part_msdos
        insmod xfs
        set root='hd0,msdos1'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'  b64fedb9-
b1e9-465d-994f-9cee287c1d57
        else
          search --no-floppy --fs-uuid --set=root b64fedb9-
b1e9-465d-994f-9cee287c1d57
        fi
        linux16 /vmlinuz-0-rescue-17e12da386e9435aa7f2b6b718079678
root=UUID=b054ccd0-c3d3-47eb-a449-f897f63f1b76 ro rd.lvm.lv=centos/swap
vconsole.font=latarcyrheb-sun16 rd.lvm.lv=centos/root crashkernel=auto
vconsole.keymap=us rhgb quiet pciehp.pciehp_force=1 pci=pcie_bus_perf
        initrd16 /initramfs-0-rescue-17e12da386e9435aa7f2b6b718079678.img
}
if [ "x$default" = 'CentOS Linux, with Linux 3.10.0-123.el7.x86_64' ]; then
default='Advanced options for CentOS Linux>CentOS Linux, with Linux
3.10.0-123.el7.x86_64'; fi;
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on RHEL 6.6/6.7 and CentOS 6.6/6.7/6.8

The following procedure uses RHEL 6.7 as an example.

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /boot/grub/grub.conf** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
title Red Hat Enterprise Linux 6 (2.6.32-573.el6.x86_64)
        root (hd0,0)
        kernel /vmlinuz-2.6.32-573.el6.x86_64 ro root=/dev/mapper/VolGroup-
lv_root rd_NO_LUKS LANG=en_US.UTF-8 rd_NO_MD rd_LVM_LV=VolGroup/lv_swap
SYSFONT=latarcyrheb-sun16 crashkernel=128M rd_LVM_LV=VolGroup/lv_root
KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet pciehp.pciehp_force=1
pci=pcie_bus_perf
        initrd /initramfs-2.6.32-573.el6.x86_64.img
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on Oracle Enterprise Linux 7.3/7.4

The following procedure uses Oracle 7.3 as an example.

Step 1 Log in to the OS as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'Oracle Linux Server (4.1.12-61.1.18.el7uek.x86_64 with Unbreakable
Enterprise Kernel) 7.3' --class oracle --class gnu-linux --class gnu --class os --
unrestricted $menuentry_id_option 'gnulinux-4.1.12-61.1.18.el7uek.x86_64-
advanced-544c61b2-0ee7-4071-b081-084072d6dd85' {
 load_video
 set gfxpayload=keep
 insmod gzio
```



```
 insmod part_msdos
 insmod xfs
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 else
     search --no-floppy --fs-uuid --set=root 4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 fi
 linux16 /vmlinuz-4.1.12-61.1.18.el7uek.x86_64 root=/dev/mapper/ol-root ro
crashkernel=auto rd.lvm.lv=ol/root rd.lvm.lv=ol/swap rhgb quiet LANG=en_US.UTF-8
pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd16 /initramfs-4.1.12-61.1.18.el7uek.x86_64.img
}
menuentry 'Oracle Linux Server (3.10.0-514.el7.x86_64 with Linux) 7.3' --class
oracle --class gnu-linux --class gnu --class os --unrestricted
$menuentry_id_option 'gnulinux-3.10.0-514.el7.x86_64-advanced-544c61b2-0ee7-4071-
b081-084072d6dd85' {
 load_video
 set gfxpayload=keep
 insmod gzio
 insmod part_msdos
 insmod xfs
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 else
     search --no-floppy --fs-uuid --set=root 4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 fi
 linux16 /vmlinuz-3.10.0-514.el7.x86_64 root=/dev/mapper/ol-root ro
crashkernel=auto rd.lvm.lv=ol/root rd.lvm.lv=ol/swap rhgb quiet LANG=en_US.UTF-8
pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd16 /initramfs-3.10.0-514.el7.x86_64.img
}
menuentry 'Oracle Linux Server (0-rescue-ced47f5d0b3a4b6c875666af895ea929 with
Linux) 7.3' --class oracle --class gnu-linux --class gnu --class os --
unrestricted $menuentry_id_option 'gnulinux-0-rescue-
ced47f5d0b3a4b6c875666af895ea929-advanced-544c61b2-0ee7-4071-b081-084072d6dd85' {
 load_video
 insmod gzio
 insmod part_msdos
 insmod xfs
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 else
     search --no-floppy --fs-uuid --set=root 4645b72b-0415-4dc8-bc3a-b67b9f2ecd4b
 fi
 linux16 /vmlinuz-0-rescue-ced47f5d0b3a4b6c875666af895ea929 root=/dev/mapper/ol-
root ro crashkernel=auto rd.lvm.lv=ol/root rd.lvm.lv=ol/swap rhgb quiet
pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd16 /initramfs-0-rescue-ced47f5d0b3a4b6c875666af895ea929.img
}
if [ "x$default" = 'Oracle Linux Server (4.1.12-61.1.18.el7uek.x86_64 with
Unbreakable Enterprise Kernel) 7.3' ]; then default='Advanced options for Oracle
Linux Server>Oracle Linux Server (4.1.12-61.1.18.el7uek.x86_64 with Unbreakable
Enterprise Kernel) 7.3'; fi;
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on Oracle Enterprise Linux 7.0/7.1/7.2

The following procedure uses Oracle Enterprise Linux 7.0 as an example.

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

⚠ CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'Oracle Linux Server, with Linux 3.10.0-123.el7.x86_64' --class oracle
--class gnu-linux --class gnu --class os --unrestricted $menuentry_id_option
'gnulinux-3.10.0-123.el7.x86_64-advanced-eba3dd96-8be3-46ce-a8d9-21699e3fbedb' {
        load_video
        set gfxpayload=keep
        insmod gzio
        insmod part_msdos
        insmod xfs
        set root='hd0,msdos1'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
fcf9f115-8f51-496e-a7a8-905f03350b4e
        else
          search --no-floppy --fs-uuid --set=root fcf9f115-8f51-496e-
a7a8-905f03350b4e
        fi
        linux16 /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=eba3dd96-8be3-46ce-
a8d9-21699e3fbedb ro crashkernel=auto  vconsole.font=latarcyrheb-sun16
rd.lvm.lv=ol/swap rd.lvm.lv=ol/root vconsole.keymap=us rhgb quiet
LANG=en_US.UTF-8 pciehp.pciehp_force=1 pci=pcie_bus_perf
        initrd16 /initramfs-3.10.0-123.el7.x86_64.img
}
menuentry 'Oracle Linux Server, with Unbreakable Enterprise Kernel
3.8.13-35.3.1.el7uek.x86_64' --class oracle --class gnu-linux --class gnu --class
os --unrestricted $menuentry_id_option 'gnulinux-3.8.13-35.3.1.el7uek.x86_64-
advanced-eba3dd96-8be3-46ce-a8d9-21699e3fbedb' {
        load_video
        set gfxpayload=keep
        insmod gzio
        insmod part_msdos
        insmod xfs
        set root='hd0,msdos1'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
fcf9f115-8f51-496e-a7a8-905f03350b4e
        else
```



```
          search --no-floppy --fs-uuid --set=root fcf9f115-8f51-496e-
a7a8-905f03350b4e
        fi
        linux16 /vmlinuz-3.8.13-35.3.1.el7uek.x86_64 root=UUID=eba3dd96-8be3-46ce-
a8d9-21699e3fbedb ro crashkernel=auto  vconsole.font=latarcyrheb-sun16
rd.lvm.lv=ol/swap rd.lvm.lv=ol/root vconsole.keymap=us rhgb quiet
pciehp.pciehp_force=1 pci=pcie_bus_perf
        initrd16 /initramfs-3.8.13-35.3.1.el7uek.x86_64.img
}
menuentry 'Oracle Linux Server, with Linux 0-rescue-
a6e5b98906f643a09908dd24f3b775f7' --class oracle --class gnu-linux --class gnu --
class os --unrestricted $menuentry_id_option 'gnulinux-0-rescue-
a6e5b98906f643a09908dd24f3b775f7-advanced-eba3dd96-8be3-46ce-a8d9-21699e3fbedb' {
        load_video
        insmod gzio
        insmod part_msdos
        insmod xfs
      set root='hd0,msdos1'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1 --hint='hd0,msdos1'
fcf9f115-8f51-496e-a7a8-905f03350b4e
        else
          search --no-floppy --fs-uuid --set=root fcf9f115-8f51-496e-
a7a8-905f03350b4e
        fi
        linux16 /vmlinuz-0-rescue-a6e5b98906f643a09908dd24f3b775f7
root=UUID=eba3dd96-8be3-46ce-a8d9-21699e3fbedb ro crashkernel=auto
vconsole.font=latarcyrheb-sun16 rd.lvm.lv=ol/swap rd.lvm.lv=ol/root
vconsole.keymap=us rhgb quiet pciehp.pciehp_force=1 pci=pcie_bus_perf
        initrd16 /initramfs-0-rescue-a6e5b98906f643a09908dd24f3b775f7.img
}
if [ "x$default" = 'Oracle Linux Server, with Unbreakable Enterprise Kernel
3.8.13-35.3.1.el7uek.x86_64' ]; then default='Advanced options for Oracle Linux
Server>Oracle Linux Server, with Unbreakable Enterprise Kernel
3.8.13-35.3.1.el7uek.x86_64'; fi;
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on SLES 12.2/12.3

The following procedure uses SLES 12.2 as an example.

Step 1 Log in to the OS as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

⚠ CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'SLES 12-SP2'  --class sles --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-simple-9e0fffba-d347-4e34-9800-59271fd5ce73' {
 load_video
 set gfxpayload=keep
 insmod gzio
 insmod part_msdos
 insmod btrfs
 set root='hd0,msdos2'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2 --hint-
efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  9e0fffba-
d347-4e34-9800-59271fd5ce73
 else
     search --no-floppy --fs-uuid --set=root 9e0fffba-d347-4e34-9800-59271fd5ce73
 fi
 echo 'Loading Linux 4.4.21-69-default ...'
 linux  /boot/vmlinuz-4.4.21-69-default root=UUID=9e0fffba-
d347-4e34-9800-59271fd5ce73  ${extra_cmdline} resume=/dev/disk/by-uuid/
532f4ffb-61b6-48aa-b850-5052ff6cd5e2 splash=silent quiet showopts
pciehp.pciehp_force=1 pci=pcie_bus_perf
 echo 'Loading initial ramdisk ...'
 initrd /boot/initrd-4.4.21-69-default
}
submenu 'Advanced options for SLES 12-SP2' --hotkey=1 $menuentry_id_option
'gnulinux-advanced-9e0fffba-d347-4e34-9800-59271fd5ce73' {
 menuentry 'SLES 12-SP2, with Linux 4.4.21-69-default' --hotkey=2 --class sles --
class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-4.4.21-69-
default-advanced-9e0fffba-d347-4e34-9800-59271fd5ce73' {
          load_video
          set gfxpayload=keep
          insmod gzio
          insmod part_msdos
          insmod btrfs
          set root='hd0,msdos2'
          if [ x$feature_platform_search_hint = xy ]; then
            search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2 --hint-
efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  9e0fffba-
d347-4e34-9800-59271fd5ce73
          else
            search --no-floppy --fs-uuid --set=root 9e0fffba-
d347-4e34-9800-59271fd5ce73
          fi
          echo 'Loading Linux 4.4.21-69-default ...'
          linux  /boot/vmlinuz-4.4.21-69-default root=UUID=9e0fffba-
d347-4e34-9800-59271fd5ce73  ${extra_cmdline} resume=/dev/disk/by-uuid/
532f4ffb-61b6-48aa-b850-5052ff6cd5e2 splash=silent quiet showopts
pciehp.pciehp_force=1 pci=pcie_bus_perf
          echo 'Loading initial ramdisk ...'
          initrd /boot/initrd-4.4.21-69-default
 }
 menuentry 'SLES 12-SP2, with Linux 4.4.21-69-default (recovery mode)' --hotkey=3
--class sles --class gnu-linux --class gnu --class os $menuentry_id_option
'gnulinux-4.4.21-69-default-recovery-9e0fffba-d347-4e34-9800-59271fd5ce73' {
          load_video
          set gfxpayload=keep
          insmod gzio
          insmod part_msdos
          insmod btrfs
          set root='hd0,msdos2'
          if [ x$feature_platform_search_hint = xy ]; then
```



```
            search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2 --hint-
efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  9e0fffba-
d347-4e34-9800-59271fd5ce73
          else
            search --no-floppy --fs-uuid --set=root 9e0fffba-
d347-4e34-9800-59271fd5ce73
          fi
          echo 'Loading Linux 4.4.21-69-default ...'
          linux  /boot/vmlinuz-4.4.21-69-default root=UUID=9e0fffba-
d347-4e34-9800-59271fd5ce73  ${extra_cmdline} pciehp.pciehp_force=1
pci=pcie_bus_perf
          echo 'Loading initial ramdisk ...'
          initrd /boot/initrd-4.4.21-69-default
 }
}
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

NOTE

Whether the file content includes "### BEGIN /etc/grub.d/20_linux_xen ###" depends on whether Xen Hypervisor is selected during OS installation.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on SLES 12/12SP1

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /boot/grub2/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
menuentry 'SLES12'  --class sles12 --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-simple-dcbfdecb-bcd9-4003-ba51-a95c7e31c56e' {
        load_video
        set gfxpayload=keep
        insmod gzio
        insmod part_msdos
        insmod btrfs
        set root='hd0,msdos2'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2 --hint-
efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
        else
          search --no-floppy --fs-uuid --set=root dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e
        fi
        echo    'Loading Linux 3.12.28-4-default ...'
        linux /boot/vmlinuz-3.12.28-4-default root=UUID=dcbfdecb-bcd9-4003-ba51-
```



```
a95c7e31c56e  ${extra_cmdline}  resume=/dev/disk/by-uuid/9a586421-38ef-43bb-9209-
afe5eede22e3 splash=silent quiet crashkernel=216M-:108M  showopts
pciehp.pciehp_force=1 pci=pcie_bus_perf
        echo    'Loading initial ramdisk ...'
        initrd  /boot/initrd-3.12.28-4-default
}
submenu 'Advanced options for SLES12' --hotkey=1 $menuentry_id_option 'gnulinux-
advanced-dcbfdecb-bcd9-4003-ba51-a95c7e31c56e' {
        menuentry 'SLES12, with Linux 3.12.28-4-default' --hotkey=2 --class
sles12 --class gnu-linux --class gnu --class os $menuentry_id_option
'gnulinux-3.12.28-4-default-advanced-dcbfdecb-bcd9-4003-ba51-a95c7e31c56e' {
                load_video
                set gfxpayload=keep
                insmod gzio
                insmod part_msdos
                insmod btrfs
                set root='hd0,msdos2'
                if [ x$feature_platform_search_hint = xy ]; then
                  search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2
--hint-efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
                else
                  search --no-floppy --fs-uuid --set=root dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e
                fi
                echo    'Loading Linux 3.12.28-4-default ...'
                linux /boot/vmlinuz-3.12.28-4-default root=UUID=dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e  ${extra_cmdline}  resume=/dev/disk/by-uuid/
9a586421-38ef-43bb-9209-afe5eede22e3 splash=silent quiet crashkernel=216M-:108M
showopts pciehp.pciehp_force=1 pci=pcie_bus_perf
                echo    'Loading initial ramdisk ...'
                initrd  /boot/initrd-3.12.28-4-default
        }
        menuentry 'SLES12, with Linux 3.12.28-4-default (recovery mode)' --
hotkey=3 --class sles12 --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-3.12.28-4-default-recovery-dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e' {
                load_video
                set gfxpayload=keep
                insmod gzio
                insmod part_msdos
                insmod btrfs
                set root='hd0,msdos2'
                if [ x$feature_platform_search_hint = xy ]; then
                  search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2
--hint-efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
                else
                  search --no-floppy --fs-uuid --set=root dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e
                fi
                echo    'Loading Linux 3.12.28-4-default ...'
                linux /boot/vmlinuz-3.12.28-4-default root=UUID=dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e  ${extra_cmdline} showopts apm=off noresume edd=off
powersaved=off nohz=off highres=off processor.max_cstate=1 nomodeset x11failsafe
crashkernel=216M-:108M pciehp.pciehp_force=1 pci=pcie_bus_perf
                echo    'Loading initial ramdisk ...'
                initrd  /boot/initrd-3.12.28-4-default
        }
}

### END /etc/grub.d/10_linux ###

### BEGIN /etc/grub.d/20_linux_xen ###
menuentry 'SLES12, with Xen hypervisor' --class sles12 --class gnu-linux --class
gnu --class os --class xen $menuentry_id_option 'xen-gnulinux-simple-dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e' {
        insmod part_msdos
        insmod btrfs
```



```
        set root='hd0,msdos2'
        if [ x$feature_platform_search_hint = xy ]; then
          search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos2 --hint-
efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --hint='hd0,msdos2'  dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
        else
          search --no-floppy --fs-uuid --set=root dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e
        fi
        echo    'Loading Xen 4.4.1_06-2.2 ...'
        if [ "$grub_platform" = "pc" -o "$grub_platform" = "" ]; then
            xen_rm_opts=
        else
            xen_rm_opts="no-real-mode edd=off"
        fi
        multiboot       /boot/xen-4.4.1_06-2.2.gz placeholder  crashkernel=216M-:
108M ${xen_rm_opts}
        echo    'Loading Linux 3.12.28-4-xen ...'
        module  /boot/vmlinuz-3.12.28-4-xen placeholder root=UUID=dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e  ${extra_cmdline}  resume=/dev/disk/by-uuid/
9a586421-38ef-43bb-9209-afe5eede22e3 splash=silent quiet crashkernel=216M-:108M
showopts pciehp.pciehp_force=1 pci=pcie_bus_perf
        echo    'Loading initial ramdisk ...'
        module  --nounzip /boot/initrd-3.12.28-4-xen
}
submenu 'Advanced options for SLES12 (with Xen hypervisor)' $menuentry_id_option
'gnulinux-advanced-dcbfdecb-bcd9-4003-ba51-a95c7e31c56e' {
        submenu 'Xen hypervisor, version 4.4.1_06-2.2' $menuentry_id_option 'xen-
hypervisor-4.4.1_06-2.2-dcbfdecb-bcd9-4003-ba51-a95c7e31c56e' {
                menuentry 'SLES12, with Xen 4.4.1_06-2.2 and Linux 3.12.28-4-xen'
--class sles12 --class gnu-linux --class gnu --class os --class xen
$menuentry_id_option 'xen-gnulinux-3.12.28-4-xen-advanced-dcbfdecb-bcd9-4003-ba51-
a95c7e31c56e' {
                        insmod part_msdos
                        insmod btrfs
                        set root='hd0,msdos2'
                        if [ x$feature_platform_search_hint = xy ]; then
                          search --no-floppy --fs-uuid --set=root --hint-
bios=hd0,msdos2 --hint-efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --
hint='hd0,msdos2'  dcbfdecb-bcd9-4003-ba51-a95c7e31c56e
                        else
                          search --no-floppy --fs-uuid --set=root dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
                        fi
                        echo    'Loading Xen 4.4.1_06-2.2 ...'
                        if [ "$grub_platform" = "pc" -o "$grub_platform" = "" ];
then
                            xen_rm_opts=
                        else
                            xen_rm_opts="no-real-mode edd=off"
                        fi
                        multiboot       /boot/xen-4.4.1_06-2.2.gz placeholder
crashkernel=216M-:108M ${xen_rm_opts}
                        echo    'Loading Linux 3.12.28-4-xen ...'
                        module  /boot/vmlinuz-3.12.28-4-xen placeholder
root=UUID=dcbfdecb-bcd9-4003-ba51-a95c7e31c56e  ${extra_cmdline}  resume=/dev/
disk/by-uuid/9a586421-38ef-43bb-9209-afe5eede22e3 splash=silent quiet
crashkernel=216M-:108M  showopts pciehp.pciehp_force=1 pci=pcie_bus_perf
                        echo    'Loading initial ramdisk ...'
                        module  --nounzip /boot/initrd-3.12.28-4-xen
                }
                menuentry 'SLES12, with Xen 4.4.1_06-2.2 and Linux 3.12.28-4-xen
(recovery mode)' --class sles12 --class gnu-linux --class gnu --class os --class
xen $menuentry_id_option 'xen-gnulinux-3.12.28-4-xen-recovery-dcbfdecb-bcd9-4003-
ba51-a95c7e31c56e' {
                        insmod part_msdos
                        insmod btrfs
                        set root='hd0,msdos2'
                        if [ x$feature_platform_search_hint = xy ]; then
```



```
                          search --no-floppy --fs-uuid --set=root --hint-
bios=hd0,msdos2 --hint-efi=hd0,msdos2 --hint-baremetal=ahci0,msdos2 --
hint='hd0,msdos2'  dcbfdecb-bcd9-4003-ba51-a95c7e31c56e
                        else
                          search --no-floppy --fs-uuid --set=root dcbfdecb-
bcd9-4003-ba51-a95c7e31c56e
                        fi
                        echo    'Loading Xen 4.4.1_06-2.2 ...'
                        if [ "$grub_platform" = "pc" -o "$grub_platform" = "" ];
then
                            xen_rm_opts=
                        else
                            xen_rm_opts="no-real-mode edd=off"
                        fi
                        multiboot       /boot/xen-4.4.1_06-2.2.gz placeholder  $
{xen_rm_opts}
                        echo    'Loading Linux 3.12.28-4-xen ...'
                        module  /boot/vmlinuz-3.12.28-4-xen placeholder
root=UUID=dcbfdecb-bcd9-4003-ba51-a95c7e31c56e single  ${extra_cmdline}
pciehp.pciehp_force=1 pci=pcie_bus_perf
                        echo    'Loading initial ramdisk ...'
                        module  --nounzip /boot/initrd-3.12.28-4-xen
                }
        }
}
### END /etc/grub.d/20_linux_xen ###
…(Some code omitted.)
```

NOTE

Whether the file content includes "### BEGIN /etc/grub.d/20_linux_xen ###" depends on whether Xen Hypervisor is selected during OS installation.

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on SLES 11 SP4

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /etc/init.d/boot.local** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 Add the following commands to the end of the boot configuration file content.

modprobe -r pciehp > /dev/null 2>&1

modprobe pciehp pciehp_force=1

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Run the **vim /boot/grub/menu.lst** command to open the boot configuration file.

Step 7 Enter **i**.

Step 8 In the boot configuration file, add **pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
###Don't change this comment - YaST2 identifier: Original name: xen###
title Xen -- SUSE Linux Enterprise Server 11 SP4 - 3.0.101-63
    root (hd0,1)
    kernel /boot/xen.gz vga=mode-0x314
    module /boot/vmlinuz-3.0.101-63-xen root=/dev/sda2 resume=/dev/sda1
splash=silent showopts vga=0x314 pci=pcie_bus_perf
    module /boot/initrd-3.0.101-63-xen

###Don't change this comment - YaST2 identifier: Original name: linux###
title SUSE Linux Enterprise Server 11 SP4 - 3.0.101-63
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.101-63-default root=/dev/sda2 resume=/dev/sda1
splash=silent crashkernel=256M-:128M showopts vga=0x314 pci=pcie_bus_perf
    initrd /boot/initrd-3.0.101-63-default

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title Failsafe -- SUSE Linux Enterprise Server 11 SP4 - 3.0.101-63
    root (hd0,1)
    kernel /boot/vmlinuz-3.0.101-63-default root=/dev/sda2 showopts ide=nodma
apm=off noresume edd=off powersaved=off nohz=off highres=off
processor.max_cstate=1 nomodeset x11failsafe vga=0x314 pci=pcie_bus_perf
    initrd /boot/initrd-3.0.101-63-default
…(Some code omitted.)
```

NOTE

Whether the file content includes "title Xen" depends on whether Xen Hypervisor is selected during OS installation.

Step 9 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 10 Restart the OS for the settings to take effect.

Step 11 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on Ubuntu 16.04/16.04.2

The following procedure uses Ubuntu16.04 as an example.

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the **vim /boot/grub/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.


```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
function gfxmode {
 set gfxpayload="${1}"
 if [ "${1}" = "keep" ]; then
     set vt_handoff=vt.handoff=7
 else
     set vt_handoff=
 fi
}
if [ "${recordfail}" != 1 ]; then
  if [ -e ${prefix}/gfxblacklist.txt ]; then
    if hwmatch ${prefix}/gfxblacklist.txt 3; then
      if [ ${match} = 0 ]; then
        set linux_gfx_mode=keep
      else
        set linux_gfx_mode=text
      fi
    else
      set linux_gfx_mode=text
    fi
  else
    set linux_gfx_mode=keep
  fi
else
  set linux_gfx_mode=text
fi
export linux_gfx_mode
menuentry 'Ubuntu' --class ubuntu --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-simple-e2afe344-3108-497b-ae46-014a53fe3cb5' {
 recordfail
 load_video
 gfxmode $linux_gfx_mode
 insmod gzio
 if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
 insmod part_msdos
 insmod ext2
 set root='hd0,msdos1'
 if [ x$feature_platform_search_hint = xy ]; then
     search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  e2afe344-3108-497b-
ae46-014a53fe3cb5
 else
     search --no-floppy --fs-uuid --set=root e2afe344-3108-497b-ae46-014a53fe3cb5
 fi
 linux  /boot/vmlinuz-4.4.0-21-generic root=UUID=e2afe344-3108-497b-
ae46-014a53fe3cb5 ro pciehp.pciehp_force=1 pci=pcie_bus_perf
 initrd /boot/initrd.img-4.4.0-21-generic
}
submenu 'Advanced options for Ubuntu' $menuentry_id_option 'gnulinux-advanced-
e2afe344-3108-497b-ae46-014a53fe3cb5' {
 menuentry 'Ubuntu, with Linux 4.4.0-21-generic' --class ubuntu --class gnu-linux
--class gnu --class os $menuentry_id_option 'gnulinux-4.4.0-21-generic-advanced-
e2afe344-3108-497b-ae46-014a53fe3cb5' {
          recordfail
          load_video
          gfxmode $linux_gfx_mode
          insmod gzio
          if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
          insmod part_msdos
          insmod ext2
          set root='hd0,msdos1'
          if [ x$feature_platform_search_hint = xy ]; then
            search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  e2afe344-3108-497b-
ae46-014a53fe3cb5
          else
            search --no-floppy --fs-uuid --set=root e2afe344-3108-497b-
ae46-014a53fe3cb5
```



```
          fi
          echo 'Loading Linux 4.4.0-21-generic ...'
          linux  /boot/vmlinuz-4.4.0-21-generic root=UUID=e2afe344-3108-497b-
ae46-014a53fe3cb5 ro pciehp.pciehp_force=1 pci=pcie_bus_perf
          echo 'Loading initial ramdisk ...'
          initrd /boot/initrd.img-4.4.0-21-generic
 }
 menuentry 'Ubuntu, with Linux 4.4.0-21-generic (recovery mode)' --class ubuntu --
class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-4.4.0-21-
generic-recovery-e2afe344-3108-497b-ae46-014a53fe3cb5' {
          recordfail
          load_video
          insmod gzio
          if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
          insmod part_msdos
          insmod ext2
          set root='hd0,msdos1'
          if [ x$feature_platform_search_hint = xy ]; then
            search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  e2afe344-3108-497b-
ae46-014a53fe3cb5
          else
            search --no-floppy --fs-uuid --set=root e2afe344-3108-497b-
ae46-014a53fe3cb5
          fi
          echo 'Loading Linux 4.4.0-21-generic ...'
          linux  /boot/vmlinuz-4.4.0-21-generic root=UUID=e2afe344-3108-497b-
ae46-014a53fe3cb5 ro recovery nomodeset pciehp.pciehp_force=1 pci=pcie_bus_perf
          echo 'Loading initial ramdisk ...'
          initrd /boot/initrd.img-4.4.0-21-generic
 }
}
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Kernel Parameter Configuration on Ubuntu 14.04/15.04

The following procedure uses Ubuntu14.04 as an example.

Step 1 Log in to the OS as an administrator.

Step 2 Run the **vim /boot/grub/grub.cfg** command to open the boot configuration file.

Step 3 Enter **i**.

Step 4 In the boot configuration file, add **pciehp.pciehp_force=1 pci=pcie_bus_perf** to the boldface part shown in the following:

⚠ CAUTION

Between the added content and its preceding content, a space is required and no line feed is allowed.

```
…(Some code omitted.)
### BEGIN /etc/grub.d/10_linux ###
```



```
function gfxmode {
set gfxpayload="${1}"
if [ "${1}" = "keep" ]; then
set vt_handoff=vt.handoff=7
else
set vt_handoff=
fi
}
if [ "${recordfail}" != 1 ]; then
if [ -e ${prefix}/gfxblacklist.txt ]; then
if hwmatch ${prefix}/gfxblacklist.txt 3; then
if [ ${match} = 0 ]; then
set linux_gfx_mode=keep
else
set linux_gfx_mode=text
fi
else
set linux_gfx_mode=text
fi
else
set linux_gfx_mode=keep
fi
else
set linux_gfx_mode=text
fi
export linux_gfx_mode
menuentry 'Ubuntu' --class ubuntu --class gnu-linux --class gnu --class os
$menuentry_id_option 'gnulinux-simple-b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8' {
recordfail
load_video
gfxmode $linux_gfx_mode
insmod gzio
insmod part_msdos
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
else
search --no-floppy --fs-uuid --set=root b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
fi
linux     /boot/vmlinuz-3.13.0-24-generic root=UUID=b7889e8e-dbc4-4fe9-bf5c-
dab9c20a93f8 ro  pciehp.pciehp_force=1 pci=pcie_bus_perf
initrd    /boot/initrd.img-3.13.0-24-generic
}
submenu 'Advanced options for Ubuntu' $menuentry_id_option 'gnulinux-advanced-
b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8' {
menuentry 'Ubuntu, with Linux 3.13.0-24-generic' --class ubuntu --class gnu-linux
--class gnu --class os $menuentry_id_option 'gnulinux-3.13.0-24-generic-advanced-
b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8' {
recordfail
load_video
gfxmode $linux_gfx_mode
insmod gzio
insmod part_msdos
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
else
search --no-floppy --fs-uuid --set=root b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
fi
echo      'Loading Linux 3.13.0-24-generic ...'
linux     /boot/vmlinuz-3.13.0-24-generic root=UUID=b7889e8e-dbc4-4fe9-bf5c-
dab9c20a93f8 ro  pciehp.pciehp_force=1 pci=pcie_bus_perf
echo      'Loading initial ramdisk ...'
initrd    /boot/initrd.img-3.13.0-24-generic
}
```

```
menuentry 'Ubuntu, with Linux 3.13.0-24-generic (recovery mode)' --class ubuntu --
class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-3.13.0-24-
generic-recovery-b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8' {
recordfail
load_video
insmod gzio
insmod part_msdos
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
search --no-floppy --fs-uuid --set=root --hint-bios=hd0,msdos1 --hint-
efi=hd0,msdos1 --hint-baremetal=ahci0,msdos1  b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
else
search --no-floppy --fs-uuid --set=root b7889e8e-dbc4-4fe9-bf5c-dab9c20a93f8
fi
echo      'Loading Linux 3.13.0-24-generic ...'
linux   /boot/vmlinuz-3.13.0-24-generic root=UUID=b7889e8e-dbc4-4fe9-bf5c-
dab9c20a93f8 ro recovery nomodeset  pciehp.pciehp_force=1 pci=pcie_bus_perf
echo      'Loading initial ramdisk ...'
initrd    /boot/initrd.img-3.13.0-24-generic
}
}
### END /etc/grub.d/10_linux ###
…(Some code omitted.)
```

Step 5 Press **Esc**, enter **:wq**, and press **Enter** to save the edits.

Step 6 Restart the OS for the settings to take effect.

Step 7 Go to **Notifying the OS**.

----End

Notifying the OS

Step 1 Log in to the OS as the **root** user.

Step 2 Stop all services accessing the ES3000 V5 SSD.

Step 3 If you have used the **mount** command to attach the SSD partitions to certain paths, use the **umount** command to detach the partitions.

Step 4 Determine the server where the SSD is installed.

- If the SSD is installed on a V3 server, go to **Step 5**.
- If the SSD is installed on a V5 server, see the orderly hot swap operation guide of each vendor.

Step 5 Notify the OS to take power off the device.

Method 1: Run the **echo 0 > /sys/bus/pci/slots/$slot/power** command.

$slot in the commands indicates the actual number of the slot where the device is attached.

Method 2: Run the **hioadm hotplug -d** *devicename* **-r** command.

Use this command if you have installed hioadm. In this command, *devicename* indicates the name of the device to be removed, such as **nvme0** and **nvme1**.

NOTE

After the previous command is executed, you can remove the SSD if the green indicator is off and the yellow indicator blinks at the frequency of 0.5 Hz.

Step 6 If the OS uses the irqbalance service to balance CPU interrupts, restart the service by using one of the following commands after performing a hot swap:

- For RHEL 6 or CentOs 6, run **service irqbalance restart**.
- For SLES 11, run **service irq_balancer restart**.
- For SLES 12, RHEL 7 or Oracle 7, run **systemctl restart irqbalance.service**.

NOTE

- For RHEL 6, CentOS 6, and SLES 12, you can add IRQBALANCE_ARGS="--hintpolicy=exact" to the end of the **/etc/sysconfig/irqbalance** file content to better balance interrupts.
- To configure the irqbalance service in other OSs, refer to documents related to the OSs.

----End

4.2.2 Orderly Hot Swap of a Drive in Windows

Operation Scenario

Notify the Windows OS before hot removing the ES3000 V5 SSD.

Preparations

- You have correctly installed the ES3000 V5 SSD on the server.
- You have installed the Windows tool package. For details, see **3.7.2 Installing the Tool Package (Windows)**.

Procedure

Windows GUI

Step 1 Log in to the OS as an administrator or administrator group member.

Step 2 Stop all services accessing the ES3000 V5 SSD.

Step 3 Double-click the **HioadmGUI_DotNet_x.0.exe** file in the tool installation directory to display the Windows GUI window, as shown in **Figure 4-1**.

- If the .NET Framework version is 2.0 or 3.5, double-click the **HioadmGUI_DotNet_2.0.exe** file in the tool installation directory.
- If the .NET Framework version is 4.0 or 4.5, double-click the **HioadmGUI_DotNet_4.0.exe** file in the tool installation directory.

Figure 4-1 Basic information

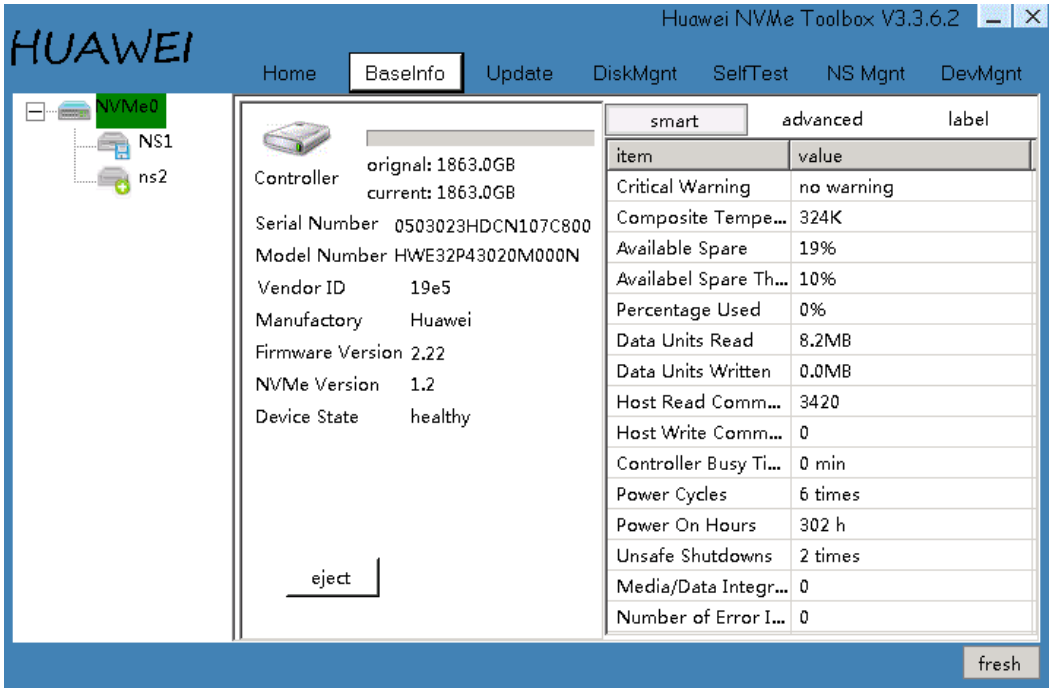

Step 4 Click **eject** in **Figure 4-1** and then click **OK** in the displayed dialog box.

If a success message is displayed, locate the ES3000 V5 SSD whose yellow indicator blinks at the frequency of 0.5 Hz and hot swap it from the server. If a failure message indicating the OS is accessing the ES3000 V5 SSD, stop related services before clicking **eject** again. You can hot remove the ES3000 V5 SSD only when a success message is displayed.

----End

Windows CLI

Step 1 Log in to the OS as an administrator or administrator group member.

Step 2 Stop all services accessing the ES3000 V5 SSD, and make all logical drives offline.

Figure 4-2 Making a logical drive offline

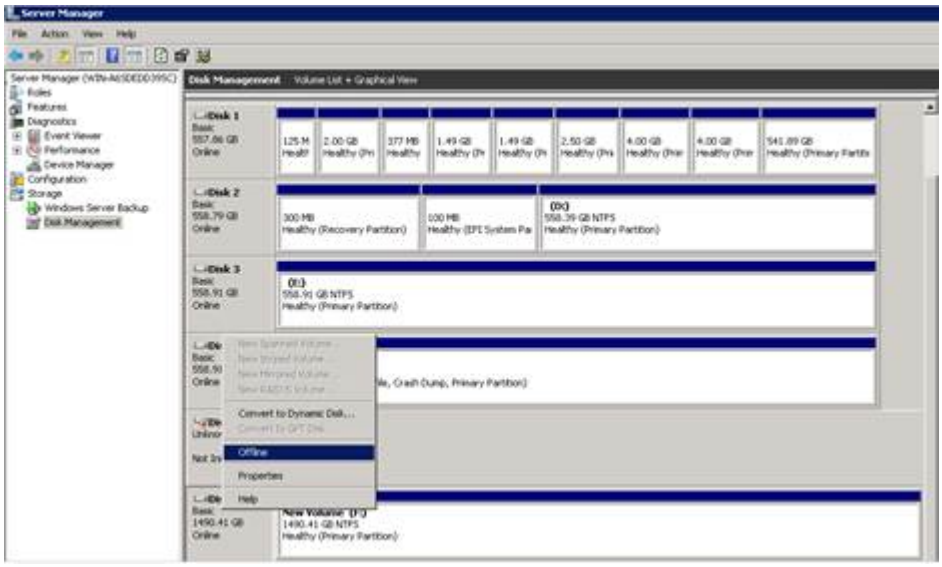

NOTICE

Do not read or write data during this operation. Otherwise, I/O errors may occur.

Step 3 Go to the directory where the tool package is located and open the Windows CLI, as shown in Figure 4-3.

Figure 4-3 Opening the Windows CLI

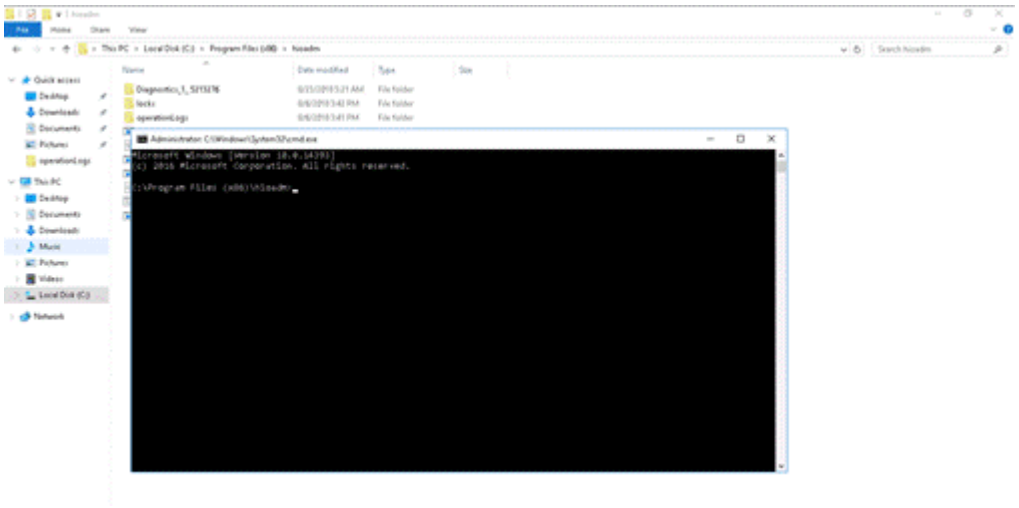

Step 4 Run the following command to instruct the OS to power off the device:

hioadm hotplug -d *devicename* **-r**

In this command, *devicename* indicates the name of the device to be removed, such as **nvme0** and **nvme1**.

NOTE

After the previous command is executed, you can remove the drive if the green indicator is off and the yellow indicator blinks at the frequency of 0.5 Hz.

----End

4.2.3 Orderly Hot Swap of a Drive in ESXi

Operation Scenario

Notify the ESXi OS before hot removing the ES3000 V5 SSD.

Preparations

- You have correctly installed the ES3000 V5 SSD on the server.
- You have installed the ESXi tool package. For details, see **3.7.3 Installing the Tool (ESXi)**.

Procedure

Hot-Removing an SSD

Step 1 Log in to the OS as an administrator or administrator group member.

Step 2 Stop all services accessing the ES3000 V5 SSD.

Step 3 Run the following command to instruct the OS to power off the device:

hioadm hotplug -d *devicename* **-r**

In this command, *devicename* indicates the name of the device to be removed, such as **nvme0** and **nvme1**.

NOTE

After the previous command is executed, you can remove the SSD if the green indicator is off and the yellow indicator blinks at the frequency of 0.5 Hz.

----End

Hot-Inserting an SSD

Step 1 Insert the ES3000 V5 into the target slot.

Step 2 Log in to the OS as an administrator or administrator group member.

Step 3 After the SSD is started, run the **hioadm info** command to check its name, such as **nvme0**.

Step 4 Run the following command to instruct the OS to power on the device:

hioadm hotplug -d *devicename* **-i**

In this command, *devicename* indicates the name of the device that has been inserted, such as **nvme0**.

----End

4.3 Surprise Hot Swap

NOTICE

Before performing surprise hot swap, ensure that:

- The server where the ES3000 V5 installed is a V5 server.
- The VMD function is enabled in the BIOS.

For details about surprise hot swap operations, see the server version requirements and operation guide of each vendor.

5 Maintenance and Upgrade

About This Chapter

This section describes how to perform hardware maintenance, software upgrades, and troubleshooting for the ES3000 V5.

5.1 Uninstalling the Driver

This section describes how to uninstall the ES3000 V5 driver on various operating systems (OSs).

5.1.1 Uninstalling the Driver in Linux

Operation Scenario

If you want to stop using or replace the ES3000 V5, uninstall the driver before proceeding.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Stop or exit the programs or files that are accessing the ES3000 V5.

Step 3 Run the uninstallation command based on the OS type.

- For RHEL, CentOS, or OEL, run the **rpm -e kmod-hiodriver** command.
- For SLES or UVP, run the **rpm -e hiodriver-kmp-*kernel feature*** command. Example: **rpm -e hiodriver-kmp-default**
- For Ubuntu or Debian, run the **dpkg -r hiodriver** command.
- For Citrix:
  a. Run the **rpm -qa | grep hiodriver** command to query the hiodriver package.
  b. Run the **rpm -e** *xxxxxx* command to uninstall the ISO packages.
  c. Run the **rpm -qa | grep hiodriver** command to check whether the hiodriver package is uninstalled.

Step 4 Check whether the following message is displayed. The message indicates that the installed NVMe driver fails to be uninstalled.

```
Warning: fail to uninstall the installed nvme kernel module. Please uninstall the
installed nvme kernel module manually or reboot the system.
```

- If yes, go to **Step 7**.
- If no, go to **Step 5**.

Step 5 Check whether the following message is displayed. The message indicates that the NVMe driver built in the OS fails to be loaded.

```
Warning: fail to install the in-kernel nvme kernel module. Please install the in-
kernel nvme kernel module manually or reboot the system.
```

- If yes, go to **Step 6**.
- If no, go to **Step 10**.

Step 6 Check whether restarting the OS will affect services.

- If yes, go to **Step 9**.
- If no, restart the OS and go to **Step 10**.

Step 7 Check whether restarting the OS will affect services.

- If yes, go to **Step 8**.
- If no, restart the OS and go to **Step 10**.

Step 8 Run the **modprobe -r nvme** command to uninstall the driver. During the uninstallation, check whether a message is displayed indicating that the driver is in use. If yes, the uninstallation fails.

- If yes, restart the OS and go to **Step 10**.
- If no, go to **Step 9**.

Step 9 Run the **modprobe nvme** command to install the open-source driver. Check whether a warning or error is displayed during the installation.

- If yes, restart the OS and go to **Step 10**.
- If no, go to **Step 10**.

Step 10 Run the **modinfo nvme** command to check the driver author information to determine whether the driver is successfully installed.

The following command output indicates that the driver is successfully uninstalled.

```
modinfo: could not find module nvme
```

The following author in the command output also indicates that the driver is successfully uninstalled.

```
author:         Matthew Wilcox <willy@linux.intel.com>
```

- If yes, no further action is required.
- If no, go to **Step 11**.

Step 11 Contact Huawei technical support for help. For details, see **H Getting Help**.

----End

5.1.2 Uninstalling the Driver in Windows

Operation Scenario

If you want to stop using or replace the ES3000 V5, uninstall its driver first.

CAUTION

Uninstall the ES3000 V5 driver by following the procedure described in this section.

Do not uninstall the ES3000 V5 driver by using Device Manager. If you have uninstalled the ES3000 V5 driver by using Device Manager, you still need to uninstall the driver by following the procedure described in this section before installing the driver again. To install the driver, see 3.6.2 Installing the ES3000 V5 Driver in Windows.

Uninstalling the Driver by Using Control Panel

Step 1 Log in to the OS as an administrator or a member of the administrator group.

Step 2 Stop the programs that are accessing the ES3000 V5.

Step 3 Choose **Start** > **Control Panel**.

The **Control Panel** window is displayed.

Step 4 Click Uninstall a program.

The **Uninstall or change a program** window is displayed, as shown in **Figure 5-1**.

Figure 5-1 Uninstall or change a program

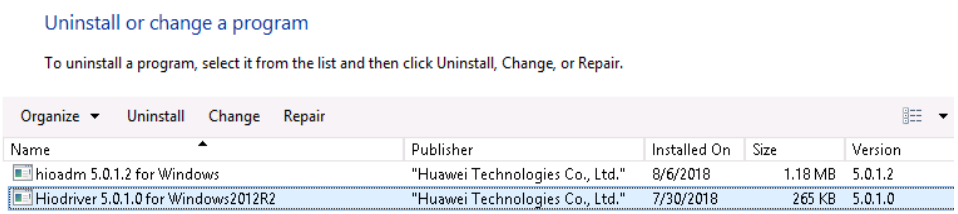

Step 5 Right-click the ES3000 V5 driver, and choose **Uninstall/Change** from the shortcut menu.

A confirmation dialog box is displayed, as shown in **Figure 5-2**.

Figure 5-2 Confirming the uninstallation

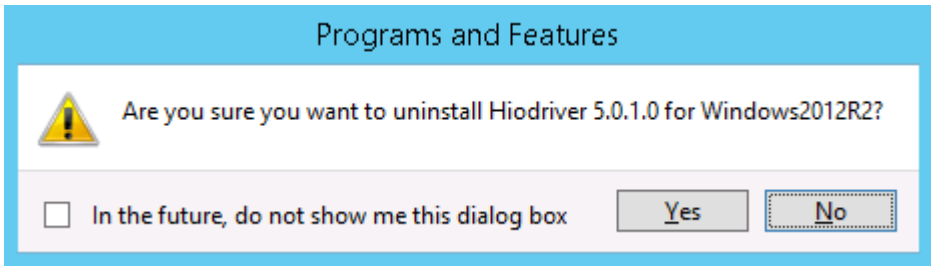

Step 6 Click **Yes** to uninstall the ES3000 V5 driver and its components.

----End

Uninstalling the Driver by Using the Installation Package

Step 1 Log in to the OS as an administrator or a member of the administrator group.

Step 2 Stop the programs that are accessing the ES3000 V5.

Step 3 Double-click the driver installation file in .msi format and select the uninstallation option.

Uninstall the driver by following on-screen instructions.

----End

5.1.3 Uninstalling the Driver in ESXi

Operation Scenario

If you want to stop using or replace the ES3000 V5, uninstall the driver first.

Procedure

Step 1 Log in to the OS as an administrator or a member of the administrator group.

Step 2 Stop or exit the programs or files that are accessing the ES3000 V5.

Step 3 Run the following command to uninstall the ES3000 V5 driver package:

esxcli software vib remove -n hiodriver

Step 4 After the driver is uninstalled, run the **reboot** command to restart the OS.

Step 5 After the restart, run the **vmkload_mod -l | grep hiodriver** command. The following command output indicates that the driver has been uninstalled:

```
[root@esxi113:~] vmkload_mod -l | grep hiodriver
[root@esxi113:~]
```

----End

5.2 Upgrading the Driver

This section describes how to upgrade the ES3000 V5 driver in different operating systems (OSs).

5.2.1 Upgrading the Driver in Linux

Operation Scenario

You can upgrade the driver after a new version is released.

The Linux driver package naming rules are as follows:

- RHEL, CentOS, or OEL driver package name: **kmod-hiodriver-*Driver version-Supported kernel version.System distributor.Hardware platform*.rpm**, such as **kmod-hiodriver-5.0.3.3-3.10.0_123.el7.centos.x86_64.rpm**, **kmod-hiodriver-5.0.3.3-3.10.0_123.el7.redhat.x86_64.rpm**, or **kmod-hiodriver-5.0.3.3-3.10.0_123.el7.oracle.x86_64.rpm**.
- SLES driver package name: **hiodriver-kmp-*Kernel feature-Driver version-Supported kernel version.System distributor.Hardware platform*.rpm**, such as **hiodriver-kmp-default-5.0.3.3-3.12.28_4.suse12.x86_64.rpm**.
- Ubuntu or Debian driver package name: **hiodriver -*Driver version-Supported kernel version.System distributor.Hardware platform*.deb**, such as **hiodriver-5.0.3.3-3.19.0_15_generic.ubuntu.amd64.deb**.
- UVP driver package name: **hiodriver-kmp-xen-*Driver version-Supported kernel version.System distributor.Hardware platform*.rpm**.

Preparations

Visit http://support.huawei.com/enterprise, choose **Product Support** > **Enterprise Data Center** > **Servers** > **Accelerator Components** > **ES3000 V5**, click the **Software Download** tab, and download the software package. The software package contains host tools such as hioadm and Huawei NVMe drivers.

Before upgrading the driver, check that:

- You have downloaded the new driver package to the server and verified the package integrity.
- You have checked the new driver version.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Go to the directory that stores the driver package, for example, **/home**.

Step 3 Run the upgrade command based on the OS type.

- For RHEL, CentOS, OEL, or SLES, run **rpm -Uvh** *driver package name*.

  An example package name for CentOS is **rpm -Uvh kmod-hiodriver-5.0.3.3-3.10.0_229.el76.centos.x86_64.rpm**.

- For Ubuntu or Debian, run **dpkg -i** *driver package name*.

  An example package name for Ubuntu is **dpkg -i hiodriver-5.0.3.3-3.13.0_24_generic.ubuntu.amd64.deb**.

Step 4 Check whether the following message indicating an incorrect upgrade package is displayed.

```
Error: The installing package does not match the current OS. Please execute
command: rpm -e kmod-hiodriver to uninstall the installing rpm.
```

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Uninstall the installed package as prompted. Upload the correct upgrade package for the OS to the server and perform **Step 2** again to install the package.

Step 6 Check whether the following message is displayed. The message indicates that the original NVMe driver fails to be uninstalled during the upgrade.

```
Warning: fail to uninstall the old nvme kernel module.  Please uninstall the old
nvme kernel module manually or reboot the system.
```

- If yes, go to **Step 9**.
- If no, go to **Step 7**.

Step 7 Check whether the following message is displayed. The message indicates that the new NVMe driver fails to be installed during the upgrade.

```
Warning: fail to install the new nvme kernel module. Please install the new nvme
kernel module manually or reboot the system.
```

- If yes, go to **Step 8**.
- If no, go to **Step 12**.

Step 8 Check whether restarting the OS will affect services.

- If yes, go to **Step 11**.
- If no, restart the OS and go to **Step 12**.

Step 9 Check whether restarting the OS will affect services.

- If yes, go to **Step 10**.
- If no, restart the OS and go to **Step 12**.

Step 10 Run the **modprobe -r nvme** command to uninstall the original driver. During the uninstallation, check whether a message is displayed indicating that the driver is in use. If yes, the uninstallation fails.

- If yes, restart the OS and go to **Step 12**.
- If no, go to **Step 11**.

Step 11 Run the **modprobe nvme** command to install the new driver. Check whether a warning or error is displayed during the installation.

- If yes, restart the OS and go to **Step 12**.
- If no, go to **Step 12**.

Step 12 Run the **modinfo nvme** command to check the version information to determine whether the driver is successfully upgraded.

If the version number in the command output is the target version number, the upgrade is successful.

The following is an example command:


```
version:        5.0.3.3
```

- If yes, go to **Step 13**.
- If no, go to **Step 15**.

Step 13 If an SSD device has been installed on the server, check whether the driver has initialized the NVMe device.

Check whether the NVMe device exists in the **/dev** directory.

- If yes, go to **Step 14**.
- If no, go to **Step 15**.

Step 14 if the OS uses the irqbalance service to balance CPU interrupts, restart the service by using one of the following commands after re-loading the driver module:

- For RHEL 6 or CentOs 6, run **service irqbalance restart**.
- For SLES 11, run **service irq_balancer restart**.
- For SLES 12 or RHEL 7, run **systemctl restart irqbalance.service**.

NOTE

- For RHEL 6, CentOS 6, and SLES 12, you can add IRQBALANCE_ARGS="--hintpolicy=exact" to the end of the **/etc/sysconfig/irqbalance** file content to better balance interrupts.
- To configure the irqbalance service in other OSs, refer to documents related to the OSs.

Check whether errors are reported after running the command for restarting the irqbalance service.

- If yes, go to **Step 15**.
- If no, no further action is required.

Step 15 Contact Huawei technical support for help. For details, see **H Getting Help**.

----End

5.2.2 Upgrading the ES3000 V5 Driver in Windows

Operation Scenario

You can upgrade the ES3000 V5 driver after a later version is released.

Preparations

Before upgrading the driver, check that:

- You have downloaded the later driver software to the server, and verified the package integrity.
- You have checked the driver version.

Procedure

Step 1 Uninstall the existing driver. For details, see **5.1.2 Uninstalling the Driver in Windows**.

Step 2 Install the new driver. For details, see **3.6.2 Installing the ES3000 V5 Driver in Windows**.

----End

5.2.3 Upgrading the Driver in ESXi

Operation Scenario

You can upgrade the driver after a new version is released.

Preparations

Before upgrading the driver, download the new driver package to the server and verify the package integrity.

Procedure

Step 1 Run the **esxcli software vib list|grep hiodriver** command to check whether the driver version is the latest..

Step 2 Install the new driver. For details, see **3.6.3 Installing the ES3000 V5 Driver in ESXi**.

NOTE

Install the new driver package directly. The system will automatically replace the old driver.

----End

5.3 Uninstalling the Tool Package

5.3.1 Uninstalling the Tool Package in Linux

Operation Scenario

Uninstall the tool package to free up space when ES3000 V5 management is not needed.

Procedure

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Run the uninstallation command based on the OS type.

- For RHEL, CentOS, OEL, or SLES, run the **rpm -e hioadm** command.
- For Ubuntu or Debian, run the **dpkg -r hioadm** command.
- For Citrix:
  a. Run the **rpm -qa | grep hioadm** command to query the hioadm package.
  b. Run the **rpm -e** *xxxxxx* command to uninstall the ISO packages.
  c. Run the **rpm -qa | grep hioadm** command to check whether the hioadm package is uninstalled.

----End

5.3.2 Uninstalling the Tool Package in Windows

Operation Scenario

Uninstall the tool package to free up space when ES3000 V5 management is not needed.

Procedure

Step 1 Log in to the operating system (OS) as an administrator or administrator group member.

Step 2 Choose **Start** > **Control Panel**.

The **Control Panel** window is displayed.

Step 3 Click **Uninstall a program**.

The **Uninstall or change a program** window is displayed, as shown in **Figure 5-3**.

Figure 5-3 Uninstall or change a program

Step 4 Find and right-click the tool package, and choose **Uninstall**.

A confirmation dialog box is displayed, as shown in **Figure 5-4**.

Figure 5-4 Confirming the uninstallation

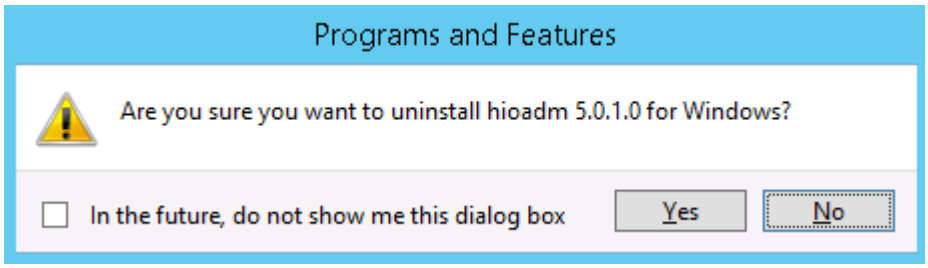

Step 5 Click **Yes**.

----End

5.3.3 Uninstalling the Tool in ESXi

The tool is integrated into the driver package. Uninstalling the driver package also uninstalls the tool. You can uninstall the driver package by referring to **5.1.3 Uninstalling the Driver in ESXi**.

5.4 Upgrading the Tool Package1

You can download the software package as follows: log in to Huawei Enterprise support website **http://support.huawei.com/enterprise**, choose **Products** > **Cloud Data Center** > **Servers** > **Intelligent Accelerator Components** > **ES3000 V5 SAS SSD Disk** > **Technical Support**, click **Download Software**, and download the required software package. The software package contains host tools (including hioadm) and Huawei's NVMe driver.

5.4.1 Upgrading the Tool Package in Linux

Operation Scenario

Upgrade the tool package to better manage the ES3000 V5 when a new version is available.

The tool package naming rules are as follows:

- For RHEL, CentOS, OEL, or SLES, the package name is hioadm-*version number-release number.hardware architecture*.rpm, for example, **hioadm-5.0.3.5-1.x86_64.rpm**.
- For Ubuntu or Debian, the package name is hioadm-*version number-release number.hardware architecture*.deb, for example, **hioadm-5.0.3.5-1.amd64.deb**.

Preparations

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the new tool package version and verified its integrity.
- You have used WinSCP or other software to upload the tool package to the server, such as the **/home** directory.

Procedure

Step 1 Log in to the operating system (OS) as an administrator.

Step 2 Go to the directory that stores the tool package, for example, **/home**.

Step 3 Run the upgrade command based on the OS type.

- For RHEL, CentOS, OEL, or SLES, run **rpm -Uvh** *tool package name*.

  Example: **rpm -Uvh hioadm-5.0.3.5-1.x86_64.rpm**
- For Ubuntu or Debian, run **dpkg -i** *tool package name*.

  Example: **dpkg -i hioadm-5.0.3.5-1.amd64.deb**

Step 4 Run the following command to obtain help information about the tool:

hioadm -h or **hioadm --help**

----End

5.4.2 Upgrading the Tool Package in Windows

Operation Scenario

Upgrade the tool package to better manage the ES3000 V5 when a new version is available.

Two upgrade methods are available:

- Uninstall the old version and install the new version.
- Upgrade the old version directly.

Preparations

- You have correctly installed the ES3000 V5 on the server.
- You have downloaded the new tool package version and verified its integrity.
- You have uploaded the tool package to the server by using Windows Secure Copy (WinSCP) or other software.

Procedure

Step 1 Uninstall the old tool package. For details, see **5.3.2 Uninstalling the Tool Package in Windows**.

Step 2 Install the new tool package. For details, see **3.7.2 Installing the Tool Package (Windows)**.

----End

5.4.3 Upgrading the Tool in ESXi

Operation Scenario

Upgrade the tool to better manage the ES3000 V5 when a new tool version is available.

Procedure

Step 1 Run the **hioadm -v** command to check whether the current tool version is the latest.

Step 2 Install the latest version. For details, see **3.7.3 Installing the Tool (ESXi)**.

NOTE

The tool upgrade procedure is the same as the tool installation procedure. The system will automatically use the new version to replace the earlier version.

----End

5.5 Upgrading the Controller Firmware

This section describes how to upgrade the ES3000 V5 controller firmware in different OSs.

NOTE

The ES3500 and ES3600P series use the same drivers and firmware.

You can also use eSight to upgrade the ES3000 V5 controller firmware. For details, see the esight user guide.

5.5.1 Upgrade Paths

Only ES3000 V5 SSDs are supported.

NOTICE

ES3000 V5 V100R001SPC100 (firmware version: 1010) can be upgraded in reset activation mode only. Versions later than 1010 can be upgraded either in reset activation or immediate activation mode.

5.5.2 Upgrading the SSD Firmware in Linux (PCIe SSD)

Operation Scenario

You can upgrade the controller firmware when a new version is released.

Preparations

Before upgrading the firmware, check that:

- Stop or exit the ES3000 V5 SSD applications or services, and unmount the file system that uses the SSD device.
- The firmware upgrade does not damage data on SSDs. However, you are advised to back up important data on SSDs.
- You have downloaded the firmware package to the server, and verified the package integrity.
- You have uploaded the controller firmware file to the server.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Run the following command to view the name of the SSD to be upgraded:

hioadm info

The command output contains information similar to the following:

```
NVMe SSD Information
      Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
      |------------ nvme0n1 (512.0GB)
|---- nvme1 (0503023HDCN107C80014)
      |------------ nvme1n1 (512.0GB)
```

Step 3 Run the following command to view the controller firmware version of the SSD. (**nvme0** is used as an example. If there are multiple PCIe SSDs, perform steps 3 to 8 for each SSD.)

hioadm updatefw -d nvme0

The command output contains information similar to the following:

```
slot version activation
1     1011
2     1011    current
```

NOTE

In the **activation** column, **current** indicates the firmware version in use, whereas **next** indicates the firmware version that takes effect after next reset.

Step 4 Go to the directory where the controller firmware resides, for example, **/home**.

NOTE

Firmware upgrade supports either an absolute path or a relative path. If you skip this step, you need to enter an absolute path in commands in next steps.

Step 5 Upgrade controller firmware. It is assumed that the upgrade file is **ES3000V5_FW_2151.bin**.

Select an upgrade mode based on the current SSD firmware version. For details, see **5.5.1 Upgrade Paths**.

Reset activation mode:

Run the **hioadm updatefw -d nvme0 -f ES3000V5_FW_2151.bin** command and perform steps 6 to 8.

Immediate activation mode:

Run the **hioadm updatefw -d nvme0 -f ES3000V5_FW_2151.bin -a 1** command and perform steps 6 and 8.

Step 6 Input **y** and press **Enter** to confirm the upgrade.

Step 7 Restart the OS after the upgrade.

Step 8 Run the following command to check whether the SSD controller firmware is upgraded:

hioadm updatefw -d nvme0

----End

5.5.3 Upgrading the SSD Firmware in Windows (PCIe SSD)

Operation Scenario

You are advised to upgrade the controller firmware immediately after a new version is released.

Preparations

Before upgrading the firmware, check that:

- You have uploaded the ES3000 V5 software package to a directory on the server, for example, the **D:** directory.
- Stop or exit the ES3000 V5 SSD applications or services, and uninstall the file system that uses the SSD device (make the to-be-upgraded SSD to be offline in the disk manager).
- The firmware upgrade does not damage data on SSDs. However, you are advised to back up important data on SSDs.
- You have downloaded the latest firmware file to a specified folder on the server and verified the file integrity.
- You have upgraded the firmware driver to a compatible version.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Choose **Start** > **Run**.

The Windows CLI is displayed.

Step 3 Run the **cd** command to go to the directory where the **hioadm.exe** file is stored.

Step 4 Run the following command to view the name of the SSD to be upgraded:

hioadm.exe info

The command output contains information similar to the following:

```
NVMe SSD Information
      Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
      |------------ nvme0n1 (512.0GB)
|---- nvme1 (0503023HDCN107C80014)
      |------------ nvme1n1 (512.0GB)
```

Step 5 Run the following command to view the controller firmware version of the SSD. (**nvme0** is used as an example. If there are multiple PCIe SSDs, perform steps 5 to 9 for each SSD.)

hioadm.exe updatefw -d nvme0

The command output contains information similar to the following:

```
slot version activation
1     1011
2     1011    current
```

NOTE

In the **activation** column, **current** indicates the firmware version in use, whereas **next** indicates the firmware version that takes effect after next reset.

Step 6 Move the firmware upgrade file, for example, **ES3000V5_FW_2151.bin**, to the path where **hioadm.exe** resides. Then select an upgrade mode based on the current SSD firmware version. For details, see **5.5.1 Upgrade Paths**.

Reset activation mode:

Run the **hioadm.exe updatefw -d nvme0 -f ES3000V5_FW_2151.bin** command and perform steps 7 to 9.

Immediate activation mode:

Run the **hioadm.exe updatefw -d nvme0 -f ES3000V5_FW_2151.bin -a 1** command and perform steps 7 and 9.

Step 7 Input **y** and press **Enter** to confirm the upgrade.

Step 8 Restart the OS after the upgrade.

Step 9 Re-log in to the OS, run the **hioadm.exe** file again, and run the following command to check whether the SSD controller firmware is upgraded:

hioadm.exe updatefw -d nvme0

----End

5.5.4 Upgrading the SSD Firmware in VMware (PCIe SSD)

Operation Scenario

You can upgrade the controller firmware when a new version is released.

Preparations

Before upgrading the firmware, check that:

- Stop or exit the ES3000 V5 SSD applications or services.
- The firmware upgrade does not damage data on SSDs. However, you are advised to back up important data on SSDs.
- You have downloaded the firmware package to the server, and verified the package integrity.
- You have uploaded the controller firmware file to the server.

Procedure

Step 1 Log in to the OS as an administrator.

Step 2 Run the following command to view the name of the SSD to be upgraded:

hioadm info

The command output contains information similar to the following:

```
NVMe SSD Information
      Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
      |------------ nvme0n1 (512.0GB)
|---- nvme1 (0503023HDCN107C80014)
      |------------ nvme1n1 (512.0GB)
```

Step 3 Run the following command to view the controller firmware version of the SSD. (**nvme0** is used as an example. If there are multiple PCIe SSDs, perform steps 3 to 8 for each SSD.)

hioadm updatefw -d nvme0

The command output contains information similar to the following:

```
slot version activation
1     1011
2     1011    current
```

NOTE

In the **activation** column, **current** indicates the firmware version in use, whereas **next** indicates the firmware version that takes effect after next reset.

Step 4 Go to the directory where the controller firmware resides, for example, **/home**.

NOTE

Firmware upgrade supports either an absolute path or a relative path. If you skip this step, you need to enter an absolute path in commands in next steps.

Step 5 Upgrade controller firmware. It is assumed that the upgrade file is **ES3000V5_FW_2151.bin**.

Select an upgrade mode based on the current SSD firmware version. For details, see **5.5.1 Upgrade Paths**.

Reset activation mode:

Run the **hioadm updatefw -d nvme0 -f ES3000V5_FW_2151.bin** command and perform steps 6 to 8.

Immediate activation mode:

Run the **hioadm updatefw -d nvme0 -f ES3000V5_FW_2151.bin -a 1** command and perform steps 6 and 8.

Step 6 Input **y** and press **Enter** to confirm the upgrade.

Step 7 Restart the OS after the upgrade.

Step 8 Run the following command to check whether the SSD controller firmware is upgraded:

hioadm updatefw -d nvme0

----End

5.6 Querying Health Status

Operation Scenario

Check the SSD health status.

Procedure

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the following command to query the health status of the SSD controller (such as nvme0):

hioadm info -d nvme0

device status in the command output indicates the SSD controller health status.

- **healty** indicates that the SSD is healthy.
- **warning** indicates that the SSD is faulty.

```
[root@localhost tool]# hioadm info -d nvme0

Namespace<1> size:  3200.6GB,    3200631791616Byte
             formatted LBA size:         512 Byte
             formatted metadata size:    0 Byte

maximum capacity                : 3200.6GB
current capacity                : 3200.6GB
volatile write cache            : Disable
serial number                   : SS8210xxxxxxxx918002
model number                    : HWE52P436T4M002N
firmware version                : 1010
NVMe version                    : 1.3
device status                   : healthy
```

----End

5.7 Information Collection

5.7.1 Collecting SSD Information

Operation Scenario

Collect SSD information when maintaining SSDs.

Procedure

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the following command to query basic information about the SSD controller (such as nvme0):

hioadm info -d nvme0

The command output is similar to the following:

```
[root@localhost tool]# hioadm info -d nvme0

Namespace<1> size:  3200.6GB,    3200631791616Byte
             formatted LBA size:         512 Byte
             formatted metadata size:    0 Byte

maximum capacity                : 3200.6GB
current capacity                : 3200.6GB
volatile write cache            : Disable
serial number                   : 0503023HDCN107C80013
model number                    : HWE32P430016M00N
firmware version                : 1010
NVMe version                    : 1.3
device status                   : healthy
```

Step 3 Run the following command to query the SMART information of the SSD controller (such as nvme0):

hioadm info -d nvme0 -s

The command output is similar to the following:

```
[root@localhost tool]# hioadm info -d nvme0 -s
critical warning              : no warning
composite temperature         : 308 degrees Kelvin (35 degrees Celsius)
available spare               : 100%
available spare threshold     : 10%
percentage used               : 0%
data units read               : 41908532.1 MB
data units written            : 31591917.9 MB
host read commands            : 661401547
host write commands           : 571198332
controller busy time          : 481 mins
power cycles                  : 414 times
power on hours                : 8852 h
unsafe shutdowns              : 250 times
media and data integrity errors: 0
number of error information log entries: 0
warning composite temperature time: 0 min
critical composite temperature time: 0 min
data status                   : OK
```

----End

5.7.2 Collecting SSD Logs

Operation Scenario

Collect SSD information when maintaining SSDs.

Procedure

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the following command to collect the SSD controller (such as nvme0) log information:

hioadm log -d nvme0 -a

The command output shows the log collecting result and file path.

```
[root@localhost tool]# hioadm log -d nvme0 -a
firmware key log saving..........OK
save firmware key log file succeeded(/opt/hio/nvme0_key.log)
firmware all log saving..........OK
save firmware all log file succeeded(/opt/hio/nvme0_fw.log)
firmware core dump log saving..........OK
save firmware core dump log file succeeded(/opt/hio/nvme0_coredump.log)
other log saving..........OK
save other logs file succeeded(/opt/hio/nvme0_other.log)
logs has been saved.
```

----End

5.7.3 One-Click Log Collection

Scenario

Collect SSD and system operation logs when maintaining SSD devices.

Prerequisite

You have downloaded the ES3000_V5_Tool_x.x.x.x.zip package from the **Software Download** tab at **http://support.huawei.com/enterprise** (**Product Support** > **Enterprise Data Center** > **Servers** > **Accelerator Components** > **ES3000 V5**).

You have already installed hioadm, which integrates one-click log collection scripts.

Linux

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the **bash /opt/hio/host_info_collect.sh** command.

Step 3 Obtain the log file **hostInfoCollect.zip** from the **/opt/hio** directory.

Step 4 (Optional) Delete the log file after use to free drive space.

----End

Windows

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Double-click **host_info_collect.bat** in **C:\Program Files (x86)\hioadm**.

Step 3 Obtain all log files from the **C:\Program Files (x86)\hioadm** directory.

Step 4 (Optional) Delete the log files after use to free drive space.

----End

ESXi

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the **/opt/hio/host_info_collect.sh** command.

Step 3 Obtain all log files from the **/scratch/hio/okc** directory.

Step 4 (Optional) Delete the log files after use to free drive space.

----End

5.8 Securely Erasing SSD Data

Operation Scenario

Collect SSD logs when maintaining SSD devices.

Preparations

Before performing operations, stop all I/O services of the SSD and perform the following:

- For Linux OSs, unmount all logical drives.
- For ESXi OSs, delete all partitions of all logical drives.
- For Windows OSs, make all logical drives offline, as shown in **Figure 5-5**.

Figure 5-5 Making all logical drives offline

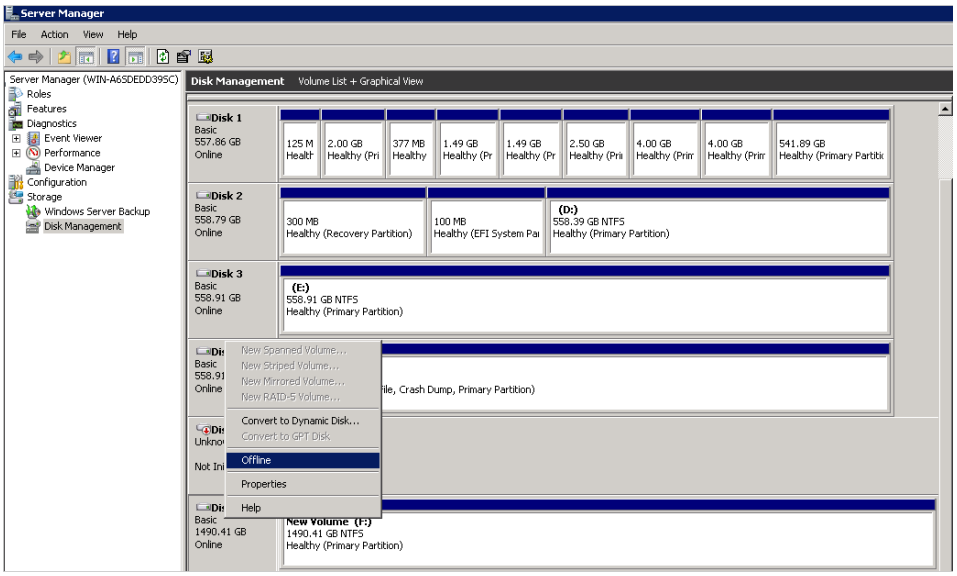

Procedure

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the following command to securely erase data from the SSD (such as nvme0):

hioadm secure -d nvme0

The warning of the risks of data erasure is displayed.

```
[root@localhost tool]# hioadm secure -d nvme0

WARNING! The secure operation will clear all data.
WARNING! You have selected the format operation.
Proceed with the secure? (Y|N):
```

Step 3 After confirming the risks, enter **y** to start the secure erasure.

The operation result is displayed.

```
[root@localhost tool]# hioadm secure -d nvme0

WARNING! The secure operation will clear all data.
WARNING! You have selected the format operation.
Proceed with the secure? (Y|N): y
Format Progress...100%
The secure operate succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

----End

5.9 Querying SSD Lifespans

Scenario

Query the lifespan of an SSD.

Procedure

Step 1 Log in to the OS as an administrator or a member in the administrator group.

Step 2 Run the following command to query the SMART information of the SSD controller (such as nvme0):

hioadm info -d nvme0 -s

The command output is similar to the following:

```
[root@localhost tool]# hioadm info -d nvme0 -s
critical warning              : no warning
composite temperature         : 308 degrees Kelvin (35 degrees Celsius)
available spare               : 100%
available spare threshold     : 10%
percentage used               : 0%

data units read               : 41908532.1 MB
data units written            : 31591917.9 MB
host read commands            : 661401547
host write commands           : 571198332
controller busy time          : 481 mins
power cycles                  : 414 times
```



```
power on hours                : 8852 h
unsafe shutdowns              : 250 times
media and data integrity errors: 0
number of error information log entries: 0
warning composite temperature time: 0 min
critical composite temperature time: 0 min
data status                   : OK
```

The **percentage used** parameter displays the used lifespan percentage of the SSD.

CAUTION

If the value of **percentage used** exceeds 100%, the SSD lifespan has expired, and Huawei's warranty service is unavailable. You need to back up data on the SSD as soon as possible and replace the SSD.

----End

5.10 Common Faults

5.10.1 Hardware Connection Faults or Other Faults

Hardware connection faults are mostly identified by indicators. If the yellow indicator of a device is on and the green indicator is off 30 seconds after a new drive is installed or during proper operating, hardware faults occur. See **5.11 Indicator Status Description** to identify the faults.

Typical fault causes are as follows:

- The drive is not securely installed.
- The drive is installed in a non-NVMe slot.
- A drive hardware port fault occurs.

5.10.2 Drive Identification Faults or Other Faults

Drive identification faults indicate scenarios where the host OS disk manager (for example, fdisk-l in Linux and the device manager in Windows) cannot properly identify ES3000 SSDs when hardware connections are correct. If drive identification faults occur, collect all logs. For details, see **5.7.3 One-Click Log Collection**.

Typical fault causes are as follows:

- The NVMe driver is not installed on the OS.
- OS driver loading fails.

5.10.3 Drive Internal Faults or Other Faults

Drive internal faults indicate scenarios where drive internal alarms occur while the OS can properly identify drives. If such faults occur, the drives do not function properly or cannot be used. This section uses an example to describe how to identify such faults.

Use Instance

Step 1 Determine the drive health status. For details, see **6.1.3 Querying Basic Information About a Device6.1.3 Querying Basic Information About a Device**.

Query the basic information about **nvme0**.

```
[root@localhost tool]# hioadm info -d nvme0

Namespace<1> size:  3200.6GB,    3200631791616Byte
             formatted LBA size:         512 Byte
             formatted metadata size:    0 Byte

maximum capacity              : 3200.6GB
current capacity              : 3200.6GB
volatile write cache          : Disable
serial number                 : 0503023HDCN107C80013
model number                  : HWE32P43032M000N
firmware version              : 1010
NVMe version                  : 1.3
device status                 : healthy
```

device status in the command output indicates the SSD controller health status.

The value **healthy** indicates that the device is healthy.

The value **warning** indicates that an exception occurs on the device. For details, see **Step 2**.

Step 2 Determine the SMART status of the drive. For details, see **6.1.3 Querying Basic Information About a Device6.1.3 Querying Basic Information About a Device**.

Query the SMART information about **nvme0**.

```
[root@localhost tool]# hioadm info -d nvme0 -s
critical warning              : no warning
composite temperature         : 308 degrees Kelvin (35 degrees Celsius)
available spare               : 100%
available spare threshold     : 10%
percentage used               : 0%
data units read               : 41908532.1 MB
data units written            : 31591917.9 MB
host read commands            : 661401547
host write commands           : 571198332
controller busy time          : 481 mins
power cycles                  : 414 times
power on hours                : 8852 h
unsafe shutdowns              : 250 times
media and data integrity errors: 0
number of error information log entries: 0
warning composite temperature time: 0 min
critical composite temperature time: 0 min
data status                   : OK
```

Critical Warning: Critical exceptions occur on the device, and emergency handling is required.

Table 5-1 Critical Warning parameters

Parameter	Description	Value
critical warning	Critical warning, such as overtemperature and insufficient redundant space.	Example: no warning There are the following types of critical warning: ● The available space is smaller than the threshold. ● The temperature exceeds the threshold. ● An internal error downgrades the NVM subsystem reliability. ● The medium is read-only. ● Device backup failed.

Typical fault causes are as follows:

- Available spare space below the threshold: Causes include excessive number of damaged blocks.
- Temperature above the overtemperature threshold or below the undertemperature threshold: Causes include drive temperature below 0°C or above 78°C.
- NVM subsystem reliability degraded due to internal errors: Causes include excessive number of failed drive granules or internal subsystem operating exceptions.
- Media placed in read only mode: Causes include capacitor failures.
- Volatile memory backup device failure: Causes include capacitor voltage below 28 V or above 35 V.

----End

Fault Diagnosis Instance

The available space is smaller than the threshold.

Step 1 Check whether the value of **available spare** in the SMART information exceeds 10%. If the value is greater than or equal to 10%, stop using the drive and back up data immediately. Otherwise, go to **Step 2**.

Step 2 Obtain drive logs by following instructions in **5.7.3 One-Click Log Collection** and contact Huawei technical support.

----End

The service life exceeds the threshold.

Step 1 Check whether the value of **percentage used** in the SMART information exceeds 100%, if the value is greater than or equal to 100%, stop using the drive and back up data immediately. Otherwise, go to **Step 2**.

Step 2 Collect drive log information. For details, see **5.7.3 One-Click Log Collection**. Contact Huawei technical support.

----End

The temperature exceeds the threshold.

Step 1 Check that the server where the drive resides provides proper cooling. If the server reports no temperature alarm and the fan modules are operating properly, go to **Step 2**.

Step 2 Check that the equipment room where the drive resides provides proper cooling. If the equipment room temperature is normal, go to **Step 3**.

Step 3 Obtain drive logs by following instructions in **5.7.3 One-Click Log Collection** and contact Huawei technical support.

----End

- **An internal error downgrades the NVM subsystem reliability.**
- **The medium is read-only.**
- **Device backup fails.**

Step 1 Obtain drive logs by following instructions in **5.7.3 One-Click Log Collection** and contact Huawei technical support.

----End

5.11 Indicator Status Description

5.11.1 ES3000 SSD Disk Indicator State Descriptions

The indicators on the front panel of the drive bay display the working status of the ES3000 V5, as shown in **Figure 5-6**. The amber/blue indicator is the fault indicator, and the green indicator is the working status indicator. **Table 5-2** describes the indicators when the SSD is installed in a Huawei server.

Figure 5-6 ES3500P V5/ES3600P V5 front view

Table 5-2 ES3000 SSD indicator state description

Activity Indicator (Green)	Fault Indicator (Amber/Blue)	State Description
Off	Off	The NVMe SSD disk cannot be detected.
On	Off	The NVMe SSD disk is detected and working properly.
2 Hz	Off	Data access to the NVMe SSD disk is in progress.

Activity Indicator (Green)	Fault Indicator (Amber/Blue)	State Description
Off	Blinking at 2 Hz	The NVMe SSD disk is in the hot swap process.
Off	Blinking at 0.5 Hz	The NVMe SSD disk completes the hot swap process and is removable.
On/Off	On	The NVMe SSD disk is faulty.

5.11.2 ES3000 SSD Card Indicator State Descriptions

Figure 5-7 ES3600C V5 front view

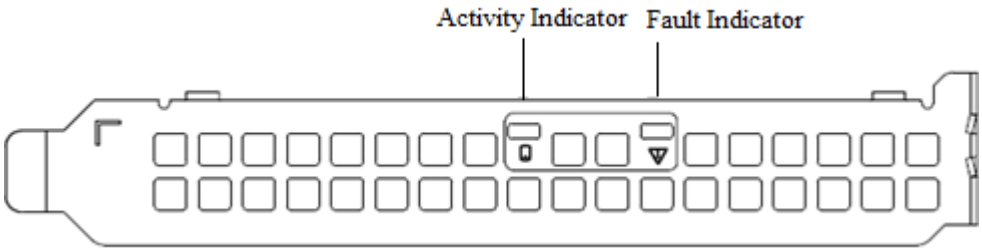

Table 5-3 ES3000 SSD card indicator state descriptions

Activity Indicator (Green)	Fault Indicator (Yellow)	Description
Off	Off	The NVMe SSD card is not powered on.
On	Off	The NVMe SSD card is working.
2 Hz	Off	Data access to the NVMe SSD card is in progress.
Off	On	Loading the NVMe SSD card firmware fails.
Off	0.5 Hz	The NVMe SSD card firmware is being initialized.
On	Blinking	Reserved
On	On	The NVMe SSD card is faulty.

5.12 Removing an ES3000 V5

Operation Scenario

- **Removing an ES3000 V5 NVMe PCIe SSD disk**
  - If the system does not support hot swap, the method for removing an NVMe PCIe SSD is the same as that for removing a common drive. For details, see **5.12.1 Removing a Non-Hot-Swappable ES3000 V5 NVMe PCIe SSD Disk**.
  - If the system supports hot swap, hot-swap an NVMe PCIe SSD disk. For details, see **5.12.2 Removing a Hot Swappable ES3000 V5 NVMe PCIe SSD Disk**.

    NOTE

    Hot swap: Users are allowed to take out and replace faulty drives, PSUs, or cards when the system is not stopping or the power is not switched off. Surprised hot swap and orderly host swap are supported. For details, see **4 Hot Swap4 Hot Swap**.
- **Removing an ES3000 V5 NVMe PCIe SSD card**
  - If PCIe card slots are on a riser card, the SSD card needs to be removed from the riser card. For details, see **5.12.3 Removing an ES3000 V5 NVMe PCIe SSD Card from a Riser Card**.
  - If PCIe card slots are on the mainboard, the SSD card needs to be removed from the mainboard. For details, see **5.12.4 Removing an ES3000 V5 NVMe PCIe SSD Card from the Mainboard**.

Preparations

Before removing an ES3000 V5 NVMe PCIe SSD, make the following preparations:

- Wear ESD gloves.

  Used to prevent ESD damage when you insert, remove, or hold the ES3000 V5 or a precision device.
- Back up data.
  - Stop the access to a hot swappable SSD disk.
  - Stop the services on the server and power off the server if you need to remove a non-hot-swappable SSD disk or an SSD card.

5.12.1 Removing a Non-Hot-Swappable ES3000 V5 NVMe PCIe SSD Disk

Prerequisites

Wear ESD gloves.

Procedure

Step 1 Determine the position of the drive in the server and power off the server.

Step 2 Press the ejector release button. See (1) in **Figure 5-8**.

Figure 5-8 Removing a drive

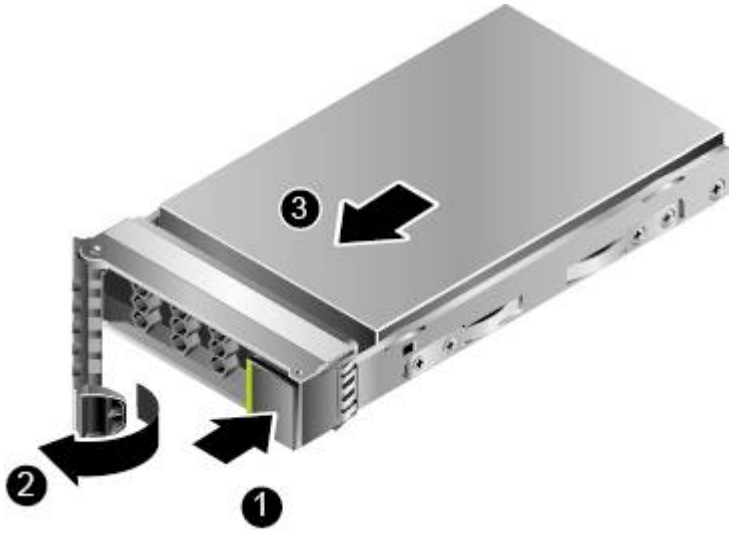

Step 3 Holding onto the drive ejector lever, pull out the drive. See (2) and (3) in **Figure 5-8**.

Step 4 Place the removed drive into an ESD bag.

----End

5.12.2 Removing a Hot Swappable ES3000 V5 NVMe PCIe SSD Disk

Prerequisites

Check that the current environment supports **4.1 OSs Supporting Hot Swap4.1 OSs Supporting Hot Swap**.

Wear ESD gloves.

Procedure

Step 1 Determine the position of the drive in the server and stop the services on the drive to be removed.

Step 2 Perform operations described in **4.2 Orderly Hot Swap**.

Step 3 Press the ejector release button. See (1) in **Figure 5-9**.

Figure 5-9 Removing a drive

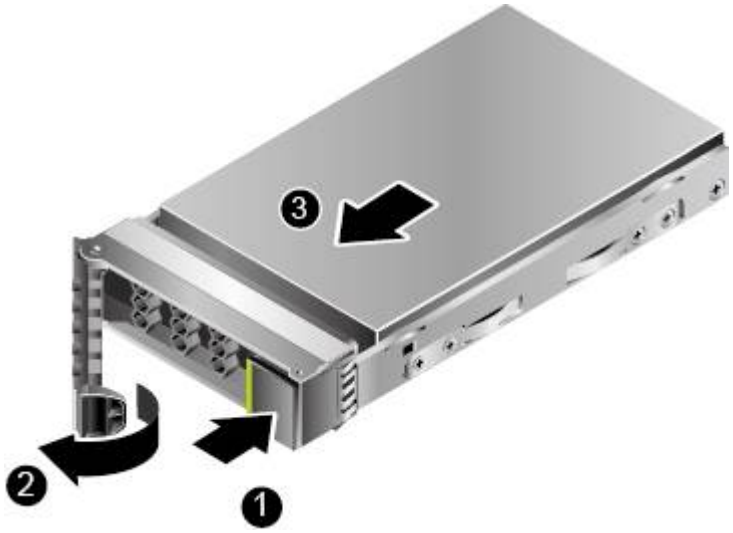

Step 4 Holding onto the drive ejector lever, pull out the drive. See (2) and (3) in **Figure 5-9**.

Step 5 Place the removed drive into an ESD bag.

----End

5.12.3 Removing an ES3000 V5 NVMe PCIe SSD Card from a Riser Card

Prerequisites

Wear ESD gloves.

Procedure

Step 1 Turn off the power to the server and remove the chassis cover.

Step 2 Remove all the cables from the riser card.

Step 3 Hold the upper part of an ES3000 V5 NVMe PCIe SSD card and remove it slowly.

Figure 5-10 Removing an ES3000 V5 NVMe PCIe SSD card

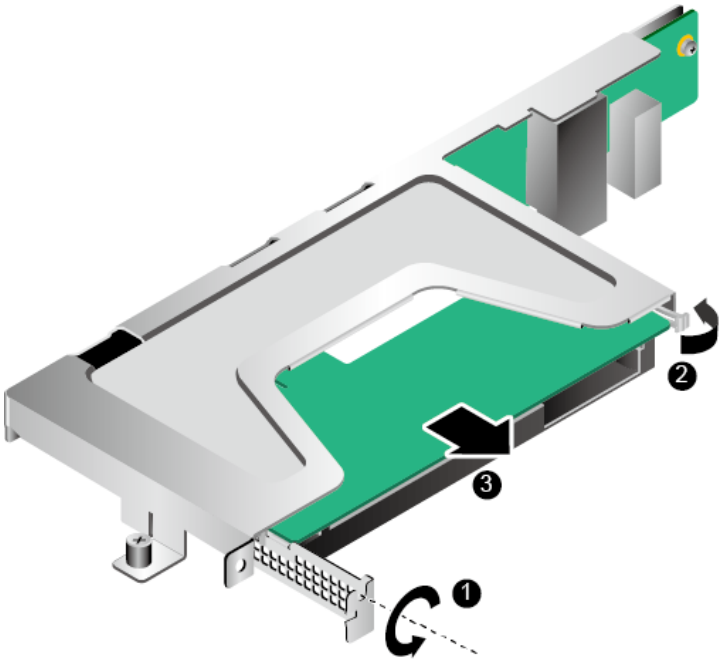

Step 4 Place the removed PCIe card in an ESD bag.

----End

5.12.4 Removing an ES3000 V5 NVMe PCIe SSD Card from the Mainboard

Prerequisites

Wear ESD gloves.

Procedure

Step 1 Turn off the power to the server and remove the chassis cover.

Step 2 Open the latch on the PCIe card.

Step 3 Hold the upper part of an ES3000 V5 NVMe PCIe SSD card and remove it upwards slowly. See **Figure 5-11**.

Figure 5-11 Removing an ES3000 V5 NVMe PCIe SSD card

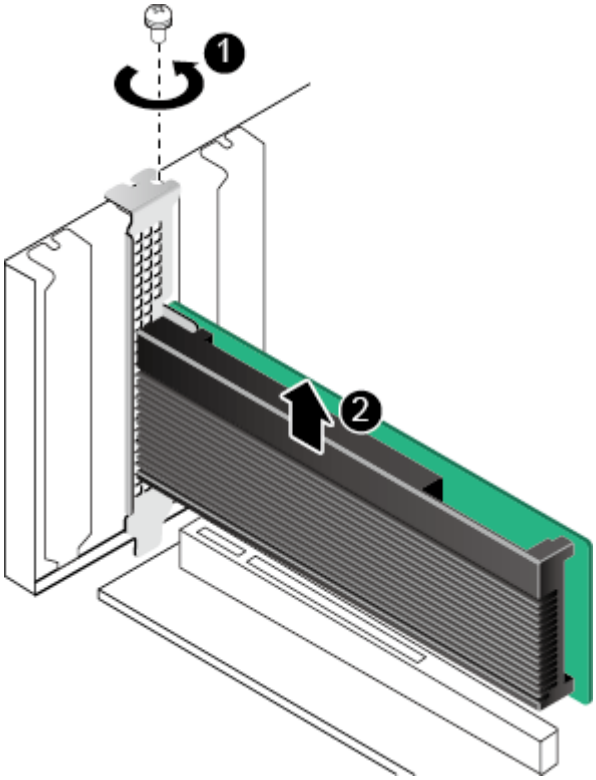

Step 4 Place the removed drive into an ESD bag.

----End

6 Command Description

About This Chapter

After installing the driver, you can use the **hioadm** tool to manage and configure the ES3000 V5. This chapter describes the **hioadm** command. The command output varies according to the actual operation result.

hioadm Tool

hioadm is an ES3000 management tool. You are advised to install this tool to facilitate environment maintenance and information collection. For details about the installation method, see **3.7 Installing the Tool Package3.7 Installing the Tool Package**. Its command format is as follows:

hioadm *<major cmd> <minor cmd1> <parameter1> <minor cmd2> <parameter2> ...*

<major cmd> can be any of the following:

- **info**: information query
- **format**: formatting
- **secure**: secure erasure of data
- **updatefw**: firmware upgrade management
- **capacity**: capacity management
- **log**: device log management
- **powerstate**: power consumption management
- **temperature**: temperature management
- **writemode**: write mode management (deprecated)
- **selftest**: diagnosis result query
- **label**: electronic label query
- **namespace**: namespace management
- **recovery**: data recovery
- **loadhist**: historical load query
- **synctime**: time synchronization management
- **restorefeatures**: factory settings recovery

- **sanitize**: device clea-up
- **hotplug**: orderly hot-swap
- **atomicwrite**: atomic write (This command is supported by ES3000 V5 2151 or later. For details about how to query the firmware version of an SSD, see **6.2.1 Querying the Firmware Version of an SSD**.)

NOTE

Target devices are SSDs (SSD controllers) and logical drives. The **-d** subcommand is used to specify the target device. For example, **nvme0** indicates a PCIe SSD, and **nvme0n1** indicates a logical drive. Currently only the **info** and **format** commands support both SSDs and logical drives. Other commands support only SSDs (if **nvme0n1** is entered, it is regarded as **nvme0**.)

Operation Logs

All setting operations performed on the ES3000 by using the hioadm tool are recorded. The log file location is as follows:

- For Linux, **/opt/hio/operationLogs/operations.log**.
- For Windows, **.\operationLogs\operations.log** in the installation directory.
- For ESXi, **/scratch/hio/operationLogs/operations.log**.

If the log file size exceeds 4 MB, the file is automatically backed up as an **operations.log.old** file, overwriting the existing **operations.log.old** file.

The log file content includes the command start time, remote client IP address, user name, command, command output time, and operation result.

The operation result is one of the following:

- **SUCCEEDED:** The command was successfully executed.
- **FAILED**: The command failed to be executed.
- **CANCELED**: The command was canceled because the user selected **No** when asked for conformation.
- **KILLED**: The command was stopped abnormally. For example, the user performed a **Ctrl+C** operation, closed the client, killed the process, or restarted the OS.

The following is an operation log in Linux:

```
[2016-01-27 17:37:10] [root@128.5.92.176] [hioadm updatefw -d nvme0 -s 3 -a 1]
[17:37:15] [SUCCEEDED].
```

6.1 Querying Information.

6.2 Querying, Upgrading, and Activating Firmware

6.3 Querying and Setting the Capacity

6.4 Formatting

6.5 Safely Erasing User Data of an SSD

6.6 Querying and Setting the Power Consumption Level

6.7 Obtaining Logs

6.8 Querying and Setting Temperature Thresholds

6.9 Querying and Setting the Write Mode

6.1 Querying Information.

6.1.1 Querying the SSD List

Function

This command is used to query the SSD list.

Format

hioadm info

Parameters

None

Usage Guidelines

None

Example

#Query the list of existing SSDs.

Linux:

```
[root@localhost tool]# hioadm info
NVMe SSD Information
      Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
      |------------ nvme0n1 (512.0GB)
      |------------ nvme0n2 (512.0GB)
|---- nvme1 (0503023HDCN107C80014)
      |------------ nvme1n1 (512.0GB)
SAS SSD Information
      Controller
```



```
|---- SSD (HS000000000000253015)
      |------------ SSD(512.0GB)
```

Windows:

```
C:\Program Files(x86)\hioadm> hioadm info
NVMe SSD Information
      Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
      |------------ nvme0n1[Disk3] (512.0GB)
      |------------ nvme0n2[Disk4] (512.0GB)
|---- nvme1 (0503023HDCN107C80014)
      |------------ nvme1n1[Disk5] (512.0GB)
SAS SSD Information
      Controller
|---- sda (HS000000000000253015)
      |------------ sda[Disk2] (512.0GB)
```

ESXi:

```
~# hioadm info
NVMe SSD Information
     Controller    Namespace
|---- nvme0 (0503023HDCN107C80013)
     |------------ nvme0n1[vmhba3:C0:T0:L0] (1.6TB)
SAS SSD Information
     Controller
```

System Response

The command output in the preceding example is explained as follows:

- Controller names: nvme0, nvme1, and sda
- Serial numbers: 0503023HDCN107C80013, 0503023HDCN107C80014 and HS000000000000253015
- Drive names: nvme0n1, nvme0n2, and sda
- Capacities: 512.0 GB

nvme0, **nvme0n1**, and **sda** in the command output are values of the parameter after **-d** in the commands.

6.1.2 Querying Basic Information About an SSD Device List

Function

This command is used to query basic information about an SSD device list.

Format

hioadm info -l

Parameters

None

Usage Guidelines

None

Example

#Query the list of existing SSDs.

```
[root@localhost tool]# hioadm info-l
Device      SN                       MN                    FW
nvme0       0503023HDCN107C80013     HWE32P43032M000N      1510
nvme1       0503023HDCN107C80014     HWE32P43016M000N      1510
sda         HS000000000000253015     HSSD-D5220AS0200      1510
```

Output Description

The command output in the preceding example is explained as follows:

- Controller names: nvme0, nvme1, and sda
- Serial numbers: 0503023HDCN107C80013, 0503023HDCN107C80014, and HS000000000000253015
- MN code: HWE32P43032M000N, HWE32P43016M000N, and HSSD-D5220AS0200
- Firmware version: 1510

6.1.3 Querying Basic Information About a Device

Function

This command is used to query basic information about an SSD (SSD controller) or logical drive. If you query information about an SSD, information about all of its logical drives is also displayed.

Format

hioadm info -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of a device	Examples: nvme0 and nvme1

Usage Guidelines

None

Example

Query the basic information about **nvme0**.

```
[root@localhost tool]# hioadm info -d nvme0

Namespace<1> size:   3200.6GB,    3200631791616Byte
             formatted LBA size:         512 Byte
             formatted metadata size:    0 Byte

maximum capacity                : 3200.6GB
```



```
current capacity                : 3200.6GB
volatile write cache            : Disable
serial number                   : 0503023HDCN107C80013
model number                    : HWE32P430016M00N
firmware version                : 1010
NVMe version                    : 1.3
device status                   : healthy
```

Query the basic information about **nvme0** when the size of formatted metadata is not 0.

```
[root@localhost tool]# hioadm info -d nvme0

Namespace<1> size:  30.0GB,    30000001024Byte
             formatted LBA size:         512 Byte
             formatted metadata size:    8 Byte
             protection type:        Type 3
             protection location:    the last eight bytes of metadata
             metadata capabilities: extended mode

maximum capacity              : 3840.7GB
current capacity              : 30.0GB
volatile write cache          : Disable
serial number                 : 024YVH10J3800180
model number                  :
firmware version              : 1510
NVMe version                  : 1.3
device status                 : healthy
```

System Response

Parameter	Description	Value
size	Available capacity of the namespace. Unit: GB In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Example: 3200.6 GB
formatted LBA size	Formatted logical block address (LBA) size of the namespace. Unit: byte	Example: 512 byte
formatted metadata size	Formatted metadata size of the namespace. Unit: byte	Example: 0 byte, 8 byte, 64 byte
protection type	Protection type. **NOTE** This parameter is displayed only when the value of **formatted metadata size** is not **0 byte**.	Example: Type 3
protection location	Location of protection information. **NOTE** This parameter is displayed only when the value of **formatted metadata size** is not **0 byte**.	Example: the last eight bytes of metadata
metadata capabilities	Metadata functions. **NOTE** This parameter is displayed only when the value of **formatted metadata size** is not **0 byte**.	Example: extended mode

Parameter	Description	Value
maximum capacity	Original maximum capacity of the controller. This value is included in basic port information, not in advanced port information.	Example: 3200.6 GB
current capacity	Current maximum capacity of the controller. This value is included in basic port information, not in advanced port information.	Example: 3200.6 GB
volatile write cache	Whether write cache is supported.	Example: Enable
serial number	Serial number.	Example: 0503023HDCN107C80013
model number	Model.	Example: HWE52P431T0L002N
firmware version	Firmware version.	Example: 1030
NVMe version	NVMe version.	Example: 1.3
device status	Device health status. Values are as follows: ● healthy ● warning ● unknown **NOTE** If an unknown status occurs, contact Huawei technical support for help. For details, see **H Getting Help**.	Example: healthy

6.1.4 Querying the SMART Information About a Device

Function

This command is used to query the SMART (Self-Monitoring, Analysis and Reporting Technology) information about an SSD (SSD controller).

Format

hioadm info -d *devicename* **-s**

Parameters

Parameter	Description	Value
devicename	Name of a device	Examples: nvme0 and nvme1

Usage Guidelines

None

Example

Query the SMART information about **nvme0**.

```
[root@localhost tool]# hioadm info -d nvme0 -s
critical warning              : no warning
composite temperature         : 308 degrees Kelvin (35 degrees Celsius)
available spare               : 100%
available spare threshold     : 10%
percentage used               : 0%
data units read               : 68.8 MB
data units written            : 0.0 MB
host read commands            : 17748
host write commands           : 0
controller busy time          : 0 mins
power cycles                  : 89 times
power on hours                : 1164 h
unsafe shutdowns              : 35 times
media and data integrity errors: 0
number of error information log entries: 0
warning composite temperature time: 0 min
critical composite temperature time: 0 min
data status                   : OK
```

System Response

Parameter	Parameter Description	Value
critical warning	Critical warning, such as overtemperature and insufficient redundant space.	Example: no warning There are the following types of critical warning: ● The available space is smaller than the threshold. ● The temperature exceeds the threshold. ● An internal error downgrades the NVM subsystem reliability. ● The medium is read-only. ● Device backup failed.
composite temperature	Temperature. Unit: degree Kelvin and Celsius	Example: 316 degrees Kelvin (43 degrees Celsius)

Parameter	Parameter Description	Value
available spare	Redundant space proportion.	Example: 100%
available spare threshold	Warning threshold of the redundant space proportion.	Example: 10%
percentage used	Percentage of used service life.	Example: 0%
data units read	Total read data volume. Unit: MB	Example: 68.8 MB
data units written	Total write data volume. Unit: MB	Example: 0.0 MB
host read commands	Number of read times.	Example: 17748
host write commands	Number of write times.	Example: 0
controller busy time	Busy period of the controller. Unit: minute	Example: 0 mins
power cycles	Number of power-ons.	Example: 89 times
power on hours	Power-on period.	Example: 1164 h
unsafe shutdowns	Number of unsafe shutdowns.	Example: 35 times
media and data integrity errors	Number of data errors.	Example: 0
number of error information log entries	Number of error log entries.	Example: 0
warning composite temperature time	Minor composite temperature alarm time.	Example: 0 mins
critical composite temperature time	Critical composite temperature alarm time.	Example: 0 mins

Parameter	Parameter Description	Value
data status	Data health status. Values are as follows: ● **OK**: The status is normal. ● **need recovery**: Data needs to be restored.	Example: OK

6.1.5 Querying Advanced Information of a Device

Function

This command is used to query the advanced information about an SSD (SSD controller).

Format

hioadm info -d *devicename* **-a**

Parameters

Parameter	Description	Value
devicename	Name of a device.	Examples: nvme0 and nvme1

Usage Guidelines

None

Example

Query the advanced information about **nvme0**.

```
[root@localhost ~]# hioadm info -d nvme0 -a
manufacturer                  : Huawei
WWN                           : 50022a1083846072
serial number                 : 025KCA10J7800101
model number                  : HWE56P436T4M002N
part number                   : 03032YRE
pcb version                   : A
controller id                 : 1
PCI slot                      : 0
highest temperature           : 324 degrees Kelvin (51 degrees Celsius)
exceed temperature threshold count: 0
exceed temperature threshold time: 0 mins
buck voltage                  : 9020 mV
capacitance                   : 1219 uF
capacitance voltage           : 30660 mV
```



```
capacitance status             : OK
capacitance charging time      : 12 times
active power                   : 5266 mW
maximum power                  : 5403 mW
number of power states support: 3
atomic write                   : OFF
bad block rate                 : 1.926386%
total blocks count             : 258048
bad blocks count               : 4971
factory bad blocks count       : 4971
flash program failed count     : 0
flash erase failed count       : 0
maximum data transfer size     : 32 pages
```

System Response

Parameter	Parameter Description	Value
manufacturer	Manufacturer.	Huawei
WWN	Unique identifier.	Example: 50022A1012345678
serial number	Serial number.	Example: 023HDF10FB800303
part number	Board number.	Example: 03023LSX
model number	Model	Example: HWE32P430016M00N
pcb version	PCB version.	Example: A
controller id	Controller ID.	Example: 1
PCI slot	PCI slot number	Example: 6 **NOTE** The value may be **--**, which indicates that the physical slot in lspci is empty.
highest temperature	Highest temperature recorded.	Example: 329 degree K (56 degree C)
exceed temperature threshold count	Number of times that the temperature threshold is exceeded.	Example: 0
exceed temperature threshold time	Period during which the temperature threshold is exceeded.	Example: 0 mins
buck voltage	Buck voltage.	Example: 10603 mV
capacitance	Capacitor capacitance.	Example: 1219 uF
capacitance voltage	Capacitor voltage.	Example: 31416 mV

Parameter	Parameter Description	Value
capacitance status	Capacitor health status. Values are as follows: ● **OK**: The capacitor is normal. ● **Invalid**: The capacitor is invalid.	Example: OK
capacitance charging time	Number of discharge/charge cycles of the capacitor.	Example: 3 times
active power	Current total power consumption.	Example: 7041 mW
maximum power	Highest power consumption recorded.	Example: 7041 mW
number of power states support	Number of power states supported.	Example: 3
bad block rate	Bad block rate.	Example: 1.776449%
total blocks count	Total number of blocks, including bad blocks.	Example: 134144
bad blocks count	Total number of bad blocks.	Example: 2383
factory bad blocks count	Number of bad blocks before delivery.	Example: 0
flash program failed count	Number of flash program failures.	Example: 0
flash erase failed count	Number of flash erasure failures.	Example: 0
maximum data transfer size	Maximum data transfer size.	Example: unlimit

6.2 Querying, Upgrading, and Activating Firmware

6.2.1 Querying the Firmware Version of an SSD

Function

This command is used to query the firmware version of an SSD.

Format

hioadm updatefw -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of an SSD	Example: nvme0 or nvme1

Usage Guidelines

None

Example

Query the firmware version of **nvme0**.

```
[root@localhost tool]# hioadm updatefw -d nvme0
slot  version activation
1     1010    next
2     1011    current
```

NOTE

A SAS SSD has only one slot, and an NVMe SSD has multiple slots. The version number to the left of **next** indicates the version that takes effect after the next reset. The version number to the left of **current** indicates the current version in use. If the **next** line is not displayed, the version does not change after the next reset.

6.2.2 Upgrading the Firmware of an SSD

Function

This command is used to upgrade the firmware of an SSD.

CAUTION

If you run the following command with **activeflag** set to **1** (that is, firmware will be activated immediately), services may be stopped. Before running this command, ensure that stopping services has no adverse impact.

Format

hioadm updatefw -d *devicename* **-f** *fwimagefile* [**-s** *slot*] [**-a** *activeflag*]

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**
fwimagefile	Path of the target firmware image file.	Example: **/home/fw_image.img**

Parameter	Description	Value
slot	Slot of the target firmware.	1 or 2 **NOTE** ● If this parameter is not specified, a slot with firmware not in use will be selected. ● If firmware in slot 1 is in use, slot 2 will be selected by default. ● A SAS SSD has only one slot, namely, slot 1.
activeflag	Firmware activation mode.	● **0**: Downloaded firmware is activated after the next reset. ● **1**: Firmware is activated immediately after being downloaded. **NOTE** If this parameter is not set, the downloaded firmware is activated after the next reset.

Usage Guidelines

None

Example

Download the **nvme0** firmware and make it activated after the next reset.

```
[root@localhost tool]# hioadm updatefw -d nvme0 -f ES3000_FW_2.08.bin -s 2 -a 0
slot  version activation
1     1010    next
2     1011    current



WARNING! You have selected slot <2> to update the firmware image.
WARNING! You have selected the update firmware operation.
Proceed with the update firmware operation? (Y|N): y
Please do not remove driver or SSD device
Loading... OK
Downloading and replacing the firmware image at slot <2> succeeded.
Activating the firmware image at slot <2> succeeded.
```

6.2.3 Activating the Firmware of an SSD

Function

This command is used to activate the firmware of an SSD.

⚠ CAUTION

If you run the following command with **activeflag** set to **1** (that is, firmware will be activated immediately), services may be interrupted. Before running this command, ensure that stopping services has no adverse impact.

Format

hioadm updatefw -d *devicename* [**-s** *slot*] [**-a** *activeflag*]

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0** and **sda**
slot	Slot of the target firmware image.	1 and 2
activeflag	Firmware activation mode.	● **0**: Activates firmware after the next reset. ● **1**: Activates firmware immediately.

Usage Guidelines

None

Example

Activate the **nvme0** firmware after the next reset.

```
[root@localhost tool]# hioadm updatefw -d nvme0 -s 1 -a 0
slot  version activation
1     1010    next
2     1011    current


WARNING! Activate immediately might abort IO.
WARNING! You have selected the update firmware operation.
Proceed with the update firmware? (Y|N): y
Activating the firmware image at slot <1> succeeded.
```

6.3 Querying and Setting the Capacity

The capacity size affects the write performance. A small capacity results in a high random write IOPS. The ES3000 allows you to set the capacity within a specified range.

NOTE

Setting the capacity may cause data loss. Save your data before this operation.

6.3.1 Querying the Capacity of an SSD

Function

This command is used to query the current capacity of an SSD.

Format

hioadm capacity -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**

Usage Guidelines

None

Example

Query the current capacity of **nvme0**.

```
[root@localhost tool]# hioadm capacity -d nvme0
current capacity : 1600.3GB (1600321314816Byte)
maximum capacity : 1600.3GB (1600321314816Byte)
```

System Response

Parameter	Parameter Description	Value
current capacity	Total available capacity (with one decimal place) of the SSD. Unit: GB. In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Example: 1600.3 GB (1600321314816 Byte)
maximum capacity	Original available capacity (with one decimal place) of the SSD. Unit: GB. In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Example: 1600.3 GB (1600321314816 Byte)

6.3.2 Setting the Capacity of an SSD

Function

This command is used to set the current capacity of an SSD.

Format

hioadm capacity -d *devicename* **-s** *capacitysize*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**
capacitysize	Target capacity. In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Integer or decimal with a unit. ● Supported units: G. ● The minimum capacity is 1 GB. The maximum capacity is the raw capacity of the disk. ● If the unit is GB, only integers are supported. **0** indicates that the target capacity is the original capacity.

Usage Guidelines

- Before running this command, stop all I/O services of the SSD, and unmount all logical drives of the SSD.
- If a Huawei-developed driver is installed, the setting takes effect immediately. If an open-source or system-provided driver is installed, the setting takes effect after the driver is uninstalled and reinstalled or the server is restarted. After the setting takes effect, the namespace settings automatically restore to the defaults, that is, only one namespace of the configured size exists and is attached to controller 1.

Example

Set the current capacity of **nvme0** to its original capacity.

```
[root@localhost tool]# hioadm capacity -d nvme0 -s 0
current capacity : 512.0GB (512000000000Byte)
maximum capacity : 1600.3GB (1600321314816Byte)
You have selected to restore original capacity.

WARNING! The set capacity operation will clear all data.
WARNING! You have selected the set capacity operation.
Proceed with the set capacity? (Y|N): y
Setting...OK
Setting capacity succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

System Response

Parameter	Parameter Description	Value
current capacity	Total available capacity (with one decimal place) of the SSD. Unit: GB or TB. In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Example: 512.0 GB (512000000000 Byte)

Parameter	Parameter Description	Value
maximum capacity	Original available capacity (with one decimal place) of the SSD. Unit: GB or TB. In Linux and ESXi, 1 GB equals 1000 MB. In Windows, 1 GB equals 1024 MB.	Example: 1600.3 GB (1600321314816Byte)

6.4 Formatting

Formatting is classified into quick formatting and secure formatting.

- Quick formatting (Secure Erase Settings=0): Erase only entries and support namespace formatting.
- Secure formatting (Secure Erase Settings=1): Erase all data and support only entire drive formatting.

CAUTION

Perform formatting with caution. Data cannot be restored after formatting.

If drivers do not support the Data Integrity Field (DIF) function, do not format drives into the DIF format. Otherwise, problems such as operating system resetting will occur.

6.4.1 Formatting a Logical Drive

Function

Perform quick formatting of a logical drive. Logical drives support only quick formatting. Secure formatting is not allowed.

Format

hioadm format -d *diskname*

Parameters

Parameter	Description	Value
diskname	Name of a logical drive.	Example: **nvme0n1**

Usage Guidelines

- Before running this command, stop all I/O services of the logical drive and perform the following:
  - For Linux, unmount the logical drive. Stop all the processes on the logical drive, and run the **fuser** command to check that these processes are all stopped. The

following information indicates that a process is using the logical drive. Stop the process.

```
[root@localhost tool]# fuser /dev/nvme0n1
/dev/nvme0n1:        14789
```

- For ESXi, delete all partitions of the logical drive.
- For Windows, make the logical drive offline, as shown in **Figure 6-1**.

Figure 6-1 Making a logical drive offline

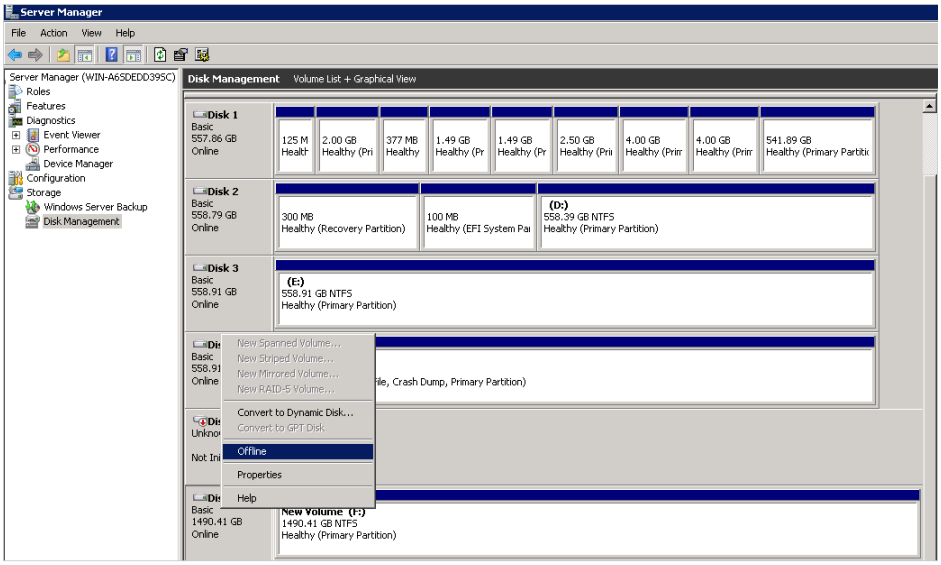

CAUTION

Do not read or write data when running this command. Otherwise, I/O errors may occur.

Example

Perform quick formatting of the logical drive **nvme0n1**.

```
[root@localhost tool]# hioadm format -d nvme0n1

WARNING! The format operation will clear all data.
WARNING! You have selected the format operation.
Proceed with the format? (Y|N): y
Formatting.... OK
The format operate succeeded.
The current LBA Format type : <0> lba=512  metadata=0
```

6.4.2 Querying Formatting Parameter Support Information

Function

Query formatting parameter support information, including LBA size and metadata size combinations, DIF verification information enablement status, type, DIF verification location, and metadata modes.

Format

hioadm format -d *devicename* **-q**

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**

Output Description

Parameter	Description	Value
LBA Format	● This parameter specifies supported combinations of LBA sizes and metadata sizes, and relative performance. ● **Data Size**: indicates the size of sector. Example: **512 bytes** ● **Metadata Size**: indicates the metadata size. Example: **0 bytes** ● **Relative Performance**: indicates relative performance. Example: **Best** ● **(in use)**: indicates the formatting combination currently in use.	Example: **LBA Format 0: Data Size: 512 bytes - Metadata Size: 0 bytes - Relative Performance: Best (in use)** This example indicates that the device supports the combination of 512 B sectors, 0 B metadata, and best relative performance, and that the combination is currently in use.

Parameter	Description	Value
Protection Information	● This parameter specifies whether DIF verification information is enabled, and supported information types. ● **Type 0**: DIF verification information is disabled and the information type is type 0. **Type 1**: DIF verification information is enabled and the information type is type 1. **Type 2**: DIF verification information is enabled and the information type is type 2. **Type 3**: DIF verification information is enabled and the information type is type 3. ● **Supported**: indicates that the current DIF verification information type is supported. **Unsupported**: indicates that the current DIF verification information type is not supported. ● **(in use)**: indicates the DIF verification information type currently in use.	Example: **Protection information is not enabled, Type 0: Supported (in use)** This example indicates that DIF verification information is disabled and the information type is type 0, and that the setting is currently in use.
Protection Information Location	● This parameter specifies supported DIF verification information locations. ● **Protection information transferred as the last eight bytes of metadata**: DIF verification information is transferred as the first eight bytes of metadata. **Protection information transferred as the first eight bytes of metadata**: DIF verification information is transferred as the last eight bytes of metadata. ● **Supported**: indicates that the current DIF verification location is supported. **Unsupported**: indicates that the current DIF verification information location is not supported. ● **(in use)**: indicates the DIF verification information location currently in use.	Example: **Protection information transferred as the last eight bytes of metadata: Supported (in use)** This example indicates that DIF verification information is transferred as the last eight bytes of metadata, and that the setting is currently in use.

Parameter	Description	Value
Metadata Capabilities	● This parameter specifies supported metadata modes. ● **Metadata being transferred as part of a separate buffer**: Metadata is transferred as part of a separate buffer (Separate Mode). **Metadata being transferred as part of an extended data LBA**: Metadata is transferred as part of an extended LBA (Extend Mode). ● **Supported**: indicates that the current metadata mode is supported. **Unsupported**: indicates that the current metadata mode is not supported. ● **(in use)**: indicates the metadata mode currently in use.	Example: **Metadata being transferred as part of an extended data LBA: Supported (in use)** This example indicates that metadata is transferred as part of an extended LBA (Extend Mode), and that the setting is currently in use.

Example

#Query **nvme0** formatting parameter support information.

```
[root@localhost tool]# hioadm format -d nvme0 -q

LBA Format:
    LBA Format 0: Data Size: 512  bytes - Metadata Size: 0  bytes - Relative
Performance: Best (in use)
    LBA Format 1: Data Size: 512  bytes - Metadata Size: 8  bytes - Relative
Performance: Better
    LBA Format 2: Data Size: 4096 bytes - Metadata Size: 64 bytes - Relative
Performance: Better
    LBA Format 3: Data Size: 4096 bytes - Metadata Size: 0  bytes - Relative
Performance: Best
    LBA Format 4: Data Size: 4096 bytes - Metadata Size: 8  bytes - Relative
Performance: Better

Protection Information:
    Protection information is not enabled, Type 0: Supported (in use)
    Protection information is enabled, Type 1: Reserved
    Protection information is enabled, Type 2: Reserved
    Protection information is enabled, Type 3: Supported

Protection Information Location:
    Protection information transferred as the last eight bytes of metadata:
Supported (in use)
    Protection information transferred as the first eight bytes of metadata:
Supported

Metadata Capabilities:
    Metadata being transferred as part of a separate buffer: Supported (in use)
    Metadata being transferred as part of an extended data LBA: Supported
```


6.4.3 Formatting an SSD

Function

This command is used to format an SSD. Secure formatting and quick formatting are supported.

Format

hioadm format -d *devicename* [**-m** *formatmode*] [**-t** *lbasize*] [**-pt** *protectiontype*] [**-pil** *location*] [**-mc** *metadatamode*]

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**
formatmode	Formatting mode.	● **0**: quick formatting (Secure Erase Settings=0) ● **1**: secure formatting (Secure Erase Settings=1) If this parameter is not specified, the default value **0** is used.
lbasize	Sector size and type.	● **0**: The sector size is 512 B, and the metadata size is 0 B. ● **1**: The sector size is 4 KB, and the metadata size is 0 B. ● **2**: The sector size is 512 B, and the metadata size is 8 B. ● **3**: The sector size is 4 KB, and the metadata size is 64 B. ● **4**: The sector size is 4 KB, and the metadata size is 8 B. ● **5**: The sector size is 520 B, and the metadata size is 0 B. ● **6**: The sector size is 4160 B, and the metadata size is 0 B. If this parameter is not specified, the sector size and type are not changed. **NOTE** The Linux operating system supports all the preceding sector types (5 and 6 are applicable only to SAS drives). The Windows operating system supports sector types 0, 1, 5, and 6 (5 and 6 are applicable only to SAS drives). The ESXi system supports only the 512 B sector type. If this parameter is not set, the sector type is not changed.

Parameter	Description	Value
protectiontype	Protection enabling and type.	● 0: indicates that the protection is not enabled. ● 1: indicates that the protection is enabled and the protection type is 1 (not supported yet). ● 2: indicates that the protection is enabled and the protection type is 2. ● 3: indicates that the protection is enabled and the protection type is 3.
location	Protection data location.	● If this value is set to **1** and the protection is enabled, protection data is transmitted as the first eight bytes of the metadata. ● If this value is set to **0** and the protection is enabled, protection data is transmitted as the latter eight bytes of the metadata.
metadatamode	Protection data mode.	● **1**: Metadata is transferred as a part of the extended data LBA. ● **0**: Metadata is transferred as a part of separate cache. **NOTE** If **Metadata Size** is set to **0**, this parameter is unavailable.

Usage Guidelines

- Before running this command, stop all I/O services of the SSD and perform the following:
  - For Linux, unmount all logical drives. Stop all the processes on the logical drives that use the SSD, and run the **fuser** command to check that these processes are all stopped. The following information (in this example, the SSD has two logical drives) indicates that a process is using an SSD logical drive. Stop the process.

    ```
    [root@localhost tool]# fuser /dev/nvme0n1 /dev/nvme0n2
    /dev/nvme0n1:        14789
    /dev/nvme0n1:        14812
    ```

  - For ESXi, delete all partitions of all logical drives.
  - For Windows, make all logical drives offline, as shown in **Figure 6-2**.

Figure 6-2 Making a logical drive offline

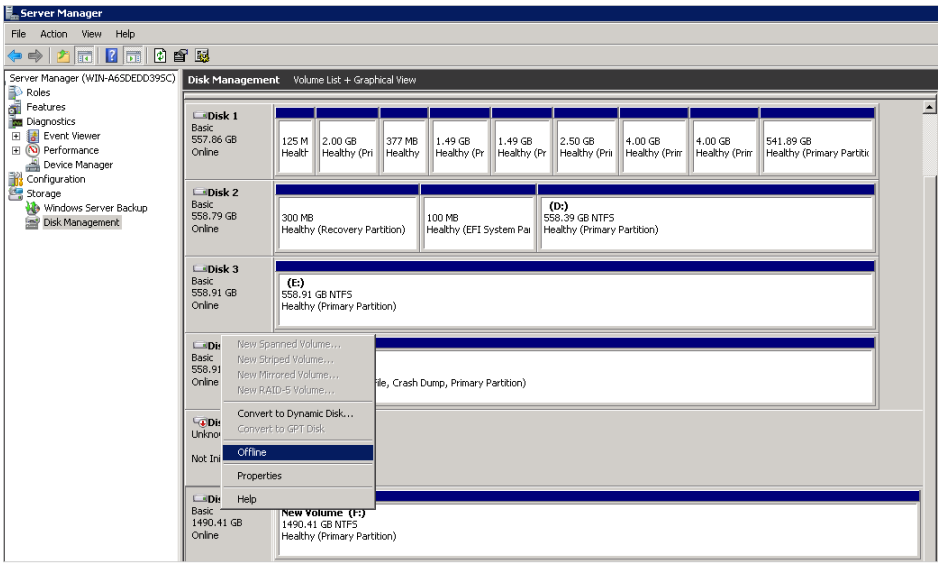

CAUTION

Do not read or write data when running this command. Otherwise, I/O errors may occur.

Example

Perform secure formatting of the **nvme0** SSD, and set the sector size to 512 B and metadata size to 0 B.

```
[root@localhost tool]# hioadm format -d nvme0 -m 1 -t 0

WARNING! The format operation will clear all data.
WARNING! You have selected the format operation.
Proceed with the format? (Y|N): y
Formatting.... OK
The format operate succeeded.
Please remove and then install nvme driver or reboot the OS.
The current LBA Format type : <0> lba=512  metadata=0
```

6.5 Safely Erasing User Data of an SSD

Function

This command is used to safely erase user data of an SSD.

After safely erasing an SSD, the sector type and quantity of logical drives are not changed.

CAUTION

Exercise caution because this operation erases data.

Format

hioadm secure -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**

Usage Guidelines

- This command supports NVMe SSDs provided by other vendors.
- Before running this command, stop all I/O services of the SSD and perform the following:
  - For Linux, unmount all logical drives. Stop all the processes on the logical drives that use the SSD, and run the **fuser** command to check that these processes are all stopped. The following information (in this example, the SSD has two logical drives) indicates that a process is using an SSD logical drive. Stop the process.

    ```
    [root@localhost tool]# fuser /dev/nvme0n1 /dev/nvme0n2
    /dev/nvme0n1:        14789
    /dev/nvme0n1:        14812
    ```

  - For ESXi, delete all partitions of all logical drives.
  - For Windows, make the logical drives offline, as shown in **Figure 6-3**.

    Figure 6-3 Making a logical drive offline

    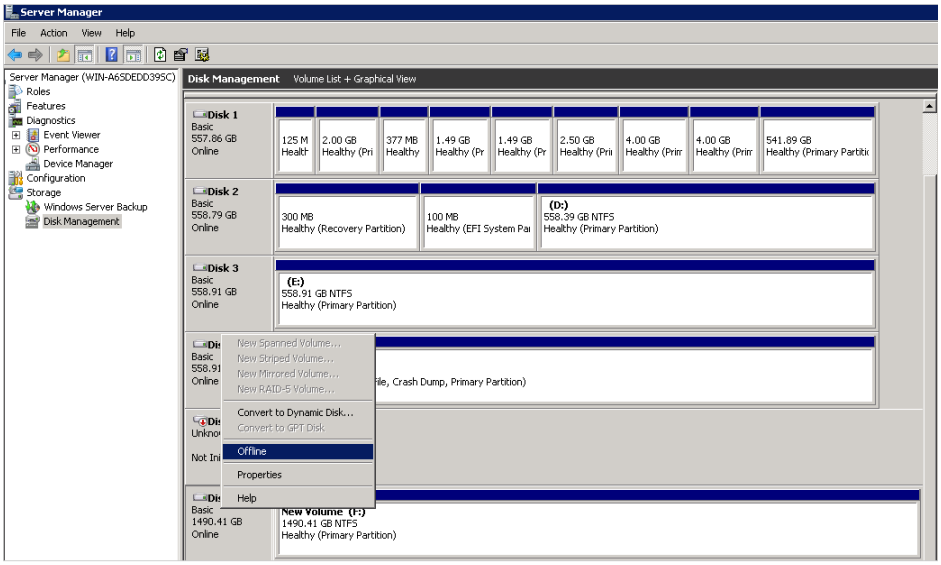

Example

Safely erase user data of the **nvme0** SSD.

```
[root@localhost tool]# hioadm secure -d nvme0

WARNING! The secure operation will clear all data.
```



```
WARNING! You have selected the secure operation.
Proceed with the secure? (Y|N): y
Format Progress.... 100%
The secure operate succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

6.6 Querying and Setting the Power Consumption Level

6.6.1 Querying the Power Consumption Level of an SSD

Function

This command is used to query the power consumption level of an SSD.

Format

hioadm powerstate -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**

Usage Guidelines

None

Example

Query the power consumption level of the **nvme0** SSD.

```
[root@localhost ~]# hioadm powerstate -d nvme0
Power states:
  [*] <0> 25W
  [ ] <1> 18W
  [ ] <2> 14W
```

6.6.2 Setting the Power Consumption Level of an SSD

Function

This command is used to set the power consumption level of an SSD.

Format

hioadm powerstate -d *devicename* **-s** *powerstate*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**
powerstate	Power consumption level.	● **0**: 25 W ● **1**: 18 W ● **2**: 14 W

NOTE

- The power consumption level affects device performance. The lower the power consumption, the poorer the performance.
- When the power consumption level is 0, the performance is not limited.
- The default power consumption level 0 ensures optimal performance. A non-default power consumption level is not recommended.

Usage Guidelines

None

Example

Set the **nvme0** SSD power consumption level to unlimited.

```
[root@localhost tool]# hioadm powerstate -d nvme0 -s 0 Set power state succeeded.
Power states:
[*] <0> 25W
[ ] <1> 18W
[ ] <2> 14W
```

6.7 Obtaining Logs

6.7.1 Obtaining Logs of a Device

Function

This command is used to obtain logs of an SSD and generate a log file.

Format

hioadm log -d *devicename* **-t** *logtype*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Parameter	Description	Value
logtype	Log type.	● **0**: firmware key log ● **1**: firmware log ● **2**: core dump log The default value is **0**.

Usage Guidelines

The path and name of the generated log file is as follows. **nvme0** is an example device name.

- Linux: The file path is **/opt/hio/**; the key log file name is **nvme0_key.log**; the full log file name is **nvme0_fw.log**; the core dump file name is **nvme0_coredump.log**.
- ESXi: The file path is **/scratch/hio/**; the key log file name is **nvme0_key.log**; the full log file name is **nvme0_fw.log**; the core dump file name is **nvme0_coredump.log**.
- Windows: The file path is the directory storing **hioadm.exe**; the key log file name is **nvme0_key**; the full log file name is **nvme0_fw**; the core dump file name is **nvme0_coredump**.

Example

Obtain the key logs of the **nvme0** SSD.

```
[root@localhost tool]# hioadm log -d nvme0 -t 0
firmware key log saving..........OK
save firmware key log file succeeded(/opt/hio/nvme0_key.log)
logs has been saved.
```

6.7.2 Obtaining All Logs of a Device At a Time

Function

Obtain and export all logs of a device to a predefined directory. Exported logs include key logs, full logs, core dump logs, and other logs of firmware.

Format

hioadm log -d *devicename* **-a**

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Usage Guidelines

- Linux: The log file path is **/opt/hio/**.

- ESXi: The log file path is **/scratch/hio/**.
- Windows: The log file path is the directory storing **hioadm.exe**.

Example

In Linux, obtain all logs of the **nvme0** SSD at a time.

```
[root@localhost tool]# hioadm log -d nvme0 -a
firmware key log saving..........OK
save firmware key log file succeeded(/opt/hio/nvme0_key.log)
firmware all log saving..........OK
save firmware all log file succeeded(/opt/hio/nvme0_fw.log)
firmware core dump log saving..........OK
save firmware core dump log file succeeded(/opt/hio/nvme0_coredump.log)
other log saving..........OK
save other logs file succeeded(/opt/hio/nvme0_other.log)
logs has been saved.
```

6.8 Querying and Setting Temperature Thresholds

6.8.1 Querying the Temperature Information About a Device

Function

This command is used to query the temperature and temperature thresholds of an SSD.

Format

hioadm temperature -d *devicename* [**-i** *index*]

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
index	Number of a temperature sensor.	● **0**: overall temperature ● **1**: controller temperature If this parameter is not specified, all sensor information is queried.

Usage Guidelines

None

Example

Query the temperature information about the **nvme0** SSD.

```
[root@localhost tool]# hioadm temperature -d nvme0
composite highest temperature:               335 degree Kelvin (62 degree
```



```
Celsius)
composite overtemperature threshold count: 0
composite overtemperature threshold time:  0 min
controller highest temperature:     341 degree Kelvin (68 degree Celsius)
controller overtemperature threshold count:  0
controller overtemperature threshold time: 0 min

sensors temperature information in degree Kelvin (degree Celsius):
              current     overtemperature threshold     undertemperature
threshold
composite    317(44)            351(78)                    273(0)
controller1  328(55)          373(100)                   273(0)
```

System Response

Parameter	Parameter Description	Value
composite highest temperature	Highest SSD temperature recorded. Unit: K and °C	Example: 363 degree K (90 degree C)
composite overtemperature threshold count	Number of SSD overtemperature occurrences.	Example: 5
composite overtemperature threshold time	Period of SSD overtemperature. Unit: minute	Example: 500 min
controller highest temperature	Highest controller temperature recorded. Unit: K and °C	Example: 363 degree K (90 degree C)
controller overtemperature threshold count	Number of controller overtemperature occurrences.	Example: 5
controller overtemperature threshold time	Period of controller overtemperature. Unit: minute	Example: 500 min
current	Current temperature. Unit: K and °C	Example: 303(30)
overtemperature threshold	Overtemperature threshold. Unit: K and °C	Example: 341(68)
undertemperature threshold	Undertemperature threshold. Unit: K and °C	Example: 273(0)

6.8.2 Setting the Temperature Thresholds of a Device

Function

This command is used to set the temperature thresholds of an SSD or controller.

Format

hioadm temperature -d *devicename* [**-i** *index*] *-o overtemperature_threshold -u undertemperature_threshold*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
index	Number of a temperature sensor.	● **0**: overall temperature ● **1** to **8**: controller sensor number **NOTE** If this parameter is not specified, all sensors are configured. SAS drives can be configured with the high temperature threshold.
overthreshold	Overtemperature threshold.	● Unit: K ● Value range: – Overall temperature: 273K to 351K (0°C to 78°C) – Controller temperature: 273K to 373K (0°C to 100°C) If this parameter is not specified, this parameter value is not changed. Either the overtemperature or undertemperature threshold must be specified. **NOTE** ES3000 SSDs have two alarm severities. The second alarm severity restricts drive I/O speed. If an overtemperature alarm occurs, check server heat dissipation and ambient temperature.
underthreshold	Undertemperature threshold.	● Unit: K ● Value range: – Overall temperature: 273K to 351K (0°C to 78°C) – Controller temperature: 273K to 373K (0°C to 100°C) If this parameter is not specified, this parameter value is not changed. Either the overtemperature or undertemperature threshold must be specified.

Usage Guidelines

None

Example

Set the overtemperature and undertemperature thresholds to **333K** and **273K** for the **nvme0** SSD.

```
[root@localhost tool]# hioadm temperature -d nvme0 -i 0 -o 333 -u 273
set overtemperature threshold succeeded.
set undertemperature threshold succeeded.
[root@localhost tool]# hioadm temperature -d nvme0
composite highest temperature:                335 degree Kelvin (62 degree
Celsius)
composite overtemperature count: 0
composite overtemperature time:  0 min
controller highest temperature:     341 degree Kelvin (68 degree Celsius)
controller overtemperature count:  0
controller overtemperature time: 0 min

sensors temperature information in degree Kelvin (degree Celsius):
               current     over threshold     under threshold
composite    303(30)       333 (60)            273(0)
controller1  303(30)       358(85)              273(0)
```

6.9 Querying and Setting the Write Mode

6.9.1 Querying the Write Mode of an SSD

Function

This command is used to query the write mode of an SSD.

Format

hioadm writemode -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Usage Guidelines

None

Example

Query the write mode of the **nvme0** SSD.

```
[root@localhost tool]# hioadm writemode -d nvme0
write mode: write through
```

6.9.2 Setting the Write Mode of an SSD

The ES3000 V5 currently does not support this operation.

6.10 Querying the Diagnosis Result of a Specified SSD

Function

Query the diagnosis result of a specified SSD.

Format

hioadm selftest -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Usage Guidelines

None

Example

Conduct diagnosis tests for the **nvme0** SSD.

```
[root@localhost host]# hioadm selftest -d nvme0
selftest result:
firmware abnormal       passed
NAND flash                passed
NOR flash                 passed
DDR UNC                    failed
ibuf UNC                   passed
capacitor                 failed
blank block               passed
life                        passed
NOR flash patrol         passed
firmware UNC              passed
system area               passed
bad block                  passed
controller                 passed
powerdown fiq              passed
format fail                passed
data check                 passed
```


System Response

Parameter	Parameter Description	Value
selftest result	Overall test result. The value is **passed** or empty. If the result is success, all the following sub-tests succeed and sub-test results are not displayed. If the result is failure, details of the following sub-tests are displayed.	Example: passed
firmware abnormal	Firmware error status test result. Values are **passed** and **failed**.	Example: passed
NAND flash	NAND flash test result. Values are **passed** and **failed**.	Example: passed
NOR flash	NOR flash test result. Values are **passed** and **failed**.	Example: passed
DDR UNC	DDR UNC test result. Values are **passed** and **failed**.	Example: failed
ibuf UNC	ibuf UNC test result. Values are **passed** and **failed**.	Example: passed
capacitor	Capacitor test result. Values are **passed** and **failed**.	Example: failed
blank block	Blank block test result. Values are **passed** and **failed**.	Example: passed
life	Life test result. Values are **passed** and **failed**.	Example: passed
NOR flash patrol	NOR patrol test result. Values are **passed** and **failed**.	Example: passed
firmware UNC	Firmware UNC test result. Values are **passed** and **failed**.	Example: passed
system area	System area test result. Values are **passed** and **failed**.	Example: passed
bad block	Bad block test result. Values are **passed** and **failed**.	Example: passed
controller	Controller validity test result. Values are **passed** and **failed**.	Example: passed
powerdown fiq	Power-off detection circuit test result. Values are **passed** and **failed**.	Example: passed
format fail	Format failure test result. Values are **passed** and **failed**.	Example: passed

Parameter	Parameter Description	Value
data check	Data check test result. Values are **passed** and **failed**.	Example: passed

6.11 Querying the Electronic Label Information About an SSD

Function

This command is used to query the electronic label information about an SSD.

Format

hioadm label -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Usage Guidelines

None

Example

Query the electronic label information about the **nvme1** SSD.

```
[root@localhost ~]# hioadm label -d nvme1
BoardType=HICASPWBD
BarCode=032YRH10J9000104
Item=03032YRH
Description=Finished Board,ES3000 V5,HICASPWBD,ES3600C-800GB-3 DWPD-PCIE 3.0 X4-
Vendor ID 19e5-Device ID 3714-1,Model number HWE56P43800M002N,HH/HL Card,NVMe SSD
Manufactured=2018-09-30
VendorName=Huawei
IssueNumber=00
CLEICode=
BOM=
Model=
/$ElabelVersion=4.0
```

6.12 Recovering Data

6.12.1 Recovering Data for an SSD

Function

This command is used to recover data for an SSD.

Format

hioadm recovery -d *devicename* [**-t** *type*]

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
type	Data recovery mode.	● **0**: common recovery ● **1**: forcible recovery **NOTE** If this parameter is not set, the common recovery mode is used by default.

Usage Guidelines

- This command may take up to 20 minutes.
- Before recovering data, run the **hioadm info -d** *devicename* **-s** command. Recover data only when **data status** in the command output is **need recovery**.

NOTE

- The status indicates that the device has critical non-hardware issues. Data recovery may not be able to recover all data.
- Common recovery will recover original user data as much as possible. To ensure data reliability, SSD disks or cards will work in read-only mode after common recovery. Forcible recovery will erase all user data. If forcible recovery is successful, SSD disks will enter the healthy state.

Example

Recover data for the **nvme0** SSD.

```
[root@localhost host]# hioadm recovery -d nvme0
you have selected normal recovery.

WARNING! The recovery operation will take tens of minutes.Please wait.
WARNING! You have selected the recovery operation.
Proceed with the recovery? (Y|N): y
Recovering succeeded.
.
```

6.12.2 Querying the SSD Data Common Recovery Progress

Function

This command is used to query the SSD data common recovery progress.

Format

hioadm recovery -d *devicename* -q

Parameter

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Example

Query the data recovery progress of the **nvme0** SSD.

```
[root@localhost host]# hioadm recovery -d nvme0 -q Recovering progress…100%
The recovery operation completed successfully.
```

6.13 Querying and Configuring Namespaces

For a device that supports multiple namespaces, the **namespace** command is used to manage namespaces.

For a device that supports only one namespace, the "The number of namespaces supported has been exceeded" message is displayed when you create a namespace.

6.13.1 Querying the Namespace Distribution of an SSD

Function

This command is used to query the namespace distribution of an SSD.

Format

hioadm namespace -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of a device	Example: **nvme0**

Usage Guidelines

None

Example

Query the namespace distribution of the **nvme0** SSD.

```
[root@localhost ~]# hioadm namespace -d nvme0
```



```
total capacity: 3840.7GB
unallocated capacity: 3740.7GB
maximum capacity to creat single namespace: 3740.7GB

namespace id  capacity  shared  controller id
1             100.0GB   NO      1
```

System Response

Parameter	Parameter Description	Value
total capacity	Total SSD capacity with one decimal place. Unit: GB.	Example: 3200.6 GB
unallocated capacity	Unallocated SSD capacity with one decimal place. Unit: GB.	Example: 3200.6 GB
maximum capacity to creat single namespace	Maximum capacity (with one decimal place) of a single namespace. Unit: GB.	Example: 800.0 GB
namespace id	All namespace IDs of an SSD.	Example: 1
capacity	Capacity (one decimal place) of a namespace. Unit: GB.	Example: 200.0.0 GB
shared	Whether the namespace is shared. The value is **no**.	no
controller id	IDs of controllers attached to a namespace.	Example: 1

6.13.2 Creating a Namespace for an SSD

Function

This command is used to create a namespace for an SSD.

Format

hioadm namespace -d *devicename* **-t** *operationtype* **-s** *size* **-f** *shareflag*

Parameters

Parameter	Description	Value
devicename	Name of a device	Example: **nvme0**
operationtype	Operation type	● **0**: Create a namespace. ● **1**: Delete a namespace. ● **2**: Attach a namespace. ● **3**: Detach a namespace. ● **6**: Set the number of namespaces. Select **0** in this command.

Parameter	Description	Value
size	Size of a namespace. Unit: G	If the unit is not specified, the default unit G is used. **NOTE** If this parameter is set to **0**, the remaining unallocated space is all used for creating a namespace.
shareflag	Whether the namespace is shared	● **0**: Exclusive ● 1: Shared

Usage Guidelines

None

Example

Create an exclusive namespace of 64 GB for the **nvme0** SSD.

```
[root@localhost tool]# hioadm namespace -d nvme0 -t 0 -s 64 -f 0
creating namespace succeeded.
namespace id: 1
```

6.13.3 Attaching a Namespace to an NVMe Controller

Function

This command is used to attach a namespace to an NVMe controller.

Format

hioadm namespace -d *devicename* **-t** *operationtype* **-i** *namespaceid* **-c** *controllerid*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
operationtype	Operation Type	● **0**: Create a namespace. ● **1**: Delete a namespace. ● **2**: Attach a namespace. ● **3**: Detach a namespace. ● **6**: Set the number of namespaces. Select **2** in this command.
namespaceid	ID of a namespace.	-

Parameter	Description	Value
controllerid	ID of an NVMe controller. The ES3000 V5 SSD has only one controller. The value of this parameter is 1.	-

Usage Guidelines

If a Huawei-developed driver is installed, the setting takes effect immediately without needing to stop I/O services. If an open-source or system-provided driver is installed, the setting takes effect after the server is restarted.

Example

Attach namespace 1 of the **nvme0** SSD to NVMe controller 1.

```
[root@localhost tool]#  hioadm namespace -d nvme0 -t 2 -i 1 -c 1
attaching the namespace succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

System Response

If "attach failed" is displayed, the operation failed. The failure causes are as follows:

- The namespace is already attached to the NVMe controller.
- The NVMe controller ID is invalid.

6.13.4 Detaching a Namespace from an NVMe Controller

Function

This command is used to detach a namespace from an NVMe controller.

Format

hioadm namespace -d *devicename* **-t** *operationtype* **-i** *namespaceid* **-c** *controllerid1 controllerid2 ...*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**

Parameter	Description	Value
operationtype	Operation type.	● **0**: Create a namespace. ● **1**: Delete a namespace. ● **2**: Attach a namespace. ● **3**: Detach a namespace. ● **6**: Set the number of namespaces. Select **3** in this command.
namespaceid	ID of a namespace.	-
controllerid	ID of an NVMe controller. The ES3000 V5 SSD has only one controller. The value of this parameter is 1.	-

Usage Guidelines

Before running this command, stop all I/O services of the namespace and perform the following operations for the namespace. You do not need to stop I/O services or perform the following operations for other namespaces.

- For Linux, unmount the logical drive.
- For ESXi, delete all partitions of the logical drive.
- For Windows, make the logical drive offline, as shown in **Figure 6-4**.

Figure 6-4 Making a logical drive offline

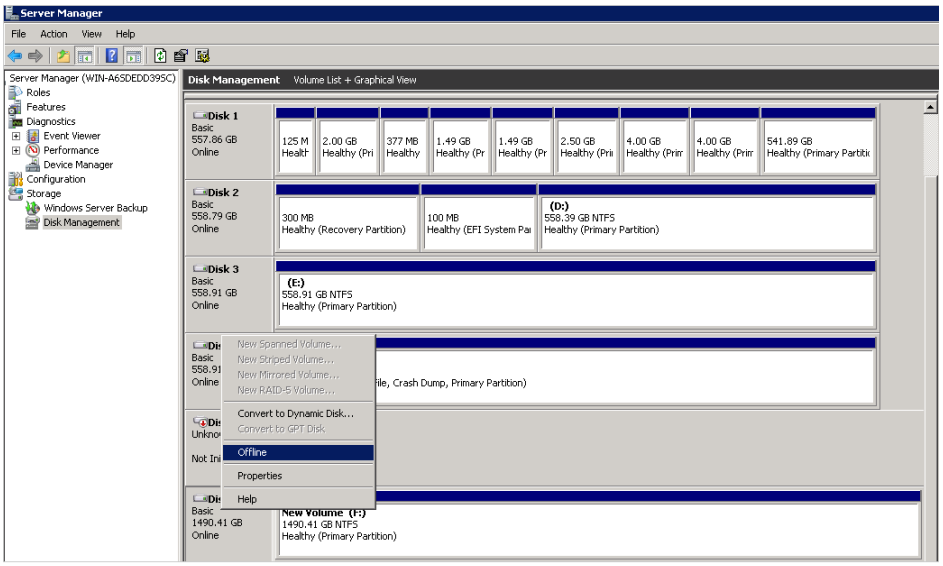

If a Huawei-developed driver is installed, the setting takes effect immediately without needing to stop I/O services. If an open-source or system-provided driver is installed, the setting takes effect after the server is restarted.

Example

Detach namespace 1 of the **nvme0** SSD to NVMe controller 1.

```
[root@localhost tool]#  hioadm namespace -d nvme0 -t 3 -i 1 -c 1
detaching the namespace succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

System Response

If "detach failed" is displayed, the operation failed. The failure causes are as follows:

- The namespace is not attached to the NVMe controller.
- The NVMe controller ID is invalid.

6.13.5 Deleting a Namespace of an SSD

Function

This command is used to delete a namespace of an SSD.

Format

hioadm namespace -d *devicename* **-t** *operationtype* **-i** *namespaceid*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
operationtype	Operation Type	● **0**: Create a namespace. ● **1**: Delete a namespace. ● **2**: Attach a namespace. ● **3**: Detach a namespace. ● **6**: Set the number of namespaces. Select **1** in this command.
namespaceid	ID of a namespace.	-

Usage Guidelines

⚠ CAUTION

Before running this command, stop all I/O operations related to the namespace.

Before running this command, detach the namespace from all controllers.

Exercise caution because deleting the namespace will also delete its storage data.

Example

Delete namespace 1 of the **nvme0** SSD.

```
[root@localhost tool]#  hioadm namespace -d nvme0 -t 1 -i 1

WARNING! The delete namespace operation will clear all data of this namespace.
WARNING! You have selected the delete namespace operation.
Proceed with the delete namespace operation? (Y|N): y
deleting the namespace succeeded.
Please remove and then install the NVMe driver or reboot the OS.
```

System Response

If "delete failed" is displayed, the operation failed. The failure cause is as follows:

- The namespace does not exist.

6.13.6 Setting the Number of Namespaces for an SSD

Function

This command is used to set the number of namespaces that can be created on a specified SSD.

Format

hioadm namespace -d *devicename* **-t** *operationtype* **-n** *num*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
operationtype	Operation type.	- **0**: Create a namespace. - **1**: Delete a namespace. - **2**: Attach a namespace. - **3**: Detach a namespace. - **6**: Set the number of namespaces. Select **6** in this command.
num	Number of namespaces.	The value ranges from 1 to 64.

Usage Guidelines

None

Example

Set the number of namespaces to be created on the **nvme0** SSD to 2.


```
[root@localhost tool]#  hioadm namespace -d nvme1 -t 6 -n 2
Setting the namespace number 2 succeeded.
```

NOTE

For the Windows operating system, when the command is executed, restart the operating system to make it take effect.

System Response

None

6.14 Querying Historical Amount of Written Data

Function

This command is used to query the amount of written data of a specified SSD. This command applies only to NVMe SSDs.

Format

hioadm loadhist -d *devicename* **-n** *days*

Parameters

Parameter	Description	Value
devicename	Name of a device.	Example: **nvme0**
days	Number of days for query.	The value ranges between 1 and 360, for example, **20**.

Usage Guidelines

None

Example

Query the amount of written data of **nvme0**.

```
[root@localhost tool]# hioadm loadhist -d nvme0 -n 20
disk write per day: 2.80
```

System Response

Parameter	Description	Value
disk write per day	Average amount of data written per day.	Percentage of the total amount, for example, **20**.

6.15 Setting and Querying Time Synchronization

6.15.1 Querying the Time of a Specified SSD

Function

This command is used to query the time of a specified SSD.

Format

hioadm synctime -d *devicename*

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0** or **sda**

Usage Guidelines

None

Example

Query the time of a specified SSD **nvme0**.

```
[root@localhost tool]# hioadm synctime -d nvme0
Current device time:Thu Mar 22 10:46:09 EDT 2018.
```

6.15.2 Synchronizing the Time of a Specified SSD

Function

This command is used to synchronize system time to a specified SSD.

Format

hioadm synctime -d *devicename* **-s**

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0** or **sda**

Usage Guidelines

None

Example

Synchronize the time of the specified SSD **nvme0**.

```
[root@localhost tool]# hioadm synctime -d nvme0 -s
Time synchronization succeeded.
Current device time:Thu Mar 22 10:51:51 EDT 2018.
[root@localhost tool]# hioadm synctime -d nvme0
Current device time: Thu Mar 22 10:51:51 EDT 2018.
```

6.16 Restoring Default Settings

Function

This command is used to restore default settings for a specified SSD. This command applies only to NVMe SSDs.

Format

hioadm restorefeatures -d *devicename*

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0**

Usage Guidelines

None

Example

Restore the default settings of a specified SSD **nvme0**.

```
linux-mazy:/ # hioadm restorefeatures -d nvme0
The restorefeatures operation will recover all features.
Arbitration
Power Management
Temperature Threshold
Error Recovery
Interrupt Coalescing
Interrupt Vector Configuration
Asynchronous Event Configuration
Software Progress Marker
Host ID
Reservation Notification Mask
```



```
Reservation Persistence

WARNING! You have selected the restorefeatures operation.
Proceed with the restorefeatures? (Y|N): y

All features are recovered.
```

6.17 Erasing a Specified SSD

Function

This command is used to erase a specified SSD.

Format

hioadm sanitize -d *devicename* [**-m** *mode*] [**-a** *action*] [**-r** *release*] [**-i** *invert*] [**-c** *count*] [**-p** *pattern*]

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0** or **sda**
mode	Indicates a restriction mode.	● 0: restricted completion mode ● 1: unrestricted completion mode
action	Indicates an operation type.	● 0: reserved ● 1: exit failure mode ● 2: block erase operation ● 3: overwrite operation ● 4: crypto erase operation **NOTE** ES3000 V5 whose firmware versions earlier than 2151 supports type 1 and 2. Firmware versions later than 2151 support types 1, 2, and 3. For details about how to query the firmware version of an SSD, see **6.2.1 Querying the Firmware Version of an SSD**.
release	Indicates whether to reassign logical blocks after the sanitize operation is completed successfully.	● 0: reassign logical blocks ● 1: no reassign logical blocks ● The default value is 0. **NOTE** If action is 1, that is, in the exit failure mode, this parameter is ignored.

Parameter	Description	Value
invert	Indicates whether the overwrite mode is inverted between two overwrites.	● 0: The overwrite mode is not inverted. ● 1: The overwrite mode is inverted. ● The default value is 0. **NOTE** This parameter is available only when action is 3 (overwrite operation). In other cases, this parameter is ignored.
count	Indicates the number of overwrite operations.	● Specifies the number of data overwrites. ● The value ranges from 1 to 16. ● The default value is 1. **NOTE** This parameter is available only when action is 3 (overwrite operation). In other cases, this parameter is ignored.
pattern	Indicates the overwrite pattern.	● A 32-bit value is used for overwriting. ● The default value is 0xFFFFFFFF. **NOTE** This parameter is available only when action is 3 (overwrite operation). In other cases, this parameter is ignored.

Usage Guidelines

None

Example

Erase a specified SSD **nvme0**.

```
linux-mazy:/ # hioadm sanitize -d nvme0 -m 0 -a 2
Set sanitize succeeded.
```

6.18 Performing an Orderly Hot Swap

6.18.1 Orderly Hot Removing a Specific SSD

Function

This command is used to orderly hot remove a specific SSD. This command applies only to NVMe SSDs.

Format

hioadm hotplug -d *devicename* -r

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0**

Usage Guidelines

- Before running this command on Linux, follow instructions in **Orderly Hot Swap of a Drive in Linux** to configure kernel parameters.
- Before running this command, stop all the I/O services on the SSD and unmount all logical drives of this SSD.

Example

Orderly hot remove a specific SSD.

```
[root@localhost tool]# hioadm hotplug -d nvme0 -r
WARNING! The hotplug operation will remove the device.
WARNING! You have selected the hotplug operation.
Proceed with the hotplug? (Y/N):Y
The hotplug remove operation succeeded, the NVMe driver has been removed from the
OS, please remove it manually.
```

6.18.2 Orderly Hot Inserting a Specific SSD

Function

This command is used to orderly hot insert a specific SSD. This command applies only to NVMe SSDs on ESXi.

Format

hioadm hotplug -d *devicename* –i

Parameters

Parameter	Description	Value
devicename	Indicates the name of an SSD.	Example: **nvme0**

Usage Guidelines

- Before running this command, hot insert the ES3000 V5 NVMe SSD to the disk slot. After the SSD is started, run the **hioadm info** command to confirm the device to which the SSD is to be inserted, for example, nvme0. Then run this command again to instruct the OS to power on the device.

Example

Orderly hot insert a specific SSD.

```
[root@localhost tool]# hioadm hotplug -d nvme0 -i
WARNING! The hotplug operation will insert the device.
WARNING! You have selected the hotplug operation.
Proceed with the hotplug? (Y/N):Y
The hotplug insert operation succeeded,the NVMe driver has been inserted to the
OS.
```

6.19 Querying or Setting the Atomic Write Status

This command is supported by ES3000 V5 2151 and later. For details about how to query the firmware version of an SSD, see **6.2.1 Querying the Firmware Version of an SSD**.

6.19.1 Querying the Atomic Write Status of a Specified SSD

Function

This command is used to query the atomic write status of a specified SSD.

Format

hioadm atomicwrite -d *devicename*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**

Usage Guidelines

None

Example

Query the atomic write status of the **nvme0** SSD.

```
[root@localhost tool]# hioadm atomicwrite -d nvme0
atomic write status: Enabled.
```

System Response

Parameter	Description	Value
Status	Atomic write status of the SSD.	Disabled/Enabled

6.19.2 Setting the Atomic Write Status for a Specified SSD

Function

This command is used to set the atomic write status for a specified SSD.

Format

hioadm atomicwrite -d *devicename* **-f** *flag*

Parameters

Parameter	Description	Value
devicename	Name of an SSD.	Example: **nvme0**
flag	Atomic write status.	● **0**: disabled ● **1**: enabled

Usage Guidelines

None

Example

Set the atomic write status of the **nvme0** SSD to **enabled**.

```
[root@localhost tool]# hioadm atomicwrite -d nvme0 -f 1
enable atomic write succeeded.
```

System Response

None

A Supported NVMe Commands

The ES3000 V5 supports all management and I/O commands required by the NVMe 1.3 protocol.

A.1 Supported NVMe Management Commands

The ES3000 V5 supports all management commands required by the NVMe 1.3 protocol. The commands are as follows:

- Delete I/O Submission Queue
- Create I/O Submission Queue
- Get Log Page
- Delete I/O Completion Queue
- Create I/O Completion Queue
- Identify
- Abort
- Set Feature
- Get Feature
- Asynchronous Event Request

The ES3000 V5 also supports the following optional management commands defined by NVMe 1.3:

- Namespace Management
- Firmware Commit
- Firmware Image Download
- Device Self-test
- Namespace Attachment
- Directive Send
- Directive Receive
- Format NVM (Huawei drivers and tools must be used if the format is DIF.)
- Sanitize

A.2 Supported NVMe I/O Commands

The ES3000 V5 supports all mandatory NVMe I/O commands:

- Flush
- Write
- Read

The ES3000 V5 also supports the following optional commands:

- Dataset Management (De-allocate only)

A.3 Supported Get Log Page Commands

The ES3000 V5 supports all **get log page** commands required by the NVMe 1.3 protocol. The commands are as follows:

- Error information (log identifier 01h)
- SMART/Health information (log identifier 02h)
- Firmware slot information (log identifier 03h)
- Changed namespace list (log identifier 04h)
- Commands Supported and Effects (Log Identifier 05h)
- Device Self-test (Log Identifier 06h)
- Sanitize Status (Log Identifier 80h)

ES3000 V5 also supports the following custom log command:

- Vendor Unique SMART/Health Information (Log Identifier C0h)
- Vendor Unique Addition SMART/Health Information (Log Identifier CAh)

A.4 SMART Properties

The ES3000 V5 supports the same SMART properties as the NVMe 1.3 protocol. **Table1**, **Table2** and **Table3** list the SMART properties.

Table A-1 SMART properties (log identifier 02h)

<table>
<tr><th>Bytes</th><th colspan="2">Description</th></tr>
<tr><td rowspan="4">0</td><td colspan="2">Critical Warning: This domain indicates a critical warning about the controller status. Critical warnings are returned to the host through asynchronous events.</td></tr>
<tr><td>Bit</td><td>Description</td></tr>
<tr><td>0</td><td>Value 1 indicates that the redundant space is smaller than the threshold.</td></tr>
<tr><td>1</td><td>Value 1 indicates that the temperature is higher or lower than a critical threshold.</td></tr>
</table>

Bytes	Description	
	2	Value **1** indicates that the component reliability is reduced due to a critical media error.
	3	Value **1** indicates that the media is read-only.
	4	Value **1** indicates that the non-volatile component of the controller fails. This domain is valid only when the controller is integrated with a non-volatile component.
	7:5	Reserved
2:1	**Temperature**: indicates the overall component temperature in unit of Kelvin. An asynchronous event is reported to the host when the temperature exceeds the threshold.	
3	**Available Spare**: indicates the percentage (0 to 100%) of the available redundant space.	
4	**Available Spare Threshold**: indicates the threshold of the available redundant space. An asynchronous event is reported to the host when the available redundant space is smaller than the threshold.	
5	**Percentage Used**: indicates the ratio of the used life to the preset life. 100 indicates that the component has reached the preset life, but may still function properly. The value can be greater than 100. The value is always set to 255 if it is greater than 254.	
31:6	Reserved	
47:32	**Data Units Read**: indicates the number of 512-byte sectors read from the controller by the host. 1 indicates that the host has read 1000 512-byte sectors, which do not include metadata. If an LBA is not 512 bytes, the controller converts it into 512-byte sectors (hexadecimal) for calculation.	
63:48	**Data Units Written**: indicates the number of 512-byte sectors read into the controller by the host. 1 indicates that the host has written 1000 512-byte sectors, which do not include metadata. If an LBA is not 512 bytes, the controller converts it into 512-byte sectors (hexadecimal) for calculation.	
79:64	**Host Read Commands**: indicates the number of read commands sent to the controller.	
95:80	**Host Write Commands**: indicates the number of write commands sent to the controller.	
111:96	**Controller Busy Time** indicates the busy period when the controller processes I/O commands. The controller is busy from command delivery (to the SQ) to the response (to the CQ). The value is in unit of minute.	
127:112	**Power Cycles**: indicates the number of power cycles (power-on and -off).	
143:128	**Power On Hours**: indicates the number of power-on hours, excluding the sleeping time.	

Bytes	Description
159:144	**Unsafe Shutdowns**: indicates the number of abnormal power-offs. The value is increased by 1 if CC.SHN is not received after an abnormal power-off.
175:160	**Media Errors**: indicates the number of unrecoverable data integrity errors detected by the controller, including incorrigible error-correcting code (ECC) errors, cyclic redundancy check (CRC) errors, and LBA tag errors.
191:176	**Number of Error Information Log Entries**: indicates the number of error log entries during the controller life cycle.
195:192	**Warning Composite Temperature Time**: indicates the time when the temperature of the entire component becomes greater than or equal to the overtemperature threshold (controller Identify information WCTEMP field) but is less than the ultra-high temperature threshold (CCTEMP field in controller identify information), in minutes.
199:196	**Critical Composite Temperature Time**: indicates the time when the temperature of the entire component becomes greater than the ultra-high temperature threshold (CCTEMP field in controller identify information), in minutes.
201:200	**Temperature Sensor 1**: indicates the value detected by temperature sensor 1. Unit: Kelvin.
203:202	**Temperature Sensor 2**: indicates the value detected by temperature sensor 2. Unit: Kelvin.
205:204	**Temperature Sensor 3**: indicates the value detected by temperature sensor 3. Unit: Kelvin.
207:206	**Temperature Sensor 4**: indicates the value detected by temperature sensor 4. Unit: Kelvin.
209:208	**Temperature Sensor 5**: indicates the value detected by temperature sensor 5. Unit: Kelvin.
211:210	**Temperature Sensor 6**: indicates the value detected by temperature sensor 6. Unit: Kelvin.
213:212	**Temperature Sensor 7**: indicates the value detected by temperature sensor 7. Unit: Kelvin.
215:214	**Temperature Sensor 8**: indicates the value detected by temperature sensor 8. Unit: Kelvin.
511:216	Reserved

Table A-2 Customized SMART properties (log identifier C0h)

Bytes	Description
1:0	Buck converter voltage. Unit: mV

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>3:2</td><td>Current power consumption. Unit: mW</td></tr>
<tr><td>5:4</td><td>Highest power consumption recorded. Unit: mW</td></tr>
<tr><td>7:6</td><td>Highest SSD temperature recorded. Unit: Kelvin</td></tr>
<tr><td>9:8</td><td>Number of SSD overtemperature occurrences</td></tr>
<tr><td>13:10</td><td>Period of SSD overtemperature. Unit: minute</td></tr>
<tr><td>15:14</td><td>Highest controller temperature recorded. Unit: Kelvin</td></tr>
<tr><td>17:16</td><td>Number of controller overtemperature occurrences</td></tr>
<tr><td>21:18</td><td>Period of controller overtemperature. Unit: minute</td></tr>
<tr><td>23:22</td><td>Capacitor's capacity. Unit: F</td></tr>
<tr><td>25:24</td><td>Capacitor's voltage. Unit: V</td></tr>
<tr><td>27:26</td><td>Number of discharge/charge cycles of the capacitor</td></tr>
<tr><td>28</td><td>Capacitor health status<table><tr><th>Value</th><th>Description</th></tr><tr><td>0h</td><td>The capacitor is working properly</td></tr><tr><td>1h</td><td>The capacitor has failed</td></tr><tr><td>FFh~2h</td><td>Reserved</td></tr></table></td></tr>
<tr><td>29</td><td>Reserved</td></tr>
<tr><td>33:30</td><td>PCIe refclock PLL lock loss count</td></tr>
<tr><td>37:34</td><td>CRC error count of PCIe port 0</td></tr>
<tr><td>41:38</td><td>Reserved</td></tr>
<tr><td>45:42</td><td>E2E validation failure count of the controller</td></tr>
<tr><td>47:46</td><td>Number of failed CEs</td></tr>
<tr><td>51:48</td><td>Total number of blocks</td></tr>
<tr><td>55:52</td><td>Total number of bad blocks</td></tr>
<tr><td>59:56</td><td>Number of new bad blocks</td></tr>
<tr><td>63:60</td><td>Total number of pages</td></tr>
<tr><td>71:64</td><td>Loss capacity size</td></tr>
<tr><td>75:72</td><td>Number of FTL table rebuilds due to power outage</td></tr>
<tr><td>79:76</td><td>Number of deep rebuilds of the FTL table</td></tr>
</table>

Bytes	Description
95:80	Number of flash read failures.
111:96	Number of flash read operations.
127:112	Reserved
143:128	Number of flash write failures.
159:144	Number of flash write operations.
175:160	Reserved
191:176	Number of flash erasure failures.
207:192	Number of flash erasure operations.
215:208	Reserved
223:216	GC read count
231:224	BS read count
235:232	Drive voltage (unit: mV)
239:236	Reserved
255:240	Reserved
271:256	Reserved
287:272	Number of FLASH RR failures
303:288	Number of FLASH RR operations
319:304	Reserved
335:320	Number of UNC errors found during background inspections
351:336	Reserved
367:352	Background inspection data volume (unit: GB)
383:368	Number of UNC errors recorded in the pending list
399:384	Reserved
415:400	Reserved
419:416	Deep rebuild flag
423:420	Reserved
427:424	Before-delivery bad block count 1
431:428	Before-delivery bad block count 2
435:432	Number of bad programming blocks
439:436	Number of bad erase blocks

Bytes	Description
443:440	Number of UNC bad blocks
447:444	Number of DIE bad blocks
451:448	Number of DDR single-bit ECC errors
459:452	I/O write count
463:460	I/O write failure count
471:464	GC write count
475:472	GC write failure count
483:476	Read correction count 1
491:484	Read correction count 2
499:492	Read correction count 3
507:500	Read correction count 4
515:508	Read correction count 5
519:516	Read correction count 6
527:520	GC read data volume (unit: GB)
531:528	Forcible data migration count
535:532	Memory error count
539:536	Free Block Num 1
543:540	Free Block Num 2
547:544	Free Block Num 3
551:548	Free Block Num 4
555:552	Free Block Num 5
559:556	Free Block Num 6
563:560	Free Block Num 7
567:564	Free Block Num 8
571:568	Average erase count
575:572	Maximum erase count
583:576	Read correction count 7
587:584	Current system time in seconds (number of seconds elapsed since January 1, 1970)
595:588	Reserved

Bytes	Description	
603:596	Reserved	
635:604	Reserved	
643:636	Detailed alarm 1	
	Bit	**Description**
	0	Firmware bootloader error
	1	Storage firmware failure
	2	BBU check exception
	3	BBU check exception
	4	Insufficient available space
	5	Lifespan alarm
	6	Temperature alarm
	7	Storage unit failure
	63:8	Reserved
651:644	Detailed alarm 2	
	Bit	**Description**
	0	Quick formatting failure (Secure Erase Settings=0)
	1	Secure formatting failure (Secure Erase Settings=1)
	2	Rebuild failure 1
	3	Rebuild failure 2
	4	Rebuild failure 3
	5	Rebuild failure 4
	6	Rebuild failure 5
	7	Power-on check failure 1
	8	Power-on check failure 2
	9	Power-on check failure 3
	10	Power-on check failure 4
	11	Internal error 1
	12	Internal error 2
	13	Internal error 3
	14	Internal error 4

Bytes	Description	
	15	BBU check exception
	16	Reserved
	17	Internal error 5
	18	Internal error 6
	19	Internal error 7
	20	Internal error 8
	21	Internal error 9
	22	Internal error 10
	23	Internal error 11
	24	Internal error 12
	25	Rebuild failure 6
	26	Internal error 13
	27	Internal error 14
	28	BBU inspection exception
	29	I/O read failure
	30	I/O write failure
	31	Reserved
	32	Reserved
	33	Deep rebuild is complete and enter the read-only mode
	34	Internal error 15
	35	Internal error 16
	63:36	Reserved
659:652	Alarm historical value: bit definition same as that of 643:636	
667:660	Exception historical value: bit definition same as that of 651:644	
671:668	Port 0 cold reset count	
675:672	Reserved	
679:676	Port 0 warm reset count	
683:680	Reserved	
687:684	Port 0 hot reset count	

<table>
<tr><th>Bytes</th><th colspan="2">Description</th></tr>
<tr><td>691:688</td><td colspan="2">Reserved</td></tr>
<tr><td>695:692</td><td colspan="2">Port 0 function reset count</td></tr>
<tr><td>699:696</td><td colspan="2">Reserved</td></tr>
<tr><td>703:700</td><td colspan="2">Port 0 subsystem reset count</td></tr>
<tr><td>707:704</td><td colspan="2">Reserved</td></tr>
<tr><td>711:708</td><td colspan="2">Port 0 controller reset count</td></tr>
<tr><td>715:712</td><td colspan="2">Reserved</td></tr>
<tr><td>719:716</td><td colspan="2">Port 0 symbol unlock count</td></tr>
<tr><td>723:720</td><td colspan="2">Reserved</td></tr>
<tr><td>727:724</td><td colspan="2">Current power consumption level</td></tr>
<tr><td rowspan="5">731:728</td><td colspan="2">Health of each slot firmware</td></tr>
<tr><td>Bit</td><td>Description</td></tr>
<tr><td>0</td><td>Reserved</td></tr>
<tr><td>1</td><td>Firmware health of slot 1</td></tr>
<tr><td>2</td><td>Firmware health of slot 2</td></tr>
<tr><td>31:43</td><td>Reserved</td></tr>
<tr><td>732</td><td colspan="2">Drive status.</td></tr>
</table>

Table A-3 Customized SMART properties (Log Identifier CAh)

Bytes	Description
0	Program Fail Count
2:1	Reserved
3	Normalized Value (Program Fail)
4	Reserved
10:5	Current Raw Value (Program Fail)
11	Reserved
12	Erase_Fail_Count
14:13	Reserved
15	Normalized Value (Erase_Fail_Count)

Bytes	Description
16	Reserved
22:17	Current Raw Value (Erase_Fail_Count)
23	Reserved
24	Wear Leveling Count
26:25	Reserved
27	Normalized Value (Wear Leveling)
28	Reserved
34:29	Current Raw Value (Wear Leveling)
35	Reserved
36	End To End Error Detection Count
38:37	Reserved
39	Normalized Value (End To End Error Detection)
40	Reserved
46:41	Current Raw Value (End To End Error Detection)
47	Reserved
48	CRC Error Count
50:49	Reserved
51	Normalized Value (CRC Error)
52	Reserved
58:53	Current Raw Value (CRC Error)
59	Reserved
60	Timed Workload Media Wear Count
62:61	Reserved
63	Normalized Value (Timed Workload Media Wear)
64	Reserved
70:65	Current Raw Value (Timed Workload Media Wear)
71	Reserved
72	Timed Workload Host Reads Count
74:73	Reserved
75	Normalized Value (Timed Workload Host Reads)

Bytes	Description
76	Reserved
82:77	Current Raw Value (Timed Workload Host Reads)
83	Reserved
84	Timed Workload Timer Count
86:85	Reserved
87	Normalized Value (Timed Workload Timer)
88	Reserved
94:89	Current Raw Value (Timed Workload Timer)
95	Reserved
96	Thermal Throttle Status Count
98:97	Reserved
99	Normalized Value (Thermal Throttle Status)
100	Reserved
106:101	Current Raw Value (Thermal Throttle Status)
107	Reserved
108	Retry Buffer Overflow Count
110:109	Reserved
111	Normalized Value (Retry Buffer Overflow)
112	Reserved
118:113	Current Raw Value (Retry Buffer Overflow)
119	Reserved
120	Pll Lock Loss Count
122:121	Reserved
123	Normalized Value (Pll Lock Loss)
124	Reserved
130:125	Current Raw Value (Pll Lock Loss)
131	Reserved
132	Nand Bytes Written Count
134:133	Reserved
135	Normalized Value (Nand Bytes Written)

Bytes	Description
136	Reserved
142:137	Current Raw Value (Nand Bytes Written)
143	Reserved
144	Host Bytes Written Count
146:145	Reserved
147	Normalized Value (Host Bytes Written)
148	Reserved
154:149	Current Raw Value (Host Bytes Written)
155	Reserved
156	F6 (System Area Life Remaining)
158:157	Reserved
159	Normalized Value(System Area Life Remaining)
160	Reserved
166:161	Current Raw Value(System Area Life Remaining)
167	Reserved

A.5 Supported Set Feature Commands

Besides SMART information, the host can use the **Get Features** command to obtain feature parameters related to working and health status. The ES3000 V5 can also use the following **Set Features** commands defined by NVMe 1.3 to change the feature parameter settings:

- Arbitration (Feature Identifier 01h)
- Power Management (Feature Identifier 02h)
- Temperature Threshold (Feature Identifier 04h)
- Error Recovery (Feature Identifier 05h)
- Number of Queues (Feature Identifier 07h)
- Interrupt Coalescing (Feature Identifier 08h)
- Interrupt Vector Configuration (Feature Identifier 09h)
- Write Atomicity (Feature Identifier 0Ah)
- Asynchronous Event Configuration (Feature Identifier 0Bh)
- Timestamp (Feature Identifier 0Eh)
- Software Progressor Marker (Feature Identifier 80h)

Table A-4 04h-Temperature Threshold-Command Dword 11

Bit	Description	
31:22	**Reserved**	
21:20	**Threshold Type Select (THSEL)**: This domain is used to set the temperature threshold type.	
	Value	**Description**
	00b	Upper temperature threshold
	01b	Lower temperature threshold
	10b~11b	Reserved
19:16	**Threshold Temperature Select (TMPSEL)**: This domain is used to set the temperature threshold sensor.	
	Value	**Description**
	0000b	Composite Temperature
	0001b	Temperature sensor 1 (controller)
	0010b	Temperature sensor 2 (not supported)
	0011b	Temperature sensor 3 (not supported)
	0100b	Temperature sensor 4 (not supported)
	0101b	Temperature sensor 5 (not supported)
	0110b	Temperature sensor 6 (not supported)
	0111b	Temperature sensor 7 (not supported)
	1000b	Temperature sensor 8 (not supported)
	1001b~1110b	Reserved
	1111b	All sensors
15:00	**Temperature Threshold (TMPTH)** indicates the SSD or controller temperature threshold in unit of Kelvin.	

The ES3000 V5 also supports the following vendor-customized **set feature** command:

Set/Get Device Time (Feature Identifier C0h)

Table A-5 Set/Get Device Time (Feature Identifier C0h)

Bit	Description
31:00	**Host Time (s)** records the number of seconds from 1970-1-1 00:00:00 until now.

A.6 Supported Vendor Specific Command

The ES3000 V5 support vendor specific command for information query and management functions.

NOTE

To be compatible with Huawei ES3000 V3 SSDs, the custom management commands (customized Opcode:0xC0 and 0xC1) are inherited. These commands are called custom inheritance commands in this section. However, you are advised to use the newly defined interfaces (customized Opcode:0xC2 and Opcode:0xC5).

Table A-6 Opcode

Opcode (07)	Opcode (06:02)	Opcode (01:00)	Opcode	O/M	Namespace Identifier Used	Command
Generic Command	Function	Data Transfer				
1b	100 00b	00b	C0h	O	NO	legacy query command
1b	100 00b	01b	C1h	O	NO	legacy configuration command
1b	100 00b	10b	C2h	O	NO	new query command
1b	100 00b	01b	C5h	O	NO	new configuration command

A.6.1 New Commands

A.6.1.1 Query Commands

The commands use the PRP1, PRP2, and Dword10 to Dword15 fields. Enter reserved value 0 for other fields that are not involved.

Table A-7 PRP1

Bit	Description
63:00	**PRP Entry 1 (PRP1):** indicates the data buffer. This buffer cannot have more than one physical discontinuity. The minimum unit is 4 KB.

Table A-8 PRP2

Bit	Description
63:00	**PRP Entry 2 (PRP2)**: indicates the second PRP entry. If PRP Entry 1 specifies enough space for the data structure, this field is reserved. If the data structure requires two PRP entries, this field represents the second PRP entry address. If data of more than two PRP entries is transferred, this field points to a PRP list.

Table A-9 Command Dword 10

Bit	Description
31:00	**Data_Len**: This field indicates the data length in unit of Dword. If data transfer is not needed, the value is **0h**.

Table A-10 Command Dword 12

Bit	Description
31:00	**Sub_OpCode:** subcommand number

A.6.1.1.1 Obtaining Logs and Statistical Information

This command is used to obtain vendor logs from the SSD. Each command sent from the host can obtain a maximum of 4KB of logs. The host obtains the total log size through the vendor log length query command and then sends commands repeatedly to obtain all logs. Besides the total data size, the commands must also specify an offset value, where 1 indicates 4 KB.

This command uses the PRP1, PRP2, Dword10, Dword12, and Dword13 fields. Enter reserved value 0 for other fields that are not involved.

This command distinguishes the type of a returned log by Sub_OpCode, as shown in the following table:

Table A-11 Sub_OpCode

Sub_OpCode	Description
0x010009h	Obtain full logs
0x020009h	Obtain key logs
0x030009h	Obtain CoreDump logs
0x040009h	Obtain statistical information

Table A-12 Command Dword 13

Bit	Description
31:00	**Offset**: log offset. 1 indicates 4 KB.

The valid log length and log offset are returned to the host through the DW0 of the CQ.

Table A-13 Completion Queue Entry Dword 0

Bit	Description
31:16	**Log_Length**: obtained log length.
15:00	**Offset**: obtained log offset value.

NOTE

Each command sent from the host can obtain a maximum of 4 KB logs. If the size exceeds 4 KB, the firmware returns the status code "NVME_SC_INVALID_FIELD = 0x2".

Full logs and key logs are composed of a large number of 64-byte log entries, whereas CoreDump logs are tiled characters.

Table A-14 Overall structure of full logs and key logs

Byte	Description
63:00	Entry 1 of full logs and key logs
127:64	Entry 2 of full logs and key logs
…	…
Length-1:Length-64	Entry n of full logs and key logs

Table A-15 Structure of full logs and key logs entries

Byte	Description
07:00	**Current_time**: number of seconds since January 1, 1970.
11:08	**Power_on_num**: number of power-on times since delivery.
27:12	**Func_name**: function name, which consists of 16 ASCII characters and does not end with \0.
29:28	**Log_id**: log ID. For details, see A-15.
63:30	**Log_content**: log plain text, which consists of 34 ACSII characters and does not end with \0. Usually the plain text ends with \r \n.

Table A-16 Log_id structure

Bit	Description
15:12	**Vension**: version number of the log format. The version must be later than 0 and earlier than 0xA.
11:11	**ConnectFlag**: indicates whether the next log and the unfinished part of the current log can be displayed together. **0** indicates yes and **1** indicates no.
10:4	**Reservd**: reserved bit
3:0	Level: log level. Logs are classified into the following levels: 0: LOG_LEVEL_DEBUG, indicating debugging information; 1: LOG_LEVEL_INFO, indicating common information, which is a normal log; 2: LOG_LEVEL_WARNING, indicating alarm information, which is an exception log. 3: LOG_LEVEL_EVENT, indicating a key process, which is a normal log. 4: LOG_LEVEL_FATAL, indicating a fatal error, which is an exception log. 5: LOG_LEVEL_OPERATE, indicating user operation logs. User operation logs are saved as full logs for easy filtering.

Table A-17 CoreDump log structure

Byte	Description
MSB:0	CoreDump logs are comprised of tiled characters.

NOTE

The process for obtaining complete logs is as follows:

1. Send the vendor-defined log length obtaining command to obtain the total length of valid logs and apply for log storage space based on the total length of valid logs.
2. Send the log obtaining command to obtain the log content. A maximum of 4 KB data in logs can be obtained each time the command is sent. Therefore, the number of times for obtaining complete logs is calculated based on the total length of the valid log. Then, logs are obtained cyclically. Each time the command for obtaining logs is sent, the length of the obtained data is specified. In addition, you need to specify an offset. One offset indicates 4 KB.
3. After obtaining complete logs, parse the logs and write them to the target file. Full logs and key logs are composed of a large number of 64-byte log entries (see A-15), whereas CoreDump logs are tiled characters. The content of full logs and key logs are parsed by item and written to the corresponding title in the target file. The content of CoreDump logs is written to the target file directly.

A.6.1.1.2 Obtaining SSD ID Parameters

This command is used to obtain SSD ID parameters, including electronic label information, part number (PN), World Wide Name (WWN), Model, and serial number (SN). The parameters are returned to the host through PRP1 and PRP2.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-18 lists parameter types.

Table A-18 Sub_OpCode

Sub_OpCode	Description
0x200009h	Electronic label information
0x330009h	PN
0x210009h	WWN
0x230009h	Model
0x220009h	SN
0x0A0005h	Flash ID

- **Electronic label information**

The subcommand is used to set the electronic label of the SSD. The host uses PRP1 and PRP2 to send electronic label information in the format described in **Table3**.

In the following table, "Field" indicates the combination of the field name, equal sign, field value, and carriage return. For example, "VendorName=Huawei\r\n" is a string, where "VendorName" is the field name and "Huawei" is the field value.

Table A-19 Electronic label information

Area	Description	Maximum Field Length (Bytes)	Maximum Field Value Length (Bytes)	Remarks
Information format area	/$[ArchivesInfo Version]\r\n	26	N/A	If a line starts with **/$**, the line is not displayed only in the network management system. **$** is an ASCII character with the character code 0x24. The data needs to be displayed during query on network elements. The current label version is 3.0, which indicates the number of fields. Usually content changes do not cause version changes.
	/$ArchivesInfoVersion=<*Version Number*>\r\n	30	6	
	\r\n	2	N/A	Two blank lines are between the information format area and the information description area.
	\r\n	2	N/A	
Information description area	[Board Properties]\r\n	20	N/A	
	BoardType=<*Model*>\r\n	30	18	Product model number
	BarCode=<*SN*>\r\n	30	20	SN on the paper label
	Item=<*BBOM Code*>\r\n	30	23	PN on the paper label
	Description=<*English Description*>\r\n	260	246	The line contains a maximum of 260 bytes.
	Manufactured=<*Manufacturing Date*>\r\n	25	10	Format: YYYY-MM-DD.

Area	Description	Maximum Field Length (Bytes)	Maximum Field Value Length (Bytes)	Remarks
	VendorName =<*Manufacturer Name*>\r\n	50	37	FRUs made by Huawei or Huawei's outsourced vendors are all branded by Huawei.
	IssueNumber =<*Issue Number*>\r\n	20	6	Item revision (FRU hardware version) on the paper label
	CLEICode=<*CLEI Code*>\r\n	25	14	This field is used by products complying with specifications of Telcordia in North America.
	BOM=<*BOM Code*>\r\n	50	44	Item of a finer granularity. This field is reserved for products whose Huawei item codes are insufficient for customers' asset management.

- **PN**

The subcommand is used to set the SSD part number (PN). The host uses PRP1 and PRP2 to send PN information in the format described in **Table4**.

Table A-20 PN

Bytes	Description
23:00	SN

- **WWN**

The subcommand is used to set the SSD WWN. The host uses PRP1 and PRP2 to send WWN information in the format described in **Table5**.

Table A-21 WWN

Bytes	Description
15:00	WWN

- **Model**

The subcommand is used to set the SSD model number. The host uses PRP1 and PRP2 to send model information in the format described in **Table6**.

Table A-22 Model

Bytes	Description
15:00	Model

- **SN**

The subcommand is used to set the SSD SN. The host uses PRP1 and PRP2 to send SN information in the format described in **Table7**.

Table A-23 SN

Bytes	Description
19:00	SN

- **FLASH ID**

This command is used to obtain the FLASH ID. This command uses PRP1, PRP2, and Dword10. The data length is specified by Dword10 (fixed at 2).

Table A-24 FLASH ID

BIT	Description
63:00	FLASH ID

A.6.1.1.3 Obtaining the Enablement Status of Functions

This command is used to obtain the enablement status of a customized function through the controller.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Sub_OpCode describes parameter types.

Table A-25 Sub_OpCode

Sub_OpCode	Description
430006h	Flash inspection enablement
F00009h	Capacitor inspection enablement

Sub_OpCode	Description
F10009h	DDR inspection enablement
EF0009h	Temperature inspection enablement
0D0005h	CE concurrency enablement
0E0005h	Read Retry enablement
180006h	Wear leveling enablement
190006h	Inter-chip RAID enablement
1A0006h	Garbage collection (GC) enablement

The enablement status of each function is returned through CQ Dword0.

Table A-26 CQ-Dword0

Bit	Description
00	**Enable Status**: **1** indicates that the function is enabled. **0** indicates that the function is disabled.
31:01	Reserved

A.6.1.1.4 Querying the Maximum Available Capacity of the SSD

This command is used to query the maximum available capacity that can be set. This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-27 Sub_OpCode

Sub_OpCode	Description
0x050009h	Query the maximum available capacity of an SSD.

The command returns the information to the host through PRP1 and PRP2. The format of the returned data is as follows:

Table A-28 Querying the maximum available capacity of the SSD

Bytes	Description
7:0	Maximum available capacity of the SSD (unit: byte)

A.6.1.1.5 Obtaining the Vendor Log Length

This command is used to obtain the length of vendor logs (including the log parsing package, full logs, and key logs).

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

This command distinguishes the type of a returned log by Sub_OpCode, as shown in the following table:

Table A-29 Sub_OpCode

Sub_OpCode	Description
0x000009h	Log parsing package
0x080009h	Full logs
0x090009h	Key logs
0x120009h	CoreDump information
0x420009h	Statistical information

The obtained valid log length is returned to the host software as **Table3**.

Table A-30 Completion Queue Entry Dword 0

Bit	Description
31:0	Valid log length (unit: byte)

A.6.1.1.6 Querying the PCB Version

This command is used to obtain the PCB version.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-31 Sub_OpCode

Sub_OpCode	Description
0x350009h	Query the PCB version.

Table A-32 CQ DW0

Bit	Description
31:0	PCB version

A.6.1.1.7 Obtaining Self-Diagnosis Results

This command is used to obtain the self-diagnosis result from the formal firmware.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-33 Sub_OpCode

Sub_OpCode	Description
0x140000h	Obtain self-diagnosis results.

The obtained cache information is returned to the host through the CQ Dword0. If a bit value is **1**, the corresponding item has a fault. If a bit value is **0**, the corresponding item is passed.

Table A-34 CQ-Dword0

Bit	Description
0	Overall test result
1	Firmware error test (front-end error check: whether the drive is in the controller drive mode, whether the drive is in read-only mode, and whether the available space is below the threshold)
2	NAND test
3	NOR connectivity test
4	DDR uncorrectable error test
5	IBUF uncorrectable error test
6	Capacitor test
7	Blank block test: insufficient blank blocks
8	Life test

Bit	Description
9	NOR patrol test
10	Firmware uncorrectable error test: four invalid firmware backups
11	System area test too many bad blocks in the NAND system area
12	Bad block test: too many bad blocks in the NAND data area
13	Controller test: invalid controller
14	Power down fast interrupt request (FIQ) test
15	Formatting failure test
16	Db poweron fail test
17	Reserved
18	Lbn miss happen test
19	Deep rebuild happen Test
24:20	Reserved
25	**1**: Firmware needs to be upgraded to restore SSDs.
26	**1**: Data loss risks exist.
27	**1**: SSDs can be restored, but risks exist.
28	**1**: SSDs can be restored.
29	**1**: Forcibly restore SSDs so that the SSDs can function properly.
30	**1**: Restart or power on and then off SSDs to restore SSDs.
31	**1**: The fault cannot be rectified. Contact Huawei support engineers.

A.6.1.1.8 Obtaining the DIEID of a Chip

This command is used to obtain the DIEID, which is the unique identifier of a chip and is used to track and manage information of the entire lifecycle (including chip manufacturing, chip test, board processing, server test, and live network operation).

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-35 Sub_OpCode

Sub_OpCode	Description
0x240009h	Obtain the chip DIEID.

Test results are returned to the host through the host's physical memory that PRP1 and PRP2 point to. The data structure is defined in the following table.

Table A-36 DIEID structure

Bytes	Description
3:0	Chip ID
7:4	Version
11:8	DIEID [31:0], indicates the first 4 bytes of DIEID.
……	-
27:24	DIEID [160:129], indicates the last 4 bytes of DIEID.
31:28	Reserved

A.6.1.1.9 Querying the Average DWPD

This interface is used to query the average DWPD in N (N <= 360) days. Dword13 indicates the number of days, and the result is returned by Dword0 of CQ.

This command uses the PRP1, PRP2, Dword10, Dword12, and Dword13 fields. Enter reserved value 0 for other fields that are not involved.

Table A-37 Sub_OpCode

Sub_OpCode	Description
0x0E0000h	Query the average DWPD value.

Table A-38 Command Dword13

Bit	Description
31:00	**Number of days**: valid value (1-360)

Table A-39 Cq Command Dword0

Bit	Description
31:00	**DWPD value**: 100 indicates 1. (If the average daily written data volume is 0.8 TB and the drive user volume is 1.6 TB, the DWPD value is 0.5 and the interface result is 50.)

A.6.1.1.10 Querying UniqueID Information

This interface is used to obtain the UniqueID information. The host issues a command for obtaining the data length, and then issues a command for obtaining the content based on the length. The content obtained is in the string format and the length does not contain '\0' at the end.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-40 Sub_OpCode

Sub_OpCode	Description
0x00005h	Obtain the length.
0x10005h	Obtain the content.

When Sub_OpCode delivered to obtain the uniqueID length is 0x00005h, the length is returned by DWORD0 of CQ.

Table A-41 Cq Command Dword0

Bit	Description
31:00	**Length:** The unit is byte.

When Sub_OpCode delivered to obtain the uniqueID content is 0x10005h, the length is returned by Dword10 and must be the same as the length obtained previously. The UniqueID content is transmitted by the memory that PRP points to, and the content is in the string format.

Table A-42 Command Dword10

Bit	Description
31:00	**Length:** The unit is Dword.

A.6.1.1.11 Querying the Namespace Pending List

This command is used to query the pending list of a specified Namespace UNC LBA.

Table A-43 Sub_OpCode

Sub_OpCode	Description
0x040000h	Query the pending list of the namespace.

This command uses the PRP1, PRP2, Dword10, Dword12, and Dword13 fields. Enter reserved value 0 for other fields that are not involved.

Table A-44 Command Dword 10

Bit	Description
31:00	**Data_Len**: Data transmission length. The unit is Dword. If **TYPE** is set to **0**, the returned **Transfer Length** is 1024 bytes. If **TYPE** is set to **1**, the length of the returned parameter is 4096 bytes.

Table A-45 Command Dword 13

Bit	Description
31:00	**Type**: If this parameter is set to **0**, the LBA address of UNC is returned. Every four bytes represent the LBA address of UNC. If **TYPE** is set to **1**, the UNC LBA address and UNC length are returned. Each item contains 16 bytes. The first 8 bytes indicate the LBA address and the last 8 bytes indicate the UNC length. The unit is sector. For example, the sector size of the 4K data logic is 512, and the corresponding length is 8.

A.6.1.1.12 Obtaining the 4-Point Digital Eye Diagram of the PCIe

This command is used to obtain the 4-point digital eye diagram of the PCIe.

Table A-46 Sub_OpCode

Sub_OpCode	Description
0x060000h	Obtain the 4-point digital eye diagram of the PCIe.

This command uses the PRP1, PRP2, Dword10, and Dword13 to Dword15 fields. Enter reserved value 0 for other fields that are not involved.

Table A-47 Command Dword 10

Bit	Description
31:00	**Data_Len**: data transmission length, in Dword. The length is 16 bytes.

Table A-48 Command Dword 13

Bit	Description
15:00	**Macro**: **0** corresponds to Macro0, and **1** corresponds to Macro0.
31:16	**Lane**: Each macro cell corresponds to two lanes: lane 0 and lane 1.

Table A-49 Command Dword 14

Bit	Description
31:00	**DataRate**: transmission rate.

Table A-50 Command Dword 15

Bit	Description
31:00	**Average Time**: average count.

Table A-51 describes the command data format.

Table A-51 Data format

Byte	Description
03:00	Upper coordinate.
07:04	Lower coordinate.
11:08	Left coordinate.
15:12	Right coordinate.

A.6.1.1.13 Obtaining the Full Digital Eye Diagram of the PCIe

This command is used to obtain the full digital eye diagram of the PCIe.

Table A-52 Sub_OpCode

Sub_OpCode	Description
0x070000h	Obtain the full digital eye diagram of the PCIe.

This command uses the PRP1, PRP2, Dword10, and Dword13 to Dword15 fields. Enter reserved value 0 for other fields that are not involved.

Table A-53 Command Dword 10

Bit	Description
31:00	**Data_Len**: data transmission length, in Dword. The length is 16 KB.

Table A-54 Command Dword 13

Bit	Description
15:00	**Macro**: **0** corresponds to Macro0, and **1** corresponds to Macro0.
31:16	**Lane**: Each macro cell corresponds to two lanes: lane 0 and lane 1.

Table A-55 Command Dword 14

Bit	Description
31:00	**DataRate**: transmission rate.

Table A-56 Command Dword 15

Bit	Description
15:00	**IsShowInner**: **1**: to scan the inner eye. **2**: not to scan the inner eye.
31:16	**ScanWinInner**: scanning layers (1 – 4)

Command data format: Each layer of the scanning result occupies 4 KB.

A.6.1.2 Configuration Command

The commands use the PRP1, PRP2, and Dword10 to Dword14 fields. Enter reserved value 0 for fields that are not involved.

Table A-57 PRP1

Bit	Description
63:00	**PRP Entry 1 (PRP1):** indicates the data buffer. This buffer cannot have more than one physical discontinuity. The minimum unit is 4 KB.

Table A-58 PRP2

Bit	Description
63:00	**PRP Entry 2 (PRP2):** indicates the second PRP entry. If PRP Entry 1 specifies enough space for the data structure, this field is reserved. Otherwise, it specifies the remainder of the data buffer. This field shall not be a pointer to a PRP List If the data structure requires two PRP entries, this field represents the second PRP entry address. If data of more than two PRP entries is transferred, this field points to a PRP list.

Table A-59 Command Dword 10

Bit	Description
31:00	**Data_Len:** indicates the data length in unit of Dword. If data transfer is not needed, the value is 0h.

Table A-60 Command Dword 12

Bit	Description
31:00	**Sub_OpCode:** subcommand number

A.6.1.2.1 Deleting SMART Information

This command is used to delete SMART information, alarm information, and other information. It is recommended that users do not use this command.

This command uses the Dword12 field. Enter reserved value 0 for other fields that are not involved.

Table A-61 Sub_OpCode

Sub_OpCode	Description
0x170000h	SMART information (including information defined by the protocol and vendor) and protocol error logs.

A.6.1.2.2 Configuring the SSD Capacity

This command is used to configure the available capacity of the entire SSD.

Table A-62 Sub_OpCode

Sub_OpCode	Description
0xA00000h	Configure the SSD capacity.

This command uses the Dword13 and Dword14 fields. Enter reserved value 0 for fields that are not involved.

If the configured capacity is not aligned by 4 KB, the value is rounded up for alignment.

Table A-63 Dword13&Dword14

Bit	Description
63:00	Capacity of the entire SSD (unit: byte)

A.6.1.2.3 Recovering Data

This command is used to try to recover the SSD from error status or restore data in error status. This command uses the Dword12 and Dword13 fields. Enter reserved value 0 for other fields that are not involved.

Table A-64 Sub_OpCode

Sub_OpCode	Description
0x230000h	Recover data.

Table A-65 Command Dword13

Bit	Description
31:00	0: Normal recovery (Restore data as much as possible and make the SSD able to process read commands, but the SSD will still be in error status.) 1: Forcible recovery (Recover the SSD from error status and erase all data.) Other: Reserved

Before recovering data, run the **hioadm info -d** *devicename* **-a** command and verify that **data status** in the command output is **need recovery**.

The may take up to 15 minutes to restore 1 TB of data.

A.6.1.2.4 Configuring the Number of Namespaces

This command is used to configure the number of namespaces (NN for short).

Table A-66 Sub_OpCode

Sub_OpCode	Description
0x010000h	Configure the NN Value.

This command involves controller-based access and can be run only in PF mode. This command uses the Dword13 field to indicate the NN value to be set. The NN value ranges from 1 to 64.

Enter reserved value 0 for fields that are not involved.

Table A-67 Command Dword13

Bit	Description
31:00	**NN Value**: NN value to be set, ranging from 1 to 64.

A.6.1.2.5 Setting SSD Capacities in the Unit of LBA

This command is used to configure the available capacity (unit: LBA) of an entire SSD. If the capacity of an SSD is set to 0, the SSD default capacity before delivery will be used.

Table A-68 Sub_OpCode

Sub_OpCode	Description
0x00000000h	Set disk capacities in the unit of LBA.

This command uses the Dword13 and Dword14 fields. Enter reserved value 0 for other fields that are not involved.

Table A-69 Dword13&Dword14

Bit	Description
63:00	LBA quantity

NOTE

The process for configuring the disk capacity in the unit of LBA is as follows:

1. Send the command to set drive capacities.

2. After the command is executed successfully, run the **fdisk -l** command to check whether the drive capacity changes. If the drive capacity does not change, restart the system or uninstall and reload the driver.

A.6.2 Legacy Command (Deprecated)

In addition to the commands defined in protocols, the ES3000 V5 supports the commands defined by vendors to provide extensive device management functions. The interfaces described in this section are historical versions and provide the same functions as the **New Commands**. These interfaces are reserved for compatibility with earlier versions. For ES3000 V5 and later SSDs, the interfaces described in **New Commands** are recommended.

A.6.2.1 Query Command

This command uses the PRP1, PRP2, and Dword10 to Dword15 fields. Enter reserved value 0 for other fields that are not involved.

Table A-70 PRP1

Bit	Description
63:00	**PRP Entry 1 (PRP1):** indicates the storage of data buffer. This buffer can be configured only in a physical discontinuity. The minimum unit is 4 KB.

Table A-71 PRP2

Bit	Description
63:00	**PRP Entry 2 (PRP2):** indicates the second PRP entry. If data transmission needs only one PRP entry, then reserve the field. If data transmission has two PRP entry, then this field represents the second PRP entry address.

Table A-72 Command Dword 10

Bit	Description
31:00	**Data_Len:** The length of data transmission is in the unit of Dword. If data transmission is not needed, then the value is 0 h.

Table A-73 Command Dword 12

Bit	Description
31:16	**Sub_OpType:** Subcommand Processing Type
15:00	**Sub_OpCode:** Subcommand Number

This command currently supports the custom Sub_Opcode as follows:

Table A-74 Sub_Opcode

Sub_OpCode	Description
01h	Obtain logs.
02h	Obtain SSD ID parameters.
03h	Obtain the function enablement status.
20h	Query the maximum available capacity of the SSD.
21h	Obtain the vendor-defined log length.
25h	Query the PCB version.
29h	Obtain the self-diagnosis result.
2Ah	Obtain the chipset DIE ID.
57h	Obtain the average DWPD value.
59h	Obtain the unique ID.
157h	Obtain the pending list.

A.6.2.1.1 Obtaining Logs and Statistical Information

This command is used to obtain vendor logs and statistics information from the SSD.

Each command sent from the host can obtain a maximum of 4KB of logs. The host obtains the total log size through the vendor log length query command (SubOpCode 0x21h) and then sends commands repeatedly to obtain all logs. Besides the total data size, the commands must also specify an offset value, where 1 indicates 4 KB.

This command uses the PRP1, PRP2, Dword10, Dword12, and Dword13 fields. Enter reserved value 0 for other fields that are not involved

Sub_OpType indicates the log type.

Table A-75 Sub_OpType

Sub_OpType	Description
0h	Log parsing package
1h	All logs
2h	Key logs
3h	CoreDump information
4h	Statistics information

This command uses the command domains of PRP1, PRP2, Dword10, and Dword13.

Table A-76 Command Dword 10

Bit	Description
31:00	**Data_Len:** The length of data transmission is in the unit of Dword. If the length of the package is not 4-byte aligned, then round up the result. The actual returned length will be shown in the Dword0 of CQ.

Table A-77 Command Dword 13

Bit	Description
31:00	**Offset:** Log offset. 1 indicates 4 KB.

The valid log length and offset are returned to the host through Dword0 of the CQ.

Table A-78 Completion Queue Entry Dword 0

Bit	Description
31:16	**Log_Length:** Obtained log length
15:00	**Offset:** obtained log offset value.

NOTE

Each command sent from the host can obtain a maximum of 4KB of logs. If the size exceeds 4KB, the firmware returns the status code "NVME_SC_INVALID_FIELD = 0x2".

A.6.2.1.2 Obtaining SSD ID Parameters

This command is used to obtain SSD ID parameters, including electronic label information, part number (PN), World Wide Name (WWN), Model, and serial number (SN). The parameters are returned to the host through PRP1 and PRP2.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table1 lists parameter types.

Table A-79 Sub_OpType

Sub_OpType	Description
00h	Electronic label information
01h	PN
02h	WWN
03h	Model
04h	SN
05h	Flash ID

- **Electronic label information**

The subcommand is used to set the electronic label of the SSD. The host uses PRP1 and PRP2 to send electronic label information in the format described in **Table 2 Electronic label information**.

In the following table, "Field" indicates the combination of the field name, equal sign, field value, and carriage return. For example, "VendorName=Huawei\r\n" is a string, where "VendorName" is the field name and "Huawei" is the field value.

Table A-80 Electronic label information

Area	Description	Maximum Field Length (Bytes)	Maximum Field Value Length (Bytes)	Remarks
Information format area	/$ [ArchivesInfo Version]\r\n	26	N/A	If a line starts with /$, the line is not displayed only in the network management system. $ is an ASCII character with the character code 0x24. The data needs to

Area	Description	Maximum Field Length (Bytes)	Maximum Field Value Length (Bytes)	Remarks
	/$ArchivesInfoVersion=<*Version Number*>\r\n	30	6	be displayed during query on network elements. The current label version is 3.0, which indicates the number of fields. Usually content changes do not cause version changes.
	\r\n	2	N/A	Two blank lines are between the information format area and the information description area.
	\r\n	2	N/A	
Information description area	[Board Properties]\r\n	20	N/A	
	BoardType=<*Model*>\r\n	30	18	Product model number
	BarCode=<*SN*>\r\n	30	20	SN on the paper label
	Item=<*BBOM Code*>\r\n	30	23	PN on the paper label
	Description=<*English Description*>\r\n	260	246	The line contains a maximum of 260 bytes.
	Manufactured=<*Manufacturing Date*>\r\n	25	10	Format: YYYY-MM-DD.
	VendorName=<*Manufacturer Name*>\r\n	50	37	FRUs made by Huawei or Huawei's outsourced vendors are all branded by Huawei.
	IssueNumber=<*Issue Number*>\r\n	20	6	Item revision (FRU hardware version) on the paper label
	CLEICode=<*CLEI Code*>\r\n	25	14	This field is used by products complying with specifications of Telcordia in North America.

Area	Description	Maximum Field Length (Bytes)	Maximum Field Value Length (Bytes)	Remarks
	BOM=<*BOM Code*>\r\n	50	44	Item of a finer granularity. This field is reserved for products whose Huawei item codes are insufficient for customers' asset management.

- **PN**

The subcommand is used to set the SSD part number (PN). The host uses PRP1 and PRP2 to send PN information in the format described in **Table3**.

Table A-81 PN

Bytes	Description
23:00	SN

- **WWN**

The subcommand is used to set the SSD WWN. The host uses PRP1 and PRP2 to send WWN information in the format described in **Table4 WWN**.

Table A-82 WWN

Bytes	Description
15:00	WWN

- **Model**

The subcommand is used to set the SSD model number. The host uses PRP1 and PRP2 to send model information in the format described in **Table5 Model**.

Table A-83 Model

Bytes	Description
15:00	Model

- **SN**

The subcommand is used to set the SSD SN. The host uses PRP1 and PRP2 to send SN information in the format described in **Table6 SN**.

Table A-84 SN

Bytes	Description
19:00	SN

- **FLASH ID**

This command is used to obtain the FLASH ID. This command uses PRP1, PRP2, and Dword10. The data length is specified by Dword10 (fixed at 2).

Table A-85 FLASH ID

BIT	Description
63:00	FLASH ID

A.6.2.1.3 Obtaining the Enablement Status of Functions

This command is used to obtain the enablement status of a customized function through the controller. The PF and VF have the permission to run this command based on controller access.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Sub_OpType indicates the function type.

Table A-86 Sub_OpType

Sub_OpType	Description
01h	Flash inspection enablement
02h	Capacitor inspection enablement
03h	DDR inspection enablement
04h	Temperature inspection enablement
06h	CE concurrency enablement
07h	Read Retry enablement
08h	Wear leveling enablement
09h	Inter-chip RAID enablement
0ah	Garbage collection (GC) enablement

Sub_OpType	Description
0bh	Aging test enablement

The enablement status of each function is returned through CQ Dword0.

Table A-87 CQ-Dword0

Bit	Description
00	**Enable Status**: **1** indicates that the function is enabled. **0** indicates that the function is disabled.
31:01	Reserved

A.6.2.1.4 Querying the Maximum Available Capacity of the SSD

This command is used to query the maximum available capacity that can be set. This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for fields that are not involved. The command returns the information to the host through PRP1 and PRP2. The format of the returned data is as follows:

Table A-88 Querying the maximum available capacity of the SS

Bytes	Description
7:0	Maximum available capacity of the SSD (unit: byte)

A.6.2.1.5 Obtaining the Vendor Log Length

This command is used to obtain the length of vendor logs (including the log parsing package, full logs, and key logs).

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Sub_OpType indicates the log type.

Table A-89 Sub_OpType

Sub_OpType	Description
0h	Log parsing package
1h	Full logs
2h	Key logs

Sub_OpType	Description
3h	CoreDump information
4h	Statistics

The obtained valid log length is returned to the host software as Dword0 in the CQ.

Table A-90 Completion Queue Entry Dword 0

Bit	Description
31:0	Valid log length (unit: byte)

A.6.2.1.6 Querying the PCB Version

This command is used to obtain the PCB version, which is returned to the host as Dword0 in the CQ.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for fields that are not involved.

Table A-91 Dword0 in the CQ

Bit	Description
31:0	PCB version

A.6.2.1.7 Obtaining Self-Diagnosis Results

If a bit value is **1**, the corresponding item has a fault. If a bit value is **0**, the corresponding item is OK.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for other fields that are not involved.

Table A-92 CQ-Dword0

Bit	Description
0	Overall test result
1	Firmware error test (front-end error check: whether the drive is in the controller drive mode, whether the drive is in read-only mode, and whether the available space is below the threshold)
2	NAND test
3	NOR connectivity test

Bit	Description
4	DDR uncorrectable error test
5	IBUF uncorrectable error test
6	Capacitor test
7	Blank block test: insufficient blank blocks
8	Life test
9	NOR patrol test
10	Firmware uncorrectable error test: four invalid firmware backups
11	System area test too many bad blocks in the NAND system area
12	Bad block test: too many bad blocks in the NAND data area
13	Controller test: invalid controller
14	Power down fast interrupt request (FIQ) test
15	Formatting failure test
16	DB poweron fail test
17	Reserved
18	Lbn miss happen test
19	Deep rebuild happen Test
24:20	Reserved
25	**1**: Firmware needs to be upgraded to restore SSD disks.
26	**1**: Data loss risks exist.
27	**1**: SSD disks can be restored, but risks exist.
28	**1**: SSD disks can be restored.
29	**1**: Forcibly restore SSD disks so that the SSD disks can function properly.
30	**1**: Restart or power on and then off SSD disks to restore SSD disks.
31	**1**: The fault cannot be rectified. Contact Huawei support engineers.

A.6.2.1.8 Obtaining the DIEID of a Chip

This command is used to obtain the DIEID, which is the unique identifier of a chip and is used to track and manage information of the entire lifecycle (including chip manufacturing, chip test, board processing, server test, and live network operation).

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for fields that are not involved.

Test results are returned to the host through the host's physical memory that PRP1 and PRP2 point to. The data structure is defined in the following table.

Table A-93 DIEID structure

Bytes	Description
3:0	Chip ID
7:4	Version
11:8	DIEID[31:0], indicates the first 4 bytes of DIEID.
…	-
27:24	DIEID[160:129], indicates the last 4 bytes of DIEID.
31:28	Reserved

A.6.2.1.9 Querying the Average DWPD

This interface is used to query the average DWPD in N (N <= 360) days. Dword13 indicates the number of days, and the result is returned by Dword0 of CQ.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for fields that are not involved.

Table A-94 Command Dword13

Bit	Description
31:00	**Number of days**: valid value (1-360)

Table A-95 Cq Command Dword0

Bit	Description
31:00	**DWPD value**: 100 indicates 1 (If the average daily written data volume is 0.8 TB and the drive user volume is 1.6 TB, the DWPD value is 0.5 and the interface result is 50).

A.6.2.1.10 Querying UniqueID Information

This interface is used to obtain the UniqueID information. The host issues a command for obtaining the data length, and then issues a command for obtaining the content based on the length. The content obtained is in the string format and the length does not contain '\0' at the end.

This command uses the PRP1, PRP2, Dword10, and Dword12 fields. Enter reserved value 0 for fields that are not involved.

Table A-96 SubOpType

SubOpType	Description
0h	Length of data to be obtained.
1h	Content to be obtained.

When SubOpType delivered to obtain the uniqueID length is 0, the length is returned by DWORD0 of CQ.

Table A-97 Cq Command Dword0

Bit	Description
31:00	**Length:** The unit is byte.

When SubOpType delivered to obtain the uniqueID content is 1, the length is returned by Dword10 and must be the same as the length obtained previously. The UniqueID content is transmitted by the memory that PRP points to, and the content is in the string format.

Table A-98 Command Dword10

Bit	Description
31:00	**Length:** The unit is Dword.

A.6.2.2 Configuration Commands

This command uses the PRP1, PRP2, and Dword10 to Dword15 fields. Enter reserved value 0 for other fields that are not involved.

Table A-99 PRP1

Bit	Description
63:00	**PRP Entry 1 (PRP1):** indicates the data buffer. This buffer cannot have more than one physical discontinuity. The minimum unit is 4 KB.

Table A-100 PRP2

Bit	Description
63:00	**PRP Entry 2 (PRP2):** indicates the second PRP entry. If PRP Entry 1 specifies enough space for the data structure, this field is reserved. Otherwise, it specifies the remainder of the data buffer. This field shall not be a pointer to a PRP List If the data structure requires two PRP entries, this field represents the second PRP entry address. If data of more than two PRP entries is transferred, this field points to a PRP list.

Table A-101 Command Dword 10

Bit	Description
31:00	**Data_Len:** indicates the data length in unit of Dword. If data transfer is not needed, the value is 0h.

Table A-102 Command Dword 12

Bit	Description
31:16	**Sub_OpType:** subcommand operation type
15:00	**Sub_OpCode:** subcommand number

This command currently supports the following Sub_Opcodes:

Table A-103 Sub_Opcode

Sub_OpCode	Description
01h	Delete vendor-defined logs and information.
02h	Configure the SSD capacity.
1Ah	Recover data
1Dh	NN value

A.6.2.2.1 Deleting SMART Information

This command is used to delete SMART information, alarm information, and other information. It is recommended that users do not use this command.

This command uses the Dword12 field. Enter reserved value 0 for other fields that are not involved.

Sub_OpType indicates the type of data to be deleted.

Table A-104 Sub_OpType

Sub_OpType	Description
00h	Vendor logs
01h	SMART information (including information defined by the protocol and vendor) and protocol error logs.
02h	All information, including vendor logs, SMART information, and protocol error logs.

A.6.2.2.2 Configuring the SSD Capacity

This command is used to configure the available capacity of the entire SSD. The command uses the Dword13 and Dword14 fields. Enter reserved value 0 for fields that are not involved. The tool on the host needs to ensure that the entire drive capacity is 4K-aligned.

Table A-105 Dword13&Dword14

Bit	Description
63:00	Capacity of the entire SSD (unit: byte)

A.6.2.2.3 Recovering Data

Run this command to recover data of the SSD. The recovery may take up to 20 minutes. Before recovering data, run the **hioadm info -d** *devicename* **-a** command and verify that **data status** in the command output is **need recovery**.

This command uses the Dword12 field. Enter reserved value 0 for fields that are not involved.

A.6.2.2.4 Configuring the Number of Namespaces

This command is used to configure the number of namespaces (NN for short), which corresponds to the Identify Controller NN field (Byte 519:516). This command involves controller-based access and can be run only in PF mode. This command uses the Dword13 field to indicate the NN value to be set. The NN value ranges from 1 to 64. Enter reserved value 0 for fields that are not involved.

Table A-106 Command Dword13

Bit	Description
31:00	**NN Value**: NN value to be set, ranging from 1 to 64.

B Out-of-Band Management

This section describes the out-of-band management functions of the ES3000 V5.

The ES3000 V5 provides the following out-of-band management functions:

- NVMe basic management commands
- VPD information
- Temperature sensors

NOTE

To use out-of-band management functions, the server must support them. To check whether a server supports out-of-band management, contact the server vendor.

B.1 NVMe Basic Management Command

The I2C address of NVMe out-of-band management commands is 0x6A (excluding the R/W bits).

Table B-1 NVMe basic management data structure

Command Code	Offset (Byte)	Description
0	00	**Length of Status**: Indicates number of additional bytes to read before encountering PEC. This value should always be 6 (06h) in implementations of this version of the spec.

Command Code	Offset (Byte)	Description
	01	**Status Flags (SFLGS)**: This field indicates the status of the NVM subsystem. **SMBus Arbitration**: Bit 7 is set **1** after a SMBus block read is complete all the way to the stop bit without bus contention and cleared to **0** if a SMBus Send Byte FFh is received on this SMBus slave address. **Drive Not Ready**: Bit 6 is set to **1** when the subsystem cannot process NVMe management commands, and the rest of the transmission may be invalid. If cleared to **0**, then the NVM subsystem is fully powered and ready to respond to management commands. This logic level intentionally identifies and prioritizes powered up and ready drives over their powered off neighbors on the same SMBus segment. **Drive Functional**: Bit 5 is set to **1** to indicate an NVM subsystem is functional. If cleared to **0**, then there is an unrecoverable failure in the NVM subsystem and the rest of the transmission may be invalid. **Reset Not Required**: Bit 4 is set to **1** to indicate the NVM subsystem does not need a reset to resume normal operation. If cleared to **0** then the NVM subsystem has experienced an error that prevents continued normal operation. A controller reset is required to resume normal operation. **Port 0 PCIe Link Active**: Bit 3 is set to **1** to indicate the first port's PCIe link is up (i.e., the Data Link Control and Management State Machine is in the DL_Active state). If cleared to **0**, then the PCIe link is down. **Port 1 PCIe Link Active**: Bit 2 is set to **1** to indicate the second port's PCIe link is up. If cleared to **0**, then the second port's PCIe link is down or not present. Bits 1 and 0 must be set to **1**.
	02	**SMART Warnings**: This field shall contain the Critical Warning field (byte 0) of the NVMe SMART / Health Information log. Each bit in this field shall be inverted from the NVMe definition (i.e., the management interface shall indicate a **0** value while the corresponding bit is **1** in the log page). Refer to the NVMe specification for bit definitions. If there are multiple controllers in the NVM subsystem, the management endpoint shall combine the Critical Warning field from every controller such that a bit in this field is: ● Cleared to **0** if any controller in the subsystem indicates a critical warning for that corresponding bit. ● Set to **1** if all controllers in the NVM subsystem do not indicate a critical warning for the corresponding bit.

<table>
<tr><th>Command Code</th><th>Offset (Byte)</th><th colspan="2">Description</th></tr>
<tr><td rowspan="11"></td><td>03</td><td colspan="2">Composite Temperature (CTemp): This field indicates the current temperature in degrees Celsius. If a temperature value is reported, it should be the same temperature as the Composite Temperature from the SMART log of hottest controller in the NVM subsystem. The reported temperature range is vendor specific, and shall not exceed the range －60°C to +127°C (－76°F to +260.6°F).</td></tr>
<tr><td rowspan="9"></td><td>Value</td><td>Detailed Description</td></tr>
<tr><td>00h-7Eh</td><td>Temperature is measured in degrees Celsius: 0°C to 126°C (32°F to 258.8°F)</td></tr>
<tr><td>7Fh</td><td>127°C (260.6°F) or higher</td></tr>
<tr><td>80h</td><td>No temperature data or temperature data is more the 5 seconds old.</td></tr>
<tr><td>81h</td><td>Temperature sensor failure</td></tr>
<tr><td>82h-C3h</td><td>Reserved</td></tr>
<tr><td>C4</td><td>The temperature is －60°C (－76°F) or lower.</td></tr>
<tr><td>C5-FFh</td><td>Temperature measured in degrees Celsius is represented in twos complement: －59°C to －1°C (－74.2°F to 30.2°F)</td></tr>
<tr><td>04</td><td colspan="2">Percentage Drive Life Used (PDLU): Contains a vendor specific estimate of the percentage of NVM subsystem NVM life used based on the actual usage and the manufacturer's prediction of NVM life. A value of 100 indicates that the estimated endurance of the NVM in the NVM subsystem has been consumed, but may not indicate an NVM subsystem failure. The value can be greater than 100. The value is always set to 255 if it is greater than 254.</td></tr>
<tr><td>06:05</td><td colspan="2">Reserved</td></tr>
<tr><td>07</td><td colspan="2">PEC: An 8-bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.</td></tr>
<tr><td>8</td><td>08</td><td colspan="2">Length of identification: Indicates number of additional bytes to read before encountering PEC. This value should always be 22 (16h) in implementations of this version of the spec.</td></tr>
</table>

Command Code	Offset (Byte)	Description
	10:09	**Vendor ID**: The 2-byte vendor ID, assigned by the PCI SIG. Should match VID in the Identify Controller command response. MSB is transmitted first.
	30:11	**Serial Number**: 20 characters that match the serial number in the NVMe Identify Controller command response. First character is transmitted first.
	31	**PEC**: An 8-bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.
32	255:32	**Vendor Specific**: Vendor-specific data.

Table B-2 Huawei-specific data structure

Command Code	Offset (Byte)	Description
32	32	Data byte length defined by the vendor before PEC. This domain must be set to 22(16h).
	33	Reserved byte. The default value is 0.
	34	Current power consumption. Unit: W
	35	Reserved byte. The default value is 0x80.
	36	Reserved byte. The default value is 0.
	37	Reserved byte. The default value is 0.
	38	Reserved byte. The default value is 0.
	46:39	**Firmware Version**: 8-byte firmware version number in ASCII code.
	54:47	**Bootloader Version**: 8-byte bootloader version number in ASCII code.
	55	**PEC**: An 8-bit CRC calculated over the slave address, command code, second slave address and returned data. Algorithm is in SMBus Specifications.
56	56	This field indicates the length of the data bytes customized by the vendor before PEC. This field is set to 30 (1Eh).
	64:57	**Detailed alarm 1. The bit definition is the same as alarm 1 in Table B-1.**
	72:65	**Detailed alarm 2. The bit definition is the same as alarm 2 in Table B-1.**

Command Code	Offset (Byte)	Description
	86:73	Reserved. The default value is **0**.
	87	**PEC**: 8-bit CRC value calculated based on the write address of the slave device, Command Code, read address of the slave device, and returned data.

B.2 VPD

The I2C address of vital product data (VPD) is 0x53 (excluding R/W bits).

The IIC address (with R bits) of the slave device is 0xa6. The information is for read only. Do not modify it with write commands, otherwise, it cannot be restored.

Perform only read operations on the VPD EEPROM. Write operations may destroy VPD data. If you must perform write operations, use the byte mode. If you must write data in page mode, write data of less than 16 bytes from the beginning of each page to prevent overwriting data.

The ES3000 V5 VPD information meets the requirements of the NVME-MI1.0 protocol.

Table B-3 VPD definition

Byte	Name
07:00	Common Header
119:08	Product Info Area
Vendor Specific:120	MultiRecord Info Area
Vendor Specific	Internal Use Area: not supported
Vendor Specific	Internal Use Area: not supported
Vendor Specific	Internal Use Area: not supported

Table B-4 Common Header

Byte	Factory Default	Description
0	01h	IPMI Format Version Number (IPMIVER): This field indicates the IPMI Format Version
1	00h	Internal Use Area Starting Offset (IUAOFF): This field indicates the starting offset in multiples of 8 bytes for the Internal Use Area. A value of 00h may be used to indicate the Internal Use Area is not present

Byte	Factory Default	Description
2	00h	Chassis Info Area Starting Offset (CIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Chassis Info Area. A value of 00h may be used to indicate the Chassis Info Area is not present
3	00h	Board Info Area Starting Offset (BIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Board Info Area. A value of 00h may be used to indicate the Board Info Area is not present
4	01h	Product Info Area Starting Offset (PIAOFF): This field indicates the starting offset in multiples of 8 bytes for the Product Info Area
5	0Fh	MultiRecord Info Area Starting Offset (MRIOFF): This field indicates the starting offset in multiples of 8 bytes for the MultiRecord Info Area
6	00h	Reserved
7	Impl Spec	Common Header Checksum (CHCHK): Checksum computed over bytes 0 through 6. The checksum is computed by adding the 8-bit value of the bytes modulo 256 and then taking the 2's complement of this sum. When the checksum and the sum of the bytes module 256 are added, the result should be 0h

Table B-5 NVMe MultiRecordArea

Byte	Factory Default	Description
00	0Bh	NVMe Record Type ID
01	2h	Bit 7 - end of list; record format version = 2
02	28h	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.
03	Impl Spec	Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 through the end of this record plus this checksum byte equals zero).
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero.
05	00h	NVMe MultiRecord Area Version Number: This field indicates the version number of this multirecord. This field shall be set to 0h in this version of the specification.

Byte	Factory Default	Description	
06	11h	Management Endpoint Form Factor (MEFF): This field indicates the form factor of the Management Endpoint.	
		Value	Definition
		0	Other - unknown
		1-15	Reserved
		16	2.5" Form Factor -unknown
		17	2.5" Form Factor - U.2 (SFF-8639) 15 mm
		18	2.5" Form Factor - U.2 (SFF-8639) 7 mm
		19 – 31	Reserved
		32	CEM add in card - unknown
		33	CEM add in card - Low Profile (HHHL)
		34	CEM add in card - Standard Height Half Length (FHHL)
		35	CEM add in card - Standard Height Full Length (FHFL)
		36-47	Reserved
		48	M.2 module - unknown
		49	M.2 module - 2230
		50	M.2 module - 2242
		51	M.2 module - 2260
		52	M.2 module - 2280
		53	M.2 module - 22110
		54-63	Reserved
		64	BGA SSD - unknown
		65-239	Reserved
		240-255	Vendor Specific
12:07	00h	Reserved	
13	00h	Initial 1.8V Power Supply Requirements: This field specifies the initial 1.8V power supply requirements in Watts prior to receiving a Set Slot Power message.	

Byte	Factory Default	Description
14	00h	Maximum 1.8V Power Supply Requirements: This field specifies the maximum 1.8V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
15	00h	Initial 3.3V Power Supply Requirements: This field specifies the initial 3.3V power supply requirements in Watts prior to receiving a Set Slot Power message.
16	00h	Maximum 3.3V Power Supply Requirements: This field specifies the maximum 3.3V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
17	00h	Reserved
18	0Ah	Maximum 3.3V aux Power Supply Requirements: This field specifies the maximum 3.3V power supply requirements in 10 mW units. A value of zero indicates that the power supply voltage is not used.
19	00h	Initial 5V Power Supply Requirements: This field specifies the initial 5V power supply requirements in Watts prior to receiving a Set Slot Power message.
20	00h	Maximum 5V Power Supply Requirements: This field specifies the maximum 5V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
21	19h	Initial 12V Power Supply Requirements: This field specifies the initial 12V power supply requirements in Watts prior to receiving a Set Slot Power message.
22	19h	Maximum 12V Power Supply Requirements: This field specifies the maximum 12V power supply requirements in Watts. A value of zero indicates that the power supply voltage is not used.
23	19h	Maximum Thermal Load: This field specifies the maximum thermal load from the NVM Subsystem in Watts.
36:24	Impl Spec	Total NVM Capacity: This field indicates the total NVM capacity of the Management Endpoint in bytes. If the NVM Subsystem supports Namespace Management, then this field should correspond to the value reported in the TNVMCAP field in the NVMe Identify Controller Data structure. A value of 0h may be used to indicate this feature is not supported.
63:37	00h	Reserved

Table B-6 NVMe PCIe PortMultiRecord Area

Byte	Factory Default	Description	
00	0Ch	NVMe PCIe Port Record Type ID	
01	82h	Bit 7 – end of list; record format version = 2h	
02	0Bh	Record Length (RLEN): This field indicates the length of the MultiRecord Area in bytes.	
03	Impl Spec	Record Checksum: This field is used to give the record data a zero checksum (i.e., the modulo 256 sum of the record data bytes from byte offset 05 through the end of this record plus this checksum byte equals zero).	
04	Impl Spec	Header Checksum: This field is used to give the record header a zero checksum (i.e., the modulo 256 sum of the preceding record bytes starting with the first byte of the header plus this checksum byte equals zero.	
05	00h	NVMe PCIe Port MultiRecord Area Version Number: This field indicates the version number of this multirecord. This field shall be set to zero in this version of the specification.	
06	00h	PCIe Port Number: This field contains the PCIe port number. This is the same value as that reported in the Port Number field in the PCIe Link Capabilities Register.	
07	01h	Port Information: This field indicates information about the PCIe Ports in the device. Bits 7 to 1 are reserved. Bit 0, if set to **1**, indicates that all PCIe ports within the device have the same capabilities (i.e., the capabilities listed in this structure are consistent across each PCIe port).	
08	07h	PCIe Link Speed: This field indicates a bit vector of link speeds supported by the PCIe port.	
		Bit	Definition
		7:3	Reserved
		2	Set to **1** if the PCIe link supports 8.0 GT/s. Otherwise cleared to **0**.
		1	Set to **1** if the PCIe link supports 5.0 GT/s. Otherwise cleared to **0**.
		0	Set to **1** if the PCIe link supports 2.5 GT/s. Otherwise cleared to **0**.

<table>
<tr><th>Byte</th><th>Factory Default</th><th colspan="2">Description</th></tr>
<tr><td rowspan="16">09</td><td rowspan="16">04h</td><td colspan="2">PCIe Maximum Link Width: The maximum PCIe link width for this NVM Subsystem port. This is the expected negotiated link width that the port link trains to if the platform supports it. A Management Controller may compare this value with the PCIe Negotiated Link Width to determine if there has been a PCIe link training issue.</td></tr>
<tr><td>Value</td><td>Definition</td></tr>
<tr><td>0</td><td>Reserved</td></tr>
<tr><td>1</td><td>PCIe x1</td></tr>
<tr><td>2</td><td>PCIe x2</td></tr>
<tr><td>3</td><td>Reserved</td></tr>
<tr><td>4</td><td>PCIe x4</td></tr>
<tr><td>5-7</td><td>Reserved</td></tr>
<tr><td>8</td><td>PCIe x8</td></tr>
<tr><td>9-11</td><td>Reserved</td></tr>
<tr><td>12</td><td>PCIe x12</td></tr>
<tr><td>13-15</td><td>Reserved</td></tr>
<tr><td>16</td><td>PCIe x16</td></tr>
<tr><td>17-31</td><td>Reserved</td></tr>
<tr><td>32</td><td>PCIe x32</td></tr>
<tr><td>33-255</td><td>Reserved</td></tr>
<tr><td>10</td><td>00h</td><td colspan="2">MCTP Support: This field contains a bit vector that specifies the level of support for the NVMe Management Interface. Bits 7 to 1 are reserved. Bit 0, if set to 1, indicates that MCTP based management commands are supported on the PCIe port.</td></tr>
<tr><td rowspan="4">11</td><td rowspan="4">01h</td><td colspan="2">Ref Clk Capability: This field contains a bit vector that specifies the PCIe clocking modes supported by the port.</td></tr>
<tr><td>Bit</td><td>Definition</td></tr>
<tr><td>7:04</td><td>Reserved</td></tr>
<tr><td>3</td><td>Set to 1 if the device automatically uses RefClk if provided and otherwise uses SRIS. Otherwise cleared to 0.</td></tr>
</table>

Byte	Factory Default	Description	
		2	Set to **1** if the PCIe link supports Separate ReClk with SSC (SRIS). Otherwise cleared to **0**.
		1	Set to **1** if the PCIe link supports Separate ReClk with no SSC (SRNS). Otherwise cleared to **0**.
		0	Set to **1** if the PCIe link supports common ReClk. Otherwise cleared to **0**.
15:12	00h	Reserved	

B.3 Temperature Register

The IC2 address of the temperature sensor is 0x1B (excluding the R/W bits). The internal register 0x05 stores the detected temperature values. In the read 16-bit data, bits 1 to 12 represent temperature values. If bit 11 represents 128ºC (262.4°F), bit 1 represents 0.125ºC (32.225°F). Convert the read values into temperature values by referring to **Table B-7**.

Table B-7 Read data format of the 0x05 register

Binary	Temperature
xxx0 0111 1101 000x	+125°C (+257°F)
xxx0 0110 0011 110x	+99.75°C (+211.55°F)
xxx0 0101 0101 000x	+85°C (+185°F)
xxx0 0010 0111 000x	+39°C (+102.2°F)
xxx0 0000 1111 110x	+15.75°C (+60.35°F)
xxx0 0000 0000 010x	+0.25°C (+32.45°F)
xxx0 0000 0000 000x	0°C (+32°F)
xxx1 1111 1111 110x	- 0.25°C (+31.55°F)
xxx1 1111 1110 000x	- 1°C (+30.2°F)
xxx1 1110 1100 000x	- 20°C (- 4°F)

C SNMP Function

C.1 Overview

This section describes the Simple Network Management Protocol (SNMP) function of the ES3000 V5.

SNMP is a request/answer-based protocol used to transmit management information between an SNMP agent and an SNMP client.

Net-SNMP is an open source code SNMP software that supports SNMP v1, SNMPv2c, SNMP v3. SNMP v3 is recommended due to its high security. Net-SNMP includes the source codes of the agent and multiple management tools and supports multiple extension modes.

The SNMP service utility of Windows provides SNMP query and trap functions, but supports only SNMP v1 and v2c. SNMP v3 is not supported. Therefore, exercise caution when using the utility.

The ES3000 V5 provides the SNMP extension agent to support the SNMP function. The extension defines the device management information base (MIB). SNMP clients can query device status through the SNMP extension agent.

C.2 Configuring the SNMP Function in Windows

Preparations

Before configuring the SNMP function, check that:

- You have correctly installed the ES3000 V5 on the server.
- You have installed the SNMP service utility in Windows. For details, see **https://msdn.microsoft.com/en-us/library/cc759570(v=ws.10).aspx**.
- You have uploaded the ES3000 V5 software package to a directory on the server, for example, **D:**\.

Installing the Software

Step 1 Log in to the OS as an administrator.

Step 2 Access the software package directory, such as the **D:** directory.

Step 3 Install the Windows driver and tool package. For details, see **3.6.2 Installing the ES3000 V5 Driver in Windows** and **3.7.2 Installing the Tool Package (Windows)**.

The SNMP module is included in the tool package. Installing the tool package also installs the SNMP module.

----End

Uninstalling the Software

Step 1 Log in to the OS as an administrator.

Step 2 Uninstall the Windows tool package. For details, see **5.3.2 Uninstalling the Tool Package in Windows**.

----End

C.3 Configuring the SNMP Function in Linux

This section describes how to configure the SNMP function in Linux.

Preparations

Before configuring the SNMP function, check that:

- You have correctly installed the ES3000 V5 on the server.
- Net-SNMP is correctly installed in Linux and you have determined the Net-SNMP software version by running the following command:
  ```
  [root]#snmpd -v
  ```
- You have determined the directory storing the **snmpd.conf** file as follows:
  ```
  [root]#net-snmp-config --snmpconfpath
  /usr/local/share/snmp
  ```

Installing the Software

Step 1 Log in to the OS as the root user.

Step 2 Access the software package directory, such as the **/root** directory.

Step 3 Install the Linux driver and tool package. For details, see **3.6.1 Installing the ES3000 V5 Driver in Linux** and **3.7.1 Installing the Tool Package (Linux)**.

Step 4 If the **snmpd.conf** file is in **/usr/local/share/snmp/snmpd.conf**, run the following command to open the file:

vim /usr/local/share/snmp/snmpd.conf

Step 5 Add content as follows:

- If the Net-SNMP version is 5.6, add the following content to the **snmpd.conf** file:

  dlmod es3000v3 /usr/lib/hioadmsnmp_5.6.so
- If the Net-SNMP version is 5.7, add the following content to the **snmpd.conf** file:

  dlmod es3000v3 /usr/lib/hioadmsnmp_5.7.so

Step 6 Press **Esc** to exit the edit mode. Enter **:wq** and press **Enter**. Save and exit the **snmpd.conf** file.

Step 7 If SNMP v3 is enabled, stop SNMP services before running the following command:

net-snmp-config --create-snmpV3-user -ro -a (Huawei123#$) -A MD5 (V3user)

Huawei123#$ is the password. *V3user* is an SNMP user name. You can change the user name and password.

Step 8 Run the following command to start SNMP services:

service snmpd start

Step 9 If SNMP services are already running, run the following command to restart SNMP services:

service snmpd restart

----End

Uninstalling the Software

Step 1 Log in to the OS as the root user.

Step 2 Uninstall the Linux tool package. For details, see **5.3.1 Uninstalling the Tool Package in Linux**.

Step 3 Open the **snmpd.conf** file described in the installation procedure and delete the following content (assuming that the Net-SNMP version is 5.7):

dlmod es3000V3 /usr/lib/hioadmsnmp_5.7.so

----End

C.4 Querying SNMP Information

This section describes the software for querying ES3000 V5 SNMP information and where to obtain the software.

Using the Net-SNMP Command Set

You need to install net-snmp-utils before using the Net-SNMP command set to query ES3000 V5 SNMP information. You can use the Linux Net-SNMP tool package.

For details about how to use the Net-SNMP command set, see **http://www.net-snmp.org**.

Using a MIB Browser

You can use a management information base (MIB) browser to query SNMP information. Multiple MIB browsers are available. The MG-SOFT MIB browser is recommended. The browser is a Windows graphical user interface (GUI) tool with the trap receiving function.

For details about how to download and use the MG-SOFT MIB browser, see **http://software.informer.com**.

Common Problems

- SELinux may forbid Net-SNMP from accessing SSDs (**/dev/nvme***) and cause the ES3000 V5 SNMP module to function abnormally, so you need to change the SELinux rights configuration.
- The firewall may cause failure to obtain SNMP and device information within the Intranet. If the failure occurs, disable the firewall or modify the firewall settings.
- Before starting Net-SNMP, you can run the following command to check whether the snmpd.conf file is loaded properly. If the file is not installed properly, the command output reports an error.

```
#snmpd -Dread_config -H 2>&1 | grep "Reading" | sort -u
```

Example

If Net-SNMP 5.7 is used, use the **snmpd.conf** file provided by the ES3000 V5 to start the snmpd program.

```
snmpd -d -Lo -c /opt/hio/snmpd.conf
```

Use the snmpwalk command tool to query ES3000 V5 through SNMPv2c.

snmpwalk -v 2c -c public localhost enterprises.2011.2

Or query through SNMP v3.

snmpwalk -V3 -u V3user-l auth -a MD5 -A Huawei123#$ 127.0.0.1 enterprises.2011.2

The following command output is an example:

```
# snmpwalk -Os -c public -v2c 127.0.0.1 enterprises.2011.2
enterprise.2011.2.319.1.0 = STRING: "5.0.1.6"
enterprise.2011.2.319.2.0 = INTEGER: 1
enterprise.2011.2.319.3.1.1.1 = Gauge32: 1
enterprise.2011.2.319.3.1.2.1 = STRING: "nvme0"
enterprise.2011.2.319.3.1.3.1 = INTEGER: 1
enterprise.2011.2.319.3.1.4.1 = STRING: "024YVJ10J4800386"
enterprise.2011.2.319.3.1.5.1 = STRING: "HWE52P432T0L002N"
enterprise.2011.2.319.3.1.6.1 = STRING: "1530"
enterprise.2011.2.319.3.1.7.1 = STRING: "1.3"
enterprise.2011.2.319.3.1.8.1 = STRING: "healthy"
enterprise.2011.2.319.3.1.9.1 = STRING: "no warning"
enterprise.2011.2.319.3.1.10.1 = STRING: "309 degree K (36 degree C)"
enterprise.2011.2.319.3.1.11.1 = Gauge32: 2000
enterprise.2011.2.319.3.1.12.1 = Gauge32: 2000
enterprise.2011.2.319.3.1.13.1 = STRING: "286720635.2MB"
enterprise.2011.2.319.3.1.14.1 = STRING: "167814299.3MB"
enterprise.2011.2.319.3.1.15.1 = INTEGER: 100
enterprise.2011.2.319.3.1.16.1 = STRING: "1.428799%"
enterprise.2011.2.319.3.1.17.1 = INTEGER: 0
enterprise.2011.2.319.3.1.18.1 = INTEGER: 2658
enterprise.2011.2.319.3.1.19.1 = STRING: "OK"
enterprise.2011.2.319.3.1.20.1 = STRING: "25W"
enterprise.2011.2.319.3.1.21.1 = STRING: "volatile write cache Disable"
```

C.5 SNMP OID

This section describes the SNMP object identifier (OID) defined by the ES3000 V5.

The SNMP OID is a value defined in the management information base (MIB) and used to identify MIB members.

Table C-1 lists the ES3000 V5 SNMP OIDs.

Table C-1 Node OIDs

Node Name	Node OID	Description	Type	Permission
driverVersion	1.3.6.1.4.1.2011.2.319.1	Driver version	OCTET STRING	Read-only
cardCount	1.3.6.1.4.1.2011.2.319.2	Number of devices	Integer32	Read-only
deviceIndex	1.3.6.1.4.1.2011.2.319.3.1.1	Device index	Unsigned32	Read-only
deviceName	1.3.6.1.4.1.2011.2.319.3.1.2	Device name	OCTET STRING	Read-only
namespacesNumber	1.3.6.1.4.1.2011.2.319.3.1.3	Number of namespaces	Integer32	Read-only
serialNumber	1.3.6.1.4.1.2011.2.319.3.1.4	Serial number	OCTET STRING	Read-only
modelNumber	1.3.6.1.4.1.2011.2.319.3.1.5	Model	OCTET STRING	Read-only
firmwareVersion	1.3.6.1.4.1.2011.2.319.3.1.6	Firmware version number	OCTET STRING	Read-only
nvmeVersion	1.3.6.1.4.1.2011.2.319.3.1.7	NVMe version number	OCTET STRING	Read-only
deviceStatus	1.3.6.1.4.1.2011.2.319.3.1.8	Device health status	OCTET STRING	Read-only
criticalWarning	1.3.6.1.4.1.2011.2.319.3.1.9	Alarm information	OCTET STRING	Read-only
compositeTemperature	1.3.6.1.4.1.2011.2.319.3.1.10	Overall temperature	OCTET STRING	Read-only
maxCapacity	1.3.6.1.4.1.2011.2.319.3.1.11	Maximum capacity	Unsigned32	Read-only
currentCapacity	1.3.6.1.4.1.2011.2.319.3.1.12	Current capacity	Unsigned32	Read-only
dataUnitsRead	1.3.6.1.4.1.2011.2.319.3.1.13	Total read data (MB)	OCTET STRING	Read-only
dataUnitsWritten	1.3.6.1.4.1.2011.2.319.3.1.14	Total write data (MB)	OCTET STRING	Read-only
availableSpare	1.3.6.1.4.1.2011.2.319.3.1.15	Idle block rate (%)	Integer32	Read-only
badBlockRate	1.3.6.1.4.1.2011.2.319.3.1.16	Bad block rate (%)	OCTET STRING	Read-only

Node Name	Node OID	Description	Type	Permission
percentageUsed	1.3.6.1.4.1.2011.2.319.3.1.17	Life used (%)	Integer32	Read-only
powerOnHours	1.3.6.1.4.1.2011.2.319.3.1.18	Total power-on time (h)	Integer32	Read-only
capacitanceStatus	1.3.6.1.4.1.2011.2.319.3.1.19	Capacitor status	OCTET STRING	Read-only
powerState	1.3.6.1.4.1.2011.2.319.3.1.20	Power status	OCTET STRING	Read-only
writeMode	1.3.6.1.4.1.2011.2.319.3.1.21	Write mode	OCTET STRING	Read-only

NOTE

The SNMP system creates a table instance for each SSD. The table node OID is 1(iso).3(org).6(dod).1(internet).4(private).1(enterprises).2011(huawei).2 (products).319(ES3000V5).3(cardTable).

D Using the VROC Function on the ES3000 V5

Using the VROC Function on the ES3000 V5.

Overview

The Intel® Virtual RAID on CPU (VROC) is a solution provided by Intel for user RAID volume management. Employing the Intel V5 Purley CPU platform, this solution uses the Volume Management Device (VMD) interface to create and manage virtual storage pools for NVMe SSDs. The Intel VROC uses the architecture to implement NVMe RAID. The Intel VROC function can be used on the ES3000 V5 to easily expand the RAID solution with flexible drives.

Usage Restriction

To use the VROC function on the ES3000 V5, the operating system, hardware, and driver must meet the following requirements:

- Supported OSs:
  - Windows Server 7 SP2, Windows Server 10, Windows Server 2012 R2, and Windows Server 2016
  - RHEL 7.3 or later
  - SLES 12 SP3 or later
- Hardware
  - V5 series CPUs are required.
  - To use the VROC function, purchase the Intel VROC KEY and install it on the server. For details about how to purchase and install the VROC KEY, see the Intel official website.
- Driver

  If the VROC is used on Windows, install the RSTe suite provided by Intel, including the Intel VMD driver and NVMe driver.

  To download the suite, visit **https://downloadcenter.intel.com/** and search for VROC in the search column.

Usage

For details about how to use the VROC, see the VROC guide of Intel.

Visit **https://downloadcenter.intel.com/**, search for VROC in the search box, and download the required software package. After decompressing the software package, view the corresponding user guide.

For Linux: **RSTe_NVMe_for_Linux_SW_User_Guide_*xx*.pdf**

For Windows: **RSTe VROC Users Manual_*xx*.pdf**

E Obtaining Device Logs by Using an Open-Source Tool

This section describes how to obtain ES3000 NVMe PCIe SSD device logs by using the open-source tool nvme-cli.

Function

You can use the python script **hiolog.py** to obtain device logs without using hioadm. The **hiolog.py** script can be obtained from the driver installation package.

Format

python hiolog.py -d *device*

Parameters

Parameter	Description	Value
device	Name of a device	Examples: nvme0 and nvme1

Usage Guidelines

Run this command under the /opt/hio/ directory. Generated logs will be saved in this directory. The structure of the /opt/hio/ directory is as follows:

```
.
        |-- /opt/hio/
            |-- 20160615
                |-- nvme0_coredump_172050.log
                |-- nvme0_fw_172050.log
                |-- nvme0_fw_172050.tmp
                |-- nvme0_key_172050.log
                |-- nvme0_key_172050.tmp
```

The .tmp files contain original log data. The .log files are exported logs.

Example

Obtain the key logs of the **nvme1** SSD.

```
[root@localhost ~]# python hiolog.py -d nvme1
Reading log fw(2359)
100%
||.....................................................................................
.....................
done
Reading log key(32)
100% ||.................................
done
Reading log coredump(39)
100% ||.....................................
done
Parsing log fw
100%
||.....................................................................................
.....................
done
Parsing log key
100%
||.....................................................................................
.....................
done

===============>
 /opt/hio/20181224/nvme1_fw_164009.log
 /opt/hio/20181224/nvme1_key_164009.log
 /opt/hio/20181224/nvme1_coredump_164009.log
<===============
```


F Introduction to Windows GUI

Introduction to Windows GUI

This section describes the Windows GUI tool in the Windows tool package.

After **3.7.2 Installing the Tool Package (Windows)**, access the Windows GUI tool as follows:

- If the .NET Framework version is 2.0 or 3.5, double-click the **HioadmGUI_DotNet_2.0.exe** file in the tool installation directory.
- If the .NET Framework version is 4.0 or 4.5, double-click the **HioadmGUI_DotNet_4.0.exe** file in the tool installation directory.

BaseInfo Page

This page shows the basic information about controllers and namespaces. If the namespace starts with NVMe, an NVMe SSD disk is used, as shown in **Figure F-1**.

Figure F-1 BaseInfo page of NVMe

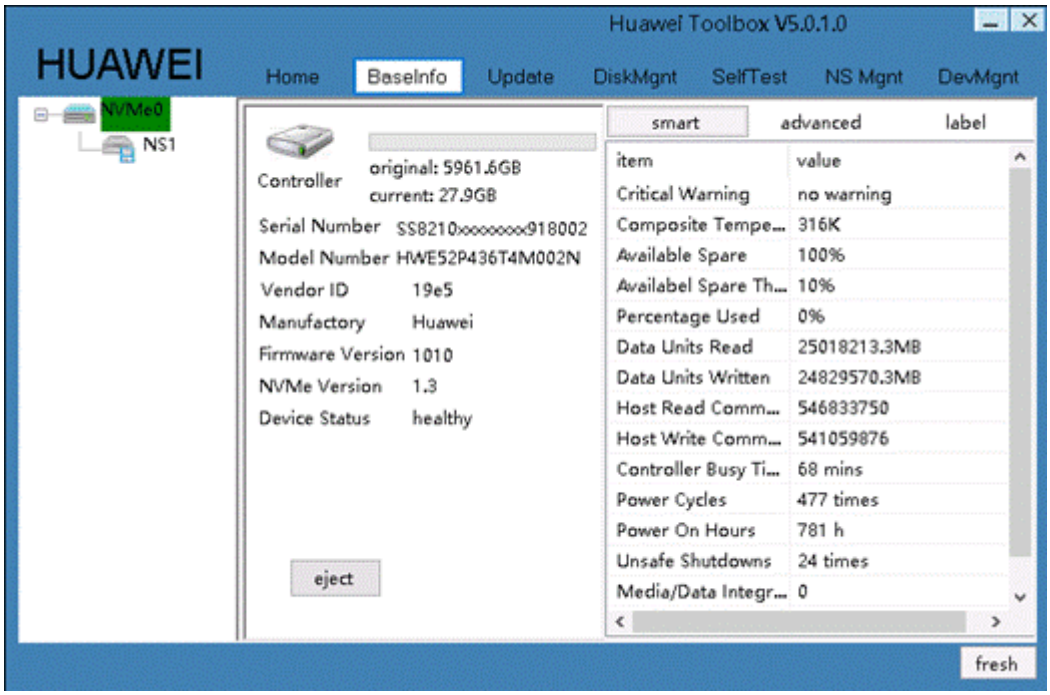

eject: Used to notify the OS to eject a device before manually removing a device.

In the left pane, **NS** indicates a namespace associated with the controller; **ns** indicates a namespace not associated with any controllers. See **Figure F-2**.

Figure F-2 BaseInfo page of NS

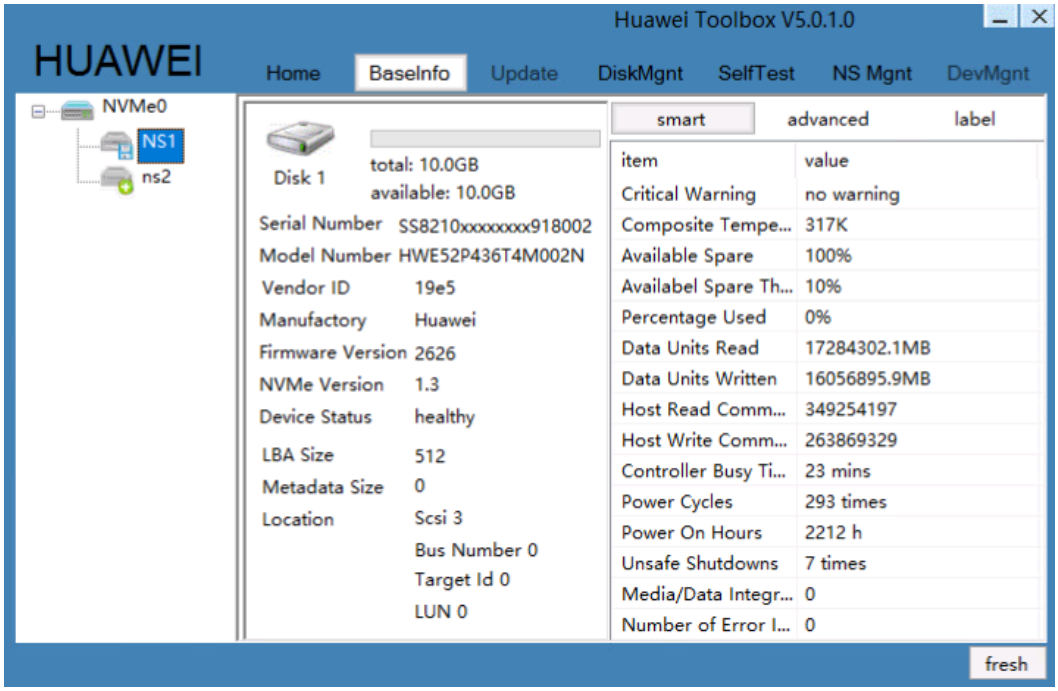

Figure F-3 BaseInfo page of ns

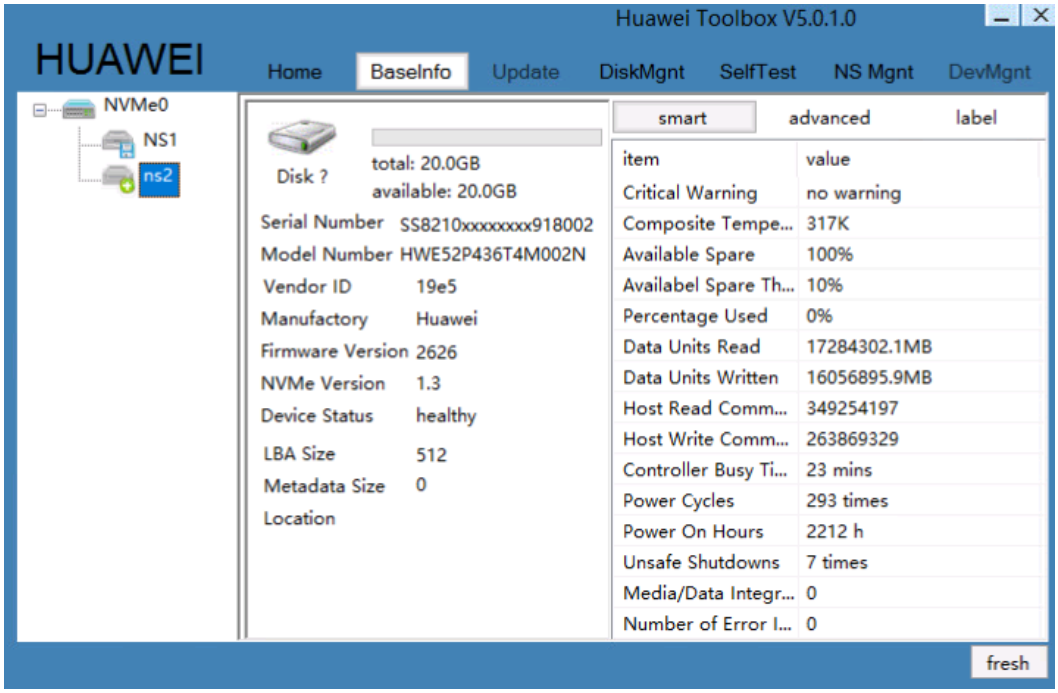

Update Page

This page is used to upgrade the ES3000 firmware, as shown in **Figure F-4** and **Figure F-5**.

Figure F-4 Update page with no firmware upgrade image selected on the ES3000

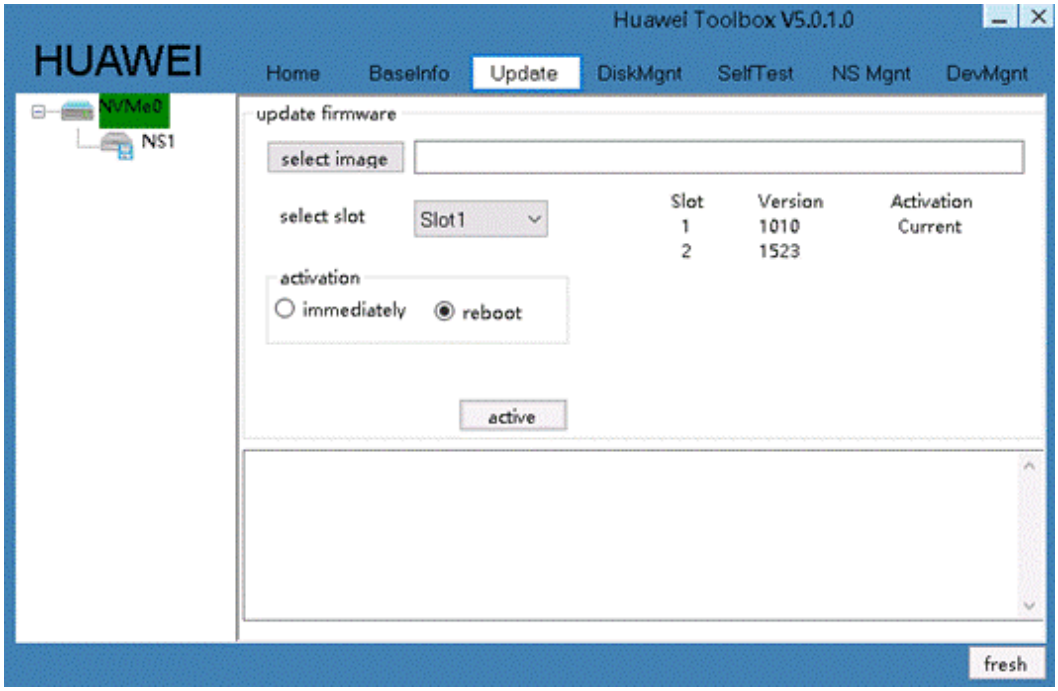

Figure F-5 Update page with a firmware upgrade image selected

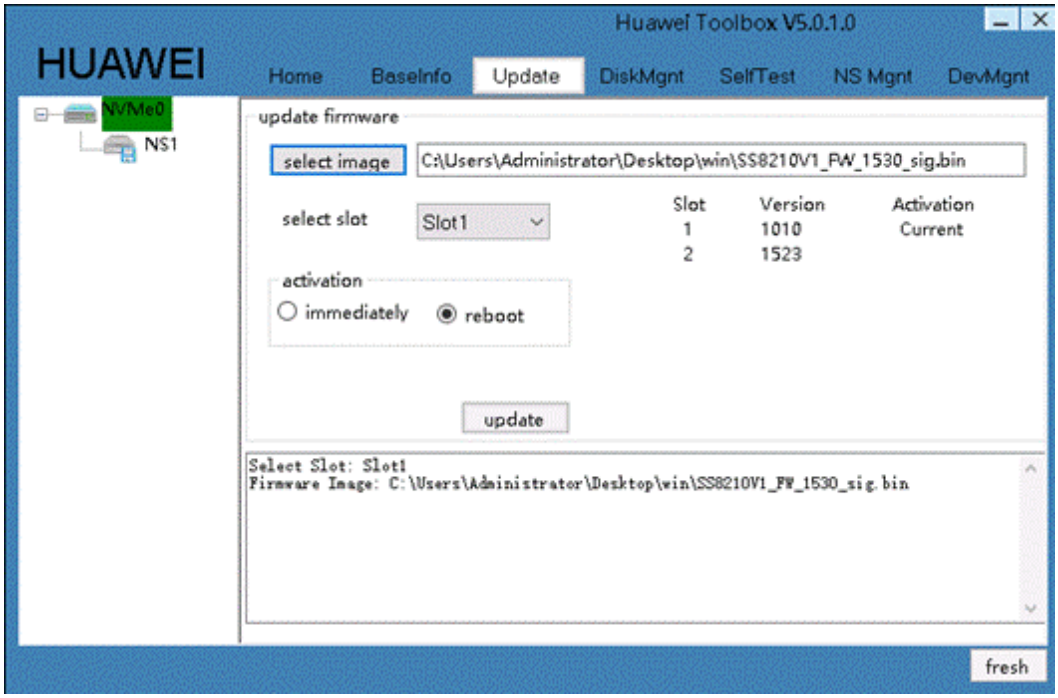

- **select image**: Select an ES3000 or SAS HSSD firmware upgrade image provided by Huawei.

- **select slot**: Only slots 1 and 2 are available for NVMe SSD.
- After a firmware image is selected, the **active** button becomes an **update** button.
- **activation**: Select an activation mode. For details, see **6.2.3 Activating the Firmware of an SSD**.

DiskMgnt Page

This page is used to format the ES3000 V5, as shown in **Figure F-6**.

Figure F-6 NVMeDiskMgnt page

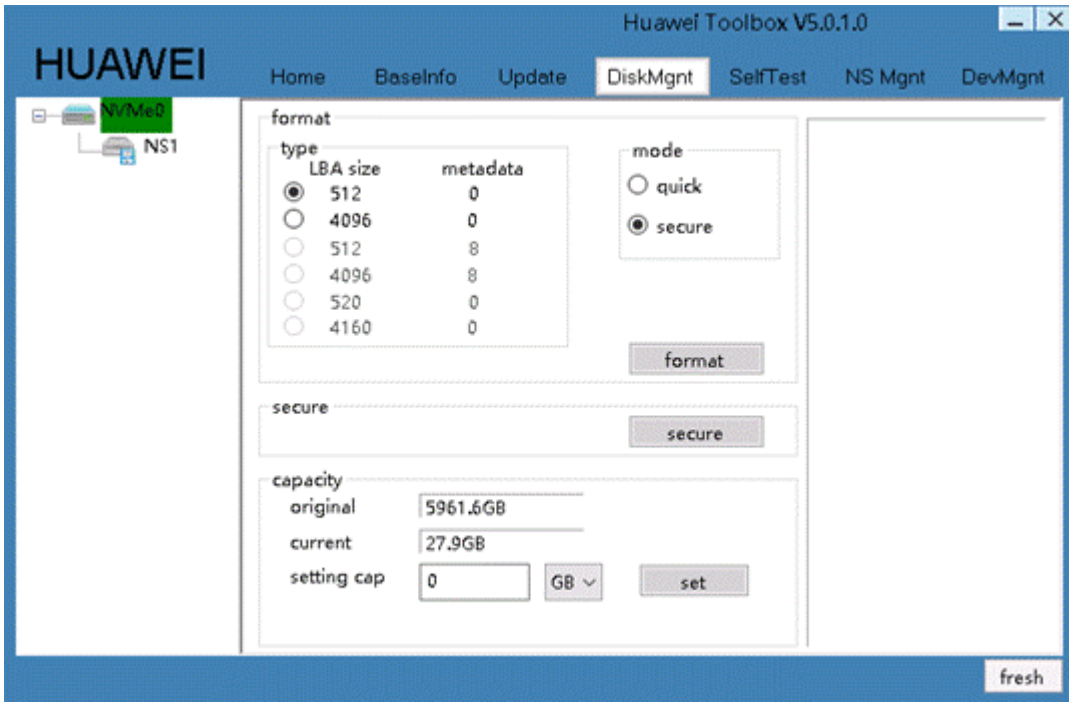

Functions on this page are the same as the operations described in **6.4 Formatting**, **6.3 Querying and Setting the Capacity**, and **6.5 Safely Erasing User Data of an SSD**.

If **setting cap** is set to **0**, the current capacity is set to the original capacity.

SelfTest Page

This page is used to diagnose the ES3000, as shown in **Figure F-7**.

Figure F-7 SelfTest page

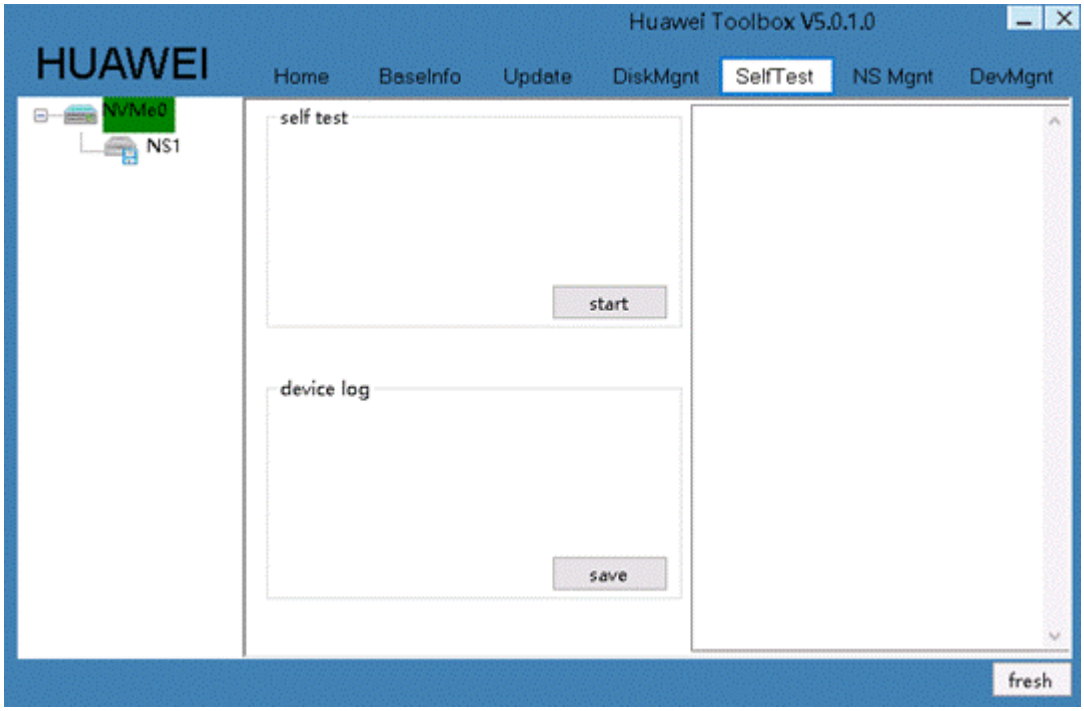

Functions on this page are the same as the operations described in **6.10 Querying the Diagnosis Result of a Specified SSD**.

NS Mgnt Page

This page is used to manage only NVMe namespaces, as shown in **Figure F-8**.

Figure F-8 NS Mgnt page

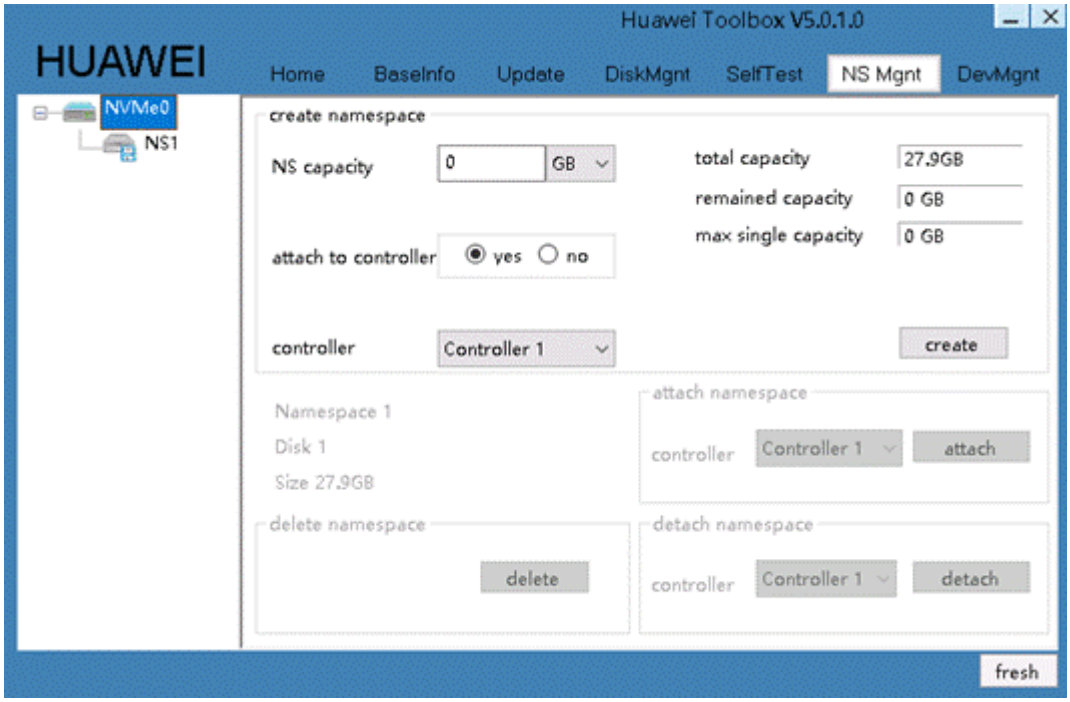

This page provides the following functions, which are the same as the operations described in **6.13 Querying and Configuring Namespaces**.

- Create a namespace. You can create a namespace with a size not exceeding the remaining capacity. You can also attach the namespace to specified controllers during the creation. See **Figure F-9**.
- Delete a namespace.
- Attach a namespace to controllers.
- Detach a namespace from controllers.

Figure F-9 Information about NS1

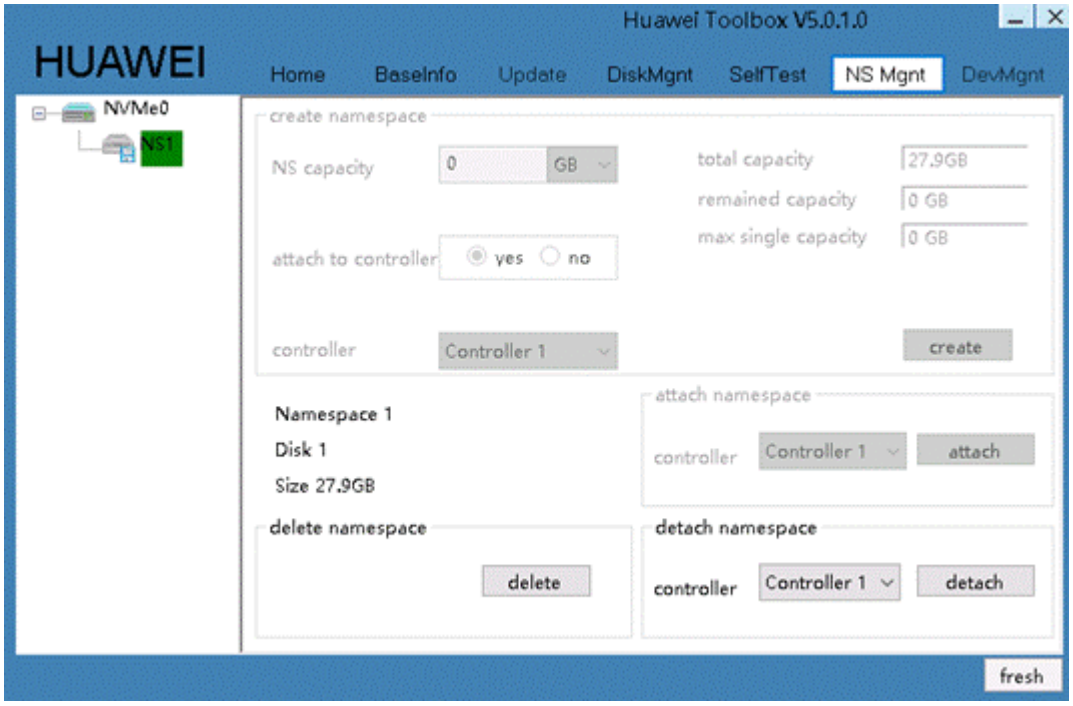

DevMgnt Page

This page is used to manage the ES3000, as shown in **Figure F-10**.

Figure F-10 NVMe DevMgnt page

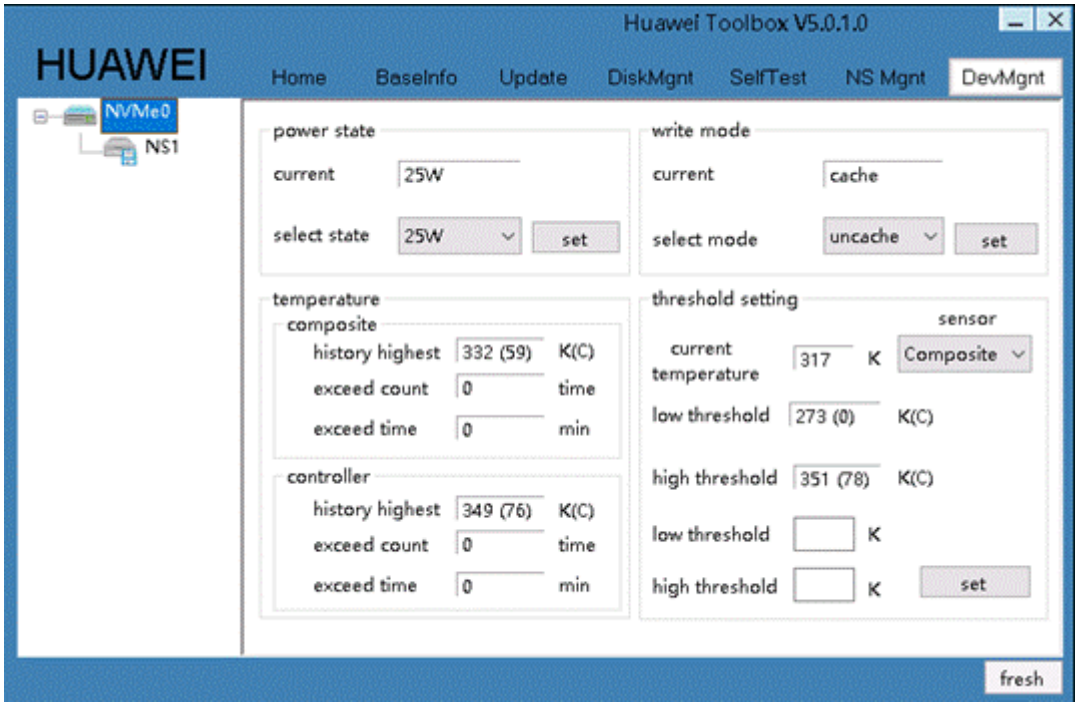

This page provides the following functions:

- **power state**: Set the power consumption level.
- **write mode**: Set the write mode.
- **temperature**: View drive temperature information, including the highest historical temperature, number of overtemperature occurrences, and overtemperature duration.
- **threshold setting**: Set temperature thresholds.
- **controller**: View controller temperature information, including the highest historical temperature, number of overtemperature occurrences, and overtemperature duration.

G Installing an OS on the ES3000

Operation Scenario

Use the **Huawei Server Compatibility Checker** to check the operating system versions and types supported by the ES3000 V5.

To install RHEL 7.0, CentOS 7.0, or OEL 7.0 in an ES3000 V5 with a 4K sector size, secure formatting the ES3000 V5 first by referring to 6.4.2 Formatting an SSD to ensure that the ES3000 V5 has no GPT partition information.

Before OS installation, set the BIOS boot mode to UEFI. If you are unfamiliar with the common OS installation procedure, see the **Huawei Server OS Installation Guide**.

NOTE

When the kdump function is used on the redhat6.8, centos6.8, and earlier operating systems, the vmcore file cannot be saved because the kexec-tools tool of the earlier version has a bug. Therefore, you need to upgrade the kexec-tools to a version later than 2.0.0-300.

To install Windows on the ES3000 V5, follow the procedure in the section because there are some special requirements.

If you format the ES3000 V5 into 4 KB sectors by following the instructions in 6.4.2 Formatting an SSD, create partitions manually when installing a Windows OS. For details about how to create partitions, see "Sample files: configure drive partitions by using Windows PE and DiskPart scripts" at **https://msdn.microsoft.com/en-us/library/windows/hardware/dn898510(v=vs.85).aspx**.

Note: When creating partitions, change the value of **select disk** to the drive number of the ES3000 V5 and change the value of **create partition efi size** to **260** in the text file. See **Figure G-1**.

Figure G-1 Creating partitions manually

```
rem == CreatePartitions-UEFI.txt ==
rem == These commands are used with DiskPart to
rem    create four partitions
rem    for a UEFI/GPT-based PC.
rem    Adjust the partition sizes to fill the drive
rem    as necessary. ==
select disk 0
clean
convert gpt
rem == 1. System partition =========================
create partition efi size=100
rem    ** NOTE: For Advanced Format 4Kn drives,
rem               change this value to size = 260 **
format quick fs=fat32 label="System"
assign letter="S"
rem == 2. Microsoft Reserved (MSR) partition =======
create partition msr size=16
rem == 3. Windows partition ========================
rem ==    a. Create the Windows partition ==========
create partition primary
rem ==    b. Create space for the recovery tools ===
shrink minimum=500
rem       ** NOTE: Update this size to match the
rem                size of the recovery tools
rem                (winre.wim) plus free space                    **
rem ==    c. Prepare the Windows partition =========
format quick fs=ntfs label="Windows"
assign letter="W"
rem === 4. Recovery tools partition ================
create partition primary
format quick fs=ntfs label="Recovery tools"
assign letter="R"
set id="de94bba4-06d1-4d40-a16a-bfd50179d6ac"
gpt attributes=0x8000000000000001
list volume
exit
```

Preparations

You have correctly installed the ES3000 V5 on the server.

Installing Windows Server 2012 64-Bit on the ES3000 V5

Step 1 Download the **hiodriver_2012_x86_64.img**driver image and mount it to the virtual floppy drive of the server.

Step 2 During the installation, select **Load driver** on the screen shown in **Figure G-2**.

Figure G-2 Selecting Load driver

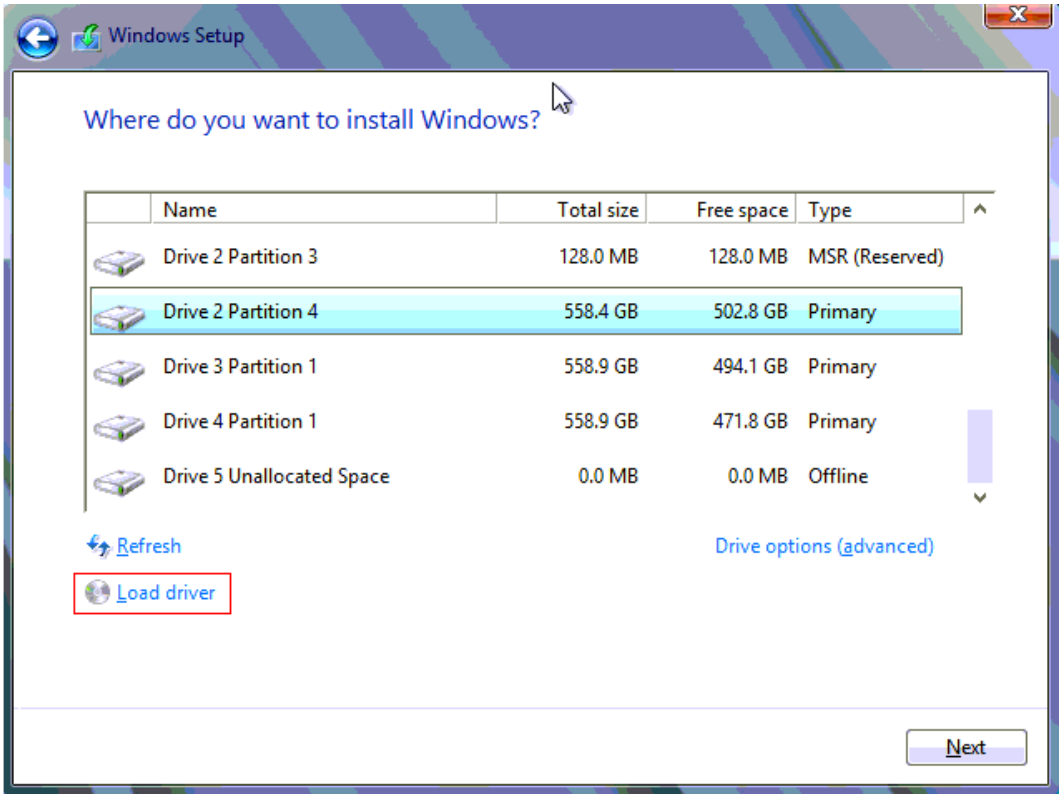

Step 3 Click **Next** to install the driver, as shown in **Figure G-3**.

Figure G-3 Installing the driver

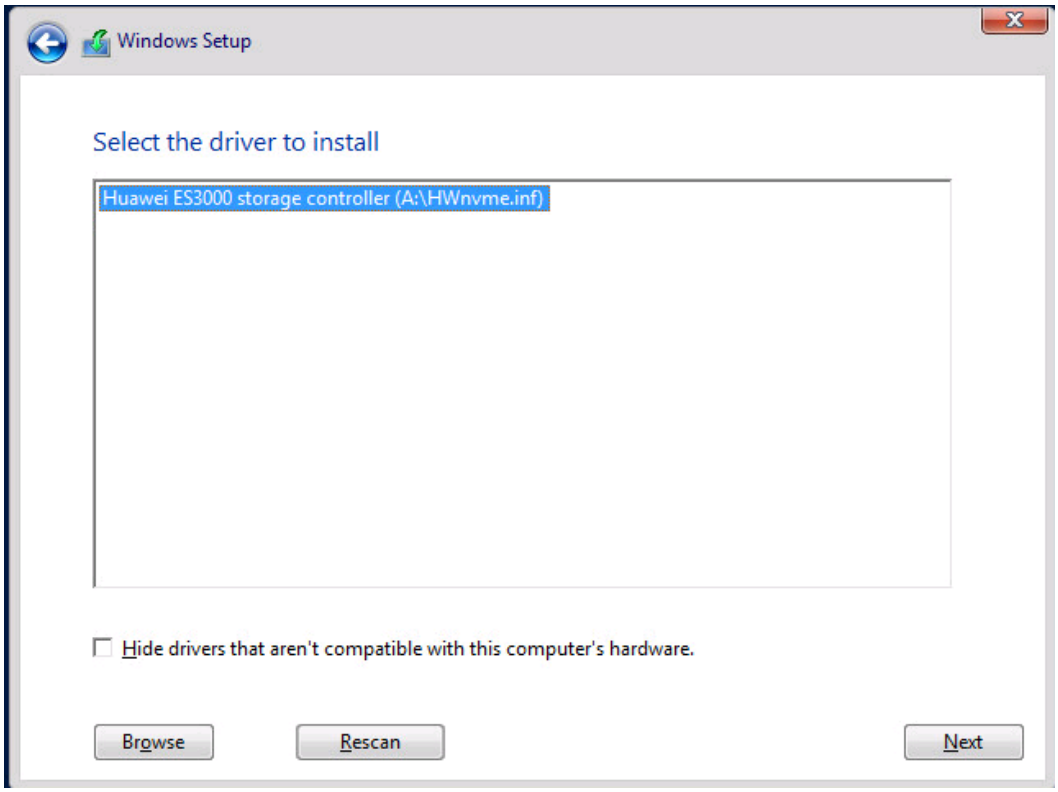

Step 4 After the driver is installed, select the ES3000 V5 for OS installation, as shown in **Figure G-4**.

Figure G-4 Selecting the ES3000 V5

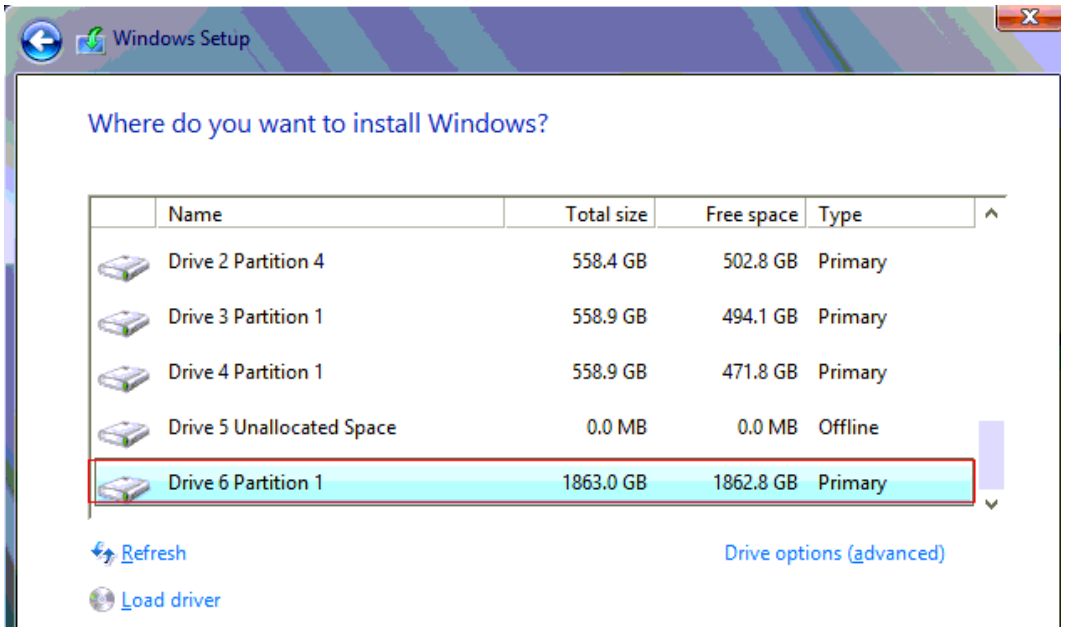

If the ES3000 V5 has previous partition information, such as OS data, the installation may fail. You can perform the following to prevent or solve the installation failure:

Step 5 On the screen shown in **Figure G-4**, press **Shift+F10** to open the CLI shown in **Figure G-5**.

Figure G-5 CLI

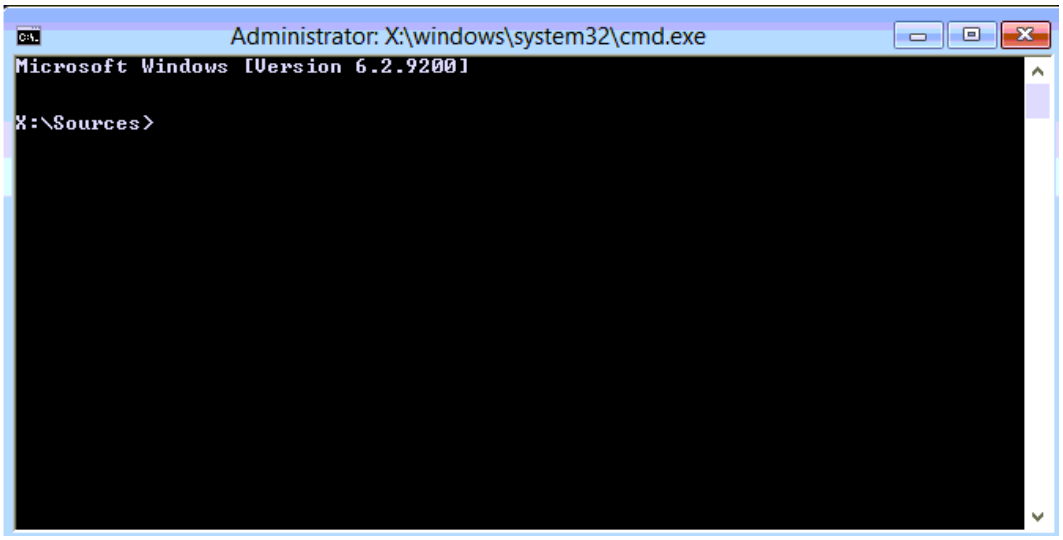

Step 6 Enter **diskpart** and press **Enter**, as shown in **Figure G-6**.

Figure G-6 Running the **diskpart** command

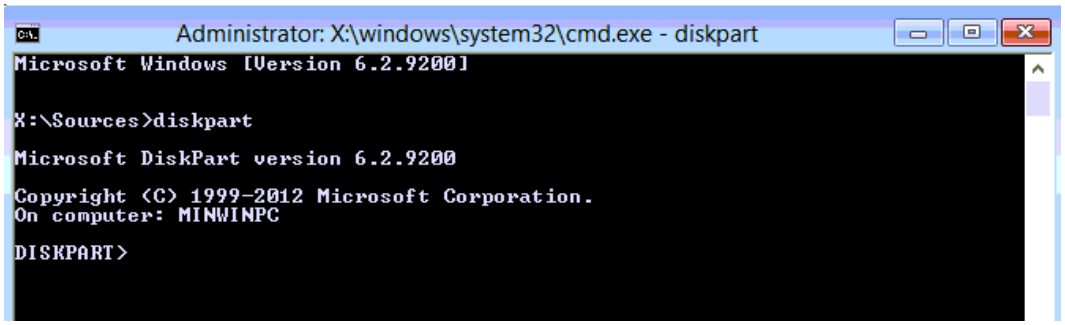

Step 7 Enter **select disk #** (# indicates the number of the ES3000 V5 drive for OS installation) and press **Enter**, as shown in **Figure G-7**.

Figure G-7 Selecting the ES3000 V5

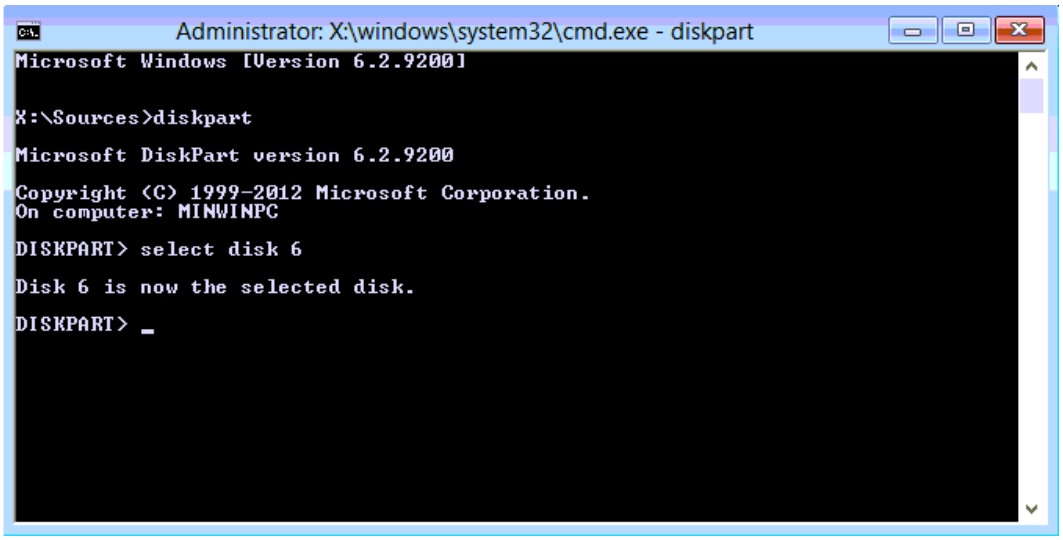

Step 8 Enter **clean** and press **Enter**, as shown in **Figure G-8**.

Figure G-8 Deleting partition information

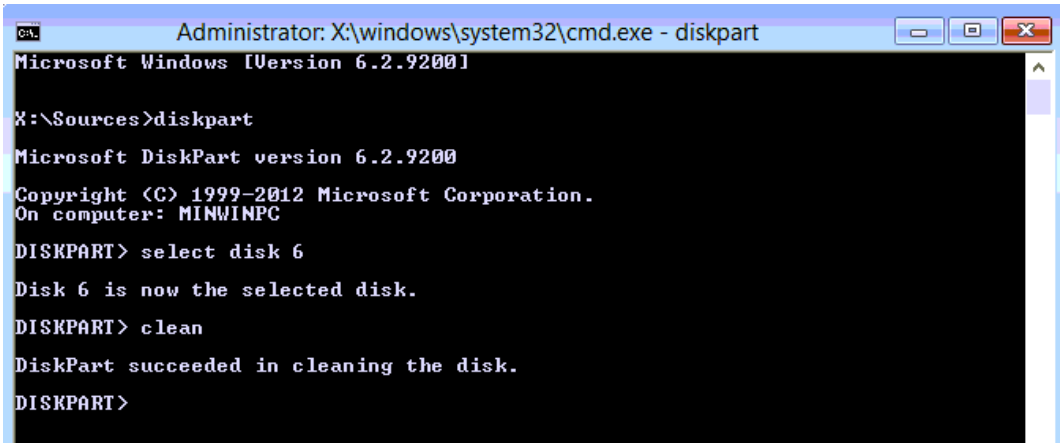

Step 9 Close the CLI and select the ES3000 V5 on the screen shown in **Figure G-4** to proceed the OS installation.

----End

Installing Windows Server 2012 R2 64-Bit on the ES3000 V5

Windows Server 2012 R2 is built with the NVMe driver. Therefore, the driver image does not need to be mounted for OS installation. Note that the ES3000 V5 is identified as 64 drives during the OS installation. Select the drive with a total size of greater than 0 for installing the OS.

⚠ CAUTION

After the OS is installed by using its built-in drivers, install the SSD driver **hiodriver_v*.*.*.*_2012R2_x84_64.msi** released by Huawei. For details about how to install, uninstall, and upgrade the Huawei SSD driver, see 3.6.2 Installing the ES3000 V5 Driver in Windows, 5.1.2 Uninstalling the Driver in Windows, and 5.2.2 Upgrading the ES3000 V5 Driver in Windows respectively.

If the ES3000 has previous partition information, such as OS data, the installation may fail. You can solve the problem in the same way as that for installing Windows Server 2012.

If the INACCESSIBLE_BOOT_DEVICE blue screen error is displayed due to other operations, perform the following procedure provided by Microsoft:

Mount the Windows Server 2012 R2 image and **hiodriverhiodriver_v*.*.*.*_2012R2_x84_64.img** file.

Boot from the CD-ROM.

Select **Repair your computer**.

Select **Troubleshoot**.

Select **Command Prompt**.

Enter **DISM /image:C:\ /Add-drivers /Driver:A:\HWnvme.inf**.

Restart the server.

H Getting Help

If you encounter any problems during routine maintenance or troubleshooting, contact Huawei technical support engineers.

H.1 Collecting Fault Information

Before troubleshooting, obtain the following information:

- Customer company and address
- Contact person and telephone number
- Time when the fault occurred
- Detailed fault symptom
- Device type and software version
- Any measures taken and effects
- Fault severity and expected rectification deadline

H.2 Preparing for Debugging

When you seek Huawei technical support, Huawei technical support engineers may assist you in performing some operations to further collect fault information or rectify the fault.

Before contacting technical support engineers, prepare the spare parts for boards and port modules, screwdrivers, screws, serial cables, and network cables.

H.3 Using Product Documentation

Huawei provides the documents delivered with the equipment. This document provides guidance for you to solve common problems that occur during routine maintenance or troubleshooting.

To better rectify the fault, you are advised to use the guide before contacting Huawei technical support engineers.

H.4 Obtaining Technical Support

Huawei's timely and efficient response is available from:

- Local branch offices
- Secondary technical support system
- Telephone support
- Remote support
- Onsite support

Technical Support Website

Obtain technical documents at **http://e.huawei.com**.

Self-Service Platform and Community

Learn more about servers and communicate with experts at:

- **HUAWEI Server Information Self-Service Platform** for specific server product documentation.
- **Huawei server intelligent Q&A system** for quick learning about products.
- **Huawei Enterprise Support Community (Servers)** for learning and discussion.

News

For notices about product life cycles, warnings, and updates, visit **Product Bulletins**.

Cases

Learn about server applications at **Knowledge Base**.

Huawei Technical Support

If a fault persists after taking the above measures, obtain technical support in the following ways:

- Contact Huawei customer service center.
  - Enterprise customers

    Send emails to **support_e@huawei.com** or visit **Global Service Hotline**.
  - Carriers

    Send emails to **support@huawei.com** or visit **Global TAC Information**.
- Contact technical support at your local Huawei office.

I Acronym or Abbreviation

Acronym or Abbreviation

C	
CCC	China Compulsory Certification
CRC	Cyclic Redundancy Check
E	
ECC	error checking and correcting
H	
HHHL	Half-height Half-length
I	
IOPS	I/O operations per second
L	
LBA	Logical Block Addressing
N	
NVMe	NVM Express
M	
MLC	multi-level cell
MTBF	mean time between failures
MIB	Management Information Base

P	
PCIe	PCI Express
R	
RAID	redundant array of independent disks
S	
SNMP	Simple Network Management Protocol
SSD	solid-state drive
V	
VROC	Intel® Virtual RAID on CPU