VANTAGEO Server

RAID User Guide (EagleStream)

Version: R1.1

VANTAGEO PRIVATE LIMITED
Corporate Address: 617, Lodha Supremus II,
Road No. 22, Wagle Estate,
Thane - 400604
URL: https://vantageo.com
E-mail: support@vantageo.com
Helpdesk - +91 18002669898

LEGAL INFORMATION

Statement on the Use of Third-Party Embedded Software:

If third-party embedded software such as Oracle, Sybase/SAP, Veritas, Microsoft, VMware, and Redhat is delivered together with this product of VANTAGEO, the embedded software must be used as only a component of this product. If this product is discarded, the licenses for the embedded software must be void either and must not be transferred. VANTAGEO will provide technical support for the embedded software of this product.

Revision History

Revision No.	Revision Date	Revision Reason
R1.1	2024-01-30	● Added "2 Applicable Server Models". ● Updated "4 VT SmartROC 3100 RAID Controller Card". ● Updated "5.1.2 Common Commands".
R1.0	2023-06-30	First edition.

Serial Number: VT20240308

Publishing Date: 2025-02-08 (R1.1)

Contents

About This Manual

Purpose

This manual describes the features of various RAID controller cards ("RAID cards" for short) installed on VANTAGEO servers, and the operations involved in initial configuration, common configuration and disk replacement.

Intended Audiences

This manual is intended for:

- Network planning engineers
- Network management engineers
- Maintenance engineers

What Is in This Manual

This manual contains the following chapters and appendixes.

Chapter 1, Basic RAID Concepts	Describes basic RAID-related concepts such as the RAID level, disk group, virtual disk, fault tolerance, and consistency check.
Chapter 2, Applicable Server Models	Describes the server models that this manual is applicable to.
Chapter 3, VT SmartIOC 2100 RAID Controller Card	Describes the capability features, initial configurations, common configurations, and typical disk replacement scenarios for a VT SmartIOC 2100 RAID controller card.
Chapter 4, VT SmartROC 3100 RAID Controller Card	Describes the capability features, initial configurations, common configurations, and typical disk replacement scenarios for a VT SmartROC 3100 RAID controller card.
Chapter 5, Appendixes	Describes the CLI tool ARCCONF for RAID controller cards, and VROC-based RAID configuration steps and common operations.

Conventions

This manual uses the following conventions.

	Notice: indicates equipment or environment safety information. Failure to comply can result in equipment damage, data loss, equipment performance degradation, environmental contamination, or other unpredictable results. Failure to comply will not result in personal injury.
	Note: provides additional information about a topic.

Chapter 1
Basic RAID Concepts

Table of Contents

1.1 RAID Levels

RAID refers to a disk array that combines multiple physical disks for the purpose of data redundancy.

It can provide higher storage performance, I/O performance and reliability than a single disk.

Commonly used RAID levels include RAID 0, RAID 1, RAID 1E, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60.

1 1.1 RAID 0

RAID 0 is also called stripe or striping. Figure 1-1 shows its data storage principle.

Figure 1-1 RAID 0 Data Storage Diagram

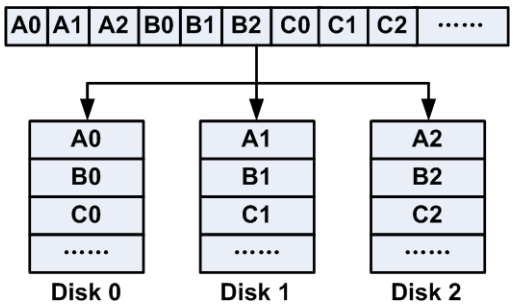

The storage principle of RAID 0 is to distribute sequential data (A0–A2, B0–B2, C0–C2, …) to several member disks.

RAID 0 provides the highest storage performance among all RAID levels, because when the system receives a data request, multiple member disks in the RAID 0 array can concurrently execute the data request. This type of concurrent data operation can make full use of the bus bandwidth, remarkably boosting the overall disk read/write performance.

Because RAID 0 provides no data redundancy, data security cannot be guaranteed. Therefore, RAID 0 can only be used in scenarios where the requirements for I/O are high, but the requirements for data security are low.

RAID 0 requires at least one disk, and the data will be lost if the disk fails.

1.1.2 RAID 1

RAID 1 is also called mirror or mirroring, see Figure 1-2.

Figure 1-2 RAID 1 Data Storage Diagram

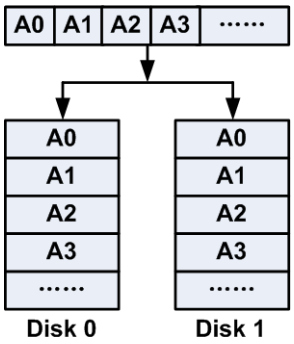

Each working disk of the RAID 1 array has a corresponding mirrored disk, and a piece of sequential data (A0, A1, A2, …) needs to be written into the working disk and the mirrored disk respectively, which is equivalent to automatic backup.

When the working disk of the RAID 1 array is faulty, the system automatically reads data from the mirrored disk.

RAID 1 requires at least two disks, and data is not lost if one disk fails.

RAID 1 provides high reliability but only a half of the total capacity. The disk usage is low. Therefore, RAID 1 is applicable to the scenarios where high fault tolerance is required, such as finance.

1.1.3 RAID 1E

RAID 1E is an enhanced version of RAID 1, which integrates data mirroring and data striping, see Figure 1-3.

Figure 1-3 RAID 1E Data Storage Diagram

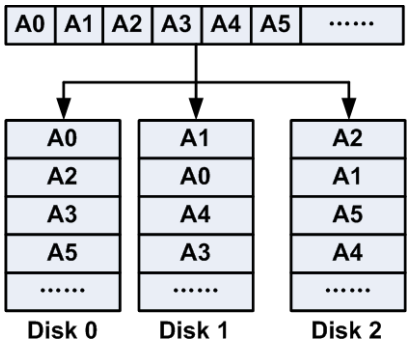

Data striping and data backup of RAID 1E spread across all the member disks. Take two stripe groups (for example, A0, A1, A2 and A2, A1 and A0) as one unit. In each unit, the data in the previous stripe group is arranged continuously, and the next stripe group is the striped mirror of the previous stripe group. That is, the data in each member disk of the previous stripe group is moved rightwards one disk and written into each member disk of the next stripe group. The data of the last member disk is written into the first member disk.

RAID 1E requires at least three disks, and data is not lost if one disk fails. When a member disk of the RAID 1E array fails, the system transfers the read/write requests to other normal member disks for processing.

Compared with RAID 1, RAID 1E also uses the mirroring mode to store data, so its effective capacity is half of the total hard disk capacity. Because RAID 1E allows access to more member disks, it has stronger data recovery capability. However, each piece of data must be written at least twice, which reduces the disk write capability.

1.1.4 RAID 5

RAID 5 is a solution that balances storage performance, data security and storage costs, see Figure 1-4.

Figure 1-4 RAID 5 Data Storage Diagram

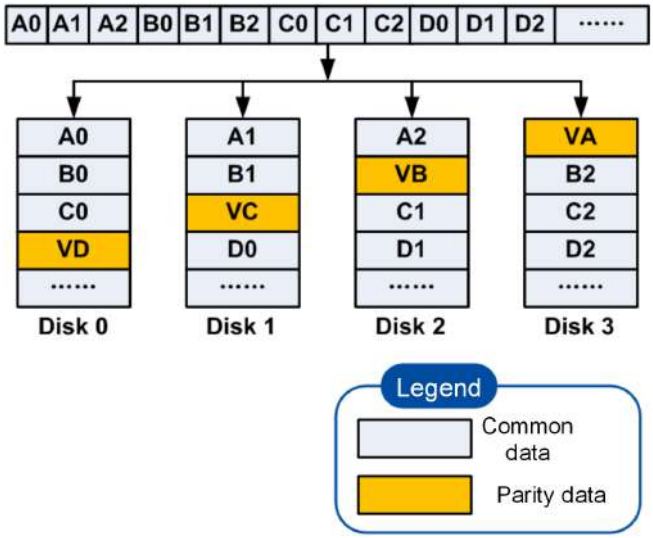

RAID 5 uses the CRC mode to distribute sequential data (A0–A2, B0–B2, C0–C2, D0–D2, …) and the corresponding parity data (VA, VD, …) to different member disks respectively.

RAID 5 requires at least three disks, and data is not lost if one disk fails. If a member disk in the RAID 5 array is faulty, the data on the faulty member disk can be rebuilt from the data on other member disks in the array through the XOR operation.

Note

For example, when the data block A0 is damaged, the recovery can be calculated with the following formula:

$$A0 = A1 \oplus A2 \oplus VA$$

RAID 5 can be used to process both a large or small amount of data. It features high speed, large capacity, and fault tolerance distribution. It is often used in various private NAS servers.

RAID 5 can be regarded as a trade-off between RAID 0 and RAID 1.

1.1.5 RAID 6

On the basis of RAID, RAID 6 is formed to further enhance data protection, see Figure 1-5.

Figure 1-5 RAID 6 Data Storage Diagram

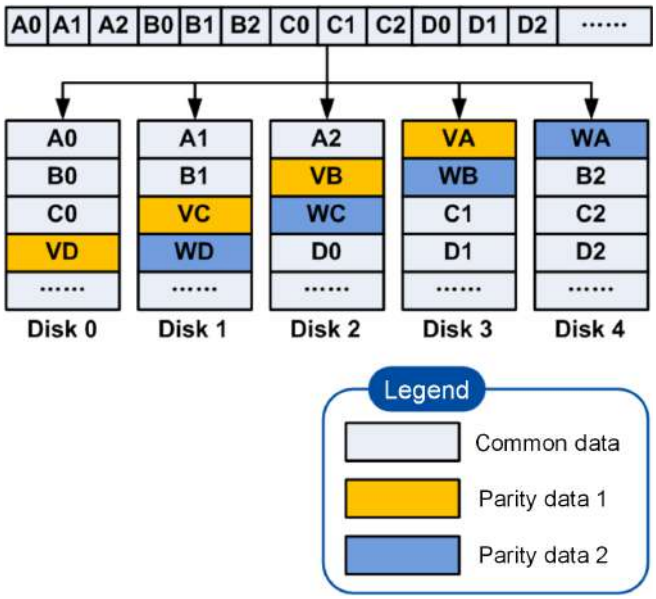

Like RAID 5, RAID 6 distributes sequential data (A0–A2, B0–B2, C0–C2, D0–D2, …) and the corresponding parity data (VA, VD, …) to different member disks respectively.

The difference is that on the basis of RAID 5, RAID 6 adds a second parity block (WA, WB, WC and WD, ……). The two independent check systems use different algorithms to implement double check and enhance fault tolerance. Therefore, when two member disks of the RAID 6 array are faulty, data security can also be guaranteed, achieving high reliability.

RAID 6 requires at least four disks, and data is not lost if two disks fail. If two member disks in a RAID 6 array is faulty, the data on the faulty member disks can be rebuilt from the data on other member disks in the array through the XOR (#) operation and the coefficient and XOR operation.

Note

For example, when two data blocks A0 and A1 are damaged at the same time, you can recover them by using the following formula:

$$A0 = A1 \oplus A2 \oplus VA$$

$$A1 = (A0 \times K0 \oplus A2 \times K2 \oplus WA) \ / \ K1$$

In the above formula, K0, K1 and K2 are polynomial values in the Galois domain, and are known coefficients.

The read performance and fault tolerance performance of RAID 6 are higher, but more disk space is allocated for parity data. Therefore, compared with RAID 5, it has greater write performance loss and higher implementation costs.

1.1.6 RAID 10

RAID 10 is a combination of RAID 1 and RAID 0, see Figure 1-6.

Figure 1-6 RAID 10 Data Storage Diagram

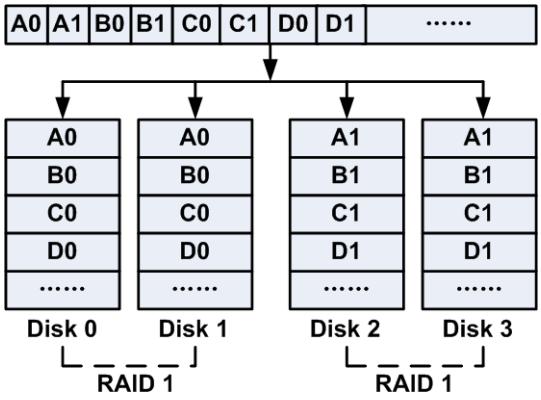

RAID 10 can be regarded as: Two member disks in a pair (Disk 0 and Disk 1, Disk 2 and Disk 3, …) form a multiple RAID 1 volume. Multiple RAID 1 volumes form a RAID 0 array, and data is distributed on each RAID 1 volume in the form of RAID 0.

After receiving the I/O data requests, the system distributes the sequential data (A0, A1, B0, B1, C0, C1, …) to two RAID 1 volumes for concurrent processing in accordance with the working mode of RAID 0. That is, the data is stored in the working disks Disk 0 and Disk 2. At the same time, in the manner of RAID 1, the system automatically copies data to mirrored disk Disk 1 when writing data into Disk 0, and copies data to the mirrored disk Disk 3 when writing data into Disk 2.

RAID 10 requires at least four disks, and data is not lost if two disks fail.

RAID 10 balances storage performance and data security. It not only provides data security (same as RAID 1) but storage performance (similar to RAID 0).

1.1.7 RAID 50

RAID 50 is a combination of RAID 5 and RAID 0, see Figure 1-7.

Figure 1-7 RAID 50 Data Storage Diagram

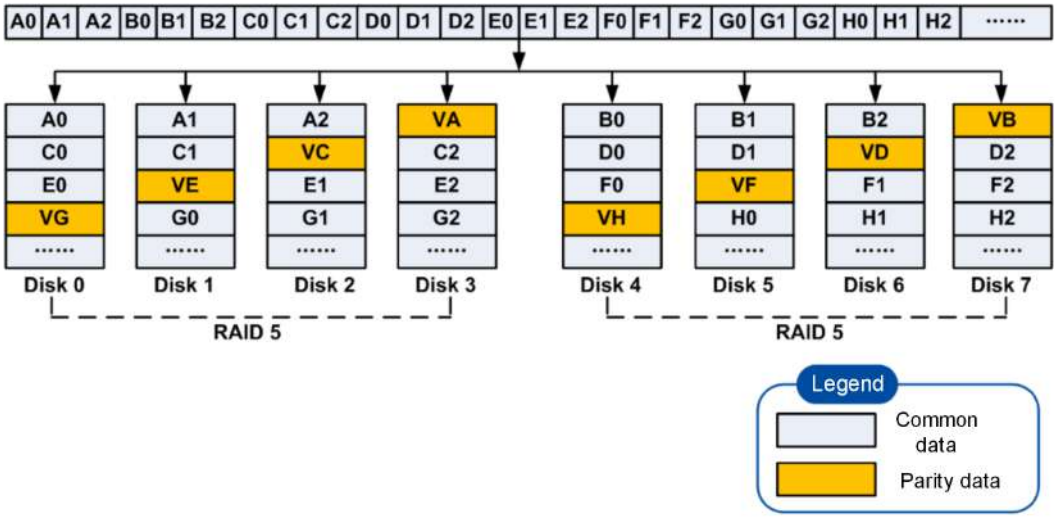

One RAID 50 array consists of multiple RAID 5 volumes. At least six hard disks are required for RAID 50. Data is distributed on each RAID 5 volume in RAID 0 mode.

RAID 50 supports the features of both RAID 0 and RAID 5:

- Like RAID 0, data is partitioned into stripes and written into all the member disks at the same time.
- Like RAID 5, data security is guaranteed through parity bits that are evenly distributed on all the member disks.

With the redundancy function provided by RAID 5, if any disk in a RAID 5 volume is faulty, the array can operate properly and restore the data from the faulty disk. The replacement of the faulty disk does not affect services. Therefore, RAID 50 allows a single disk failure in multiple RAID 5 volumes at the same time, greatly improving the fault tolerance capability.

At the same time, RAID 50 inherits the high storage performance of RAID 0, and stores the data distributed in multiple RAID 5 volumes. It can process data concurrently, thus providing better read/write performance.

1.1.8 RAID 60

RAID 60 is a combination of RAID 6 and RAID 0, see Figure 1-8.

Figure 1-8 RAID 60 Data Storage Diagram

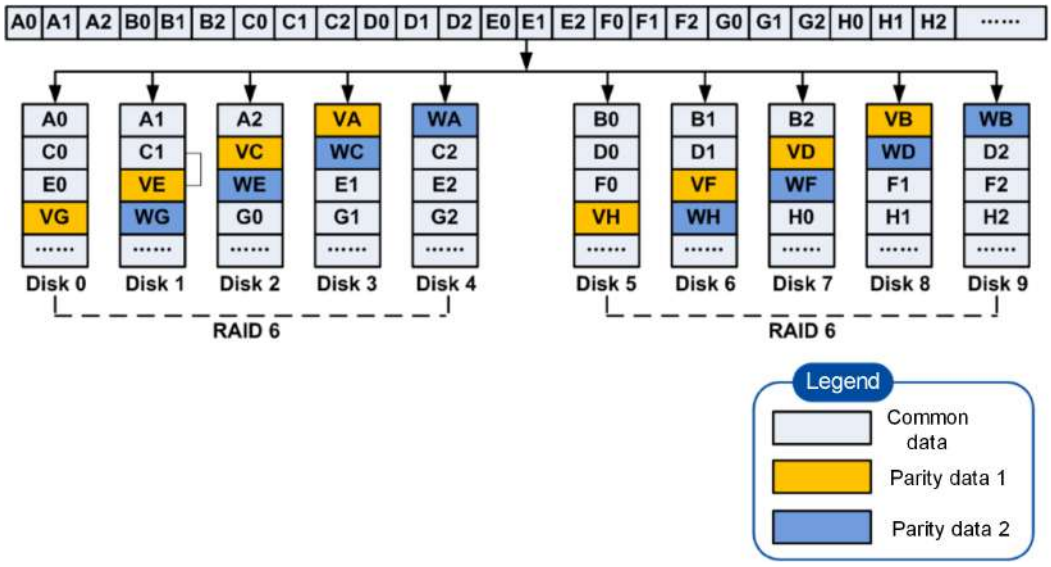

One RAID 60 array consists of multiple RAID 6 volumes. At least eight hard disks are required for RAID 60. Data is distributed on each RAID 6 volume in RAID 0 mode.

RAID 60 supports the features of both RAID 0 and RAID 6:

- Like RAID 0, data is partitioned into stripes and written into all the member disks at the same time.
- Like RAID 6, RAID 60 ensures data security by using two parity blocks that are evenly distributed on all the member disks.

If two member disks of the RAID 6 array are faulty, with the redundancy provided by RAID 6, the array can operate properly and rebuild data of the faulty disk. The replacement of faulty disks does not affect services. Therefore, RAID 60 can maintain the security of data in the faulty disks when dual disks are faulty in multiple RAID 6 volumes at the same time.

In addition, RAID 60 inherits the high storage performance of RAID 0, and stores the data distributed in multiple RAID 6 volumes. It can process data concurrently, thus providing better read/write performance.

1.2 RAID Performance Comparison

Fault Tolerance Capabilities

For a description of the fault tolerance capability comparison between the commonly used RAID levels, refer to Table 1-1. For the RAID containing multiple subgroups, RAID 10 and RAID 50 allow the number of faulty disks to be the same as the number of subgroups, but each subgroup must contain only one faulty disk. RAID 60 allows the number of faulty disks to be twice that of subgroups. Each subgroup can contain a maximum of two faulty disks.

Table 1-1 Fault Tolerance Capability Comparison Between Commonly Used RAID Levels

RAID Level	Fault Tolerance Capacity Description
RAID 0	It does not support fault tolerance. A fault in any member disk will affect data security.
RAID 1	It supports fault tolerance through data redundancy. When a member disk is faulty, the data on the corresponding mirrored disk can be used to run the system and rebuild the faulty disk.
RAID 1E	It supports fault tolerance through data redundancy. When a member disk is faulty, normal disks can be used to run the system and rebuild the faulty disk. If there are more than four member disks, two member disks are allowed to be faulty. However, the prerequisites are that the two faulty disks cannot be adjacent to each other, and the two faulty disks cannot be the first member disk and the last member disk.
RAID 5	It supports fault tolerance through the distributed parity data. When a member disk is faulty, the RAID controller card uses parity data to rebuild all the lost data.
RAID 6	It supports fault tolerance through the distributed parity data. When two member disks are faulty, the data on the normal member disks and the corresponding two sets of parity data can be used to calculate and rebuild the faulty disk.
RAID 10	It supports fault tolerance through data redundancy. Each RAID 1 volume allows one member disk to be faulty. The data in the corresponding mirrored disk in the group can be used to run the system and rebuild the faulty disk.
RAID 50	It supports fault tolerance through the distributed parity data. Each RAID 5 volume allows one member disk to be faulty. The data on normal disks and the corresponding one set of parity data can be used to calculate and rebuild the faulty disk.
RAID 60	It supports fault tolerance through the distributed parity data. Each RAID 6 volume allows two member disks to be faulty. The data on normal disks and the corresponding two sets of parity data can be used to calculate and rebuild the faulty disks.

I/O Performance

A RAID array can be used as an independent storage unit or multiple virtual units by the system. Because multiple disks can be accessed at the same time, the I/O performance for a RAID array is higher than that for a common disk. For a description of the I/O performance comparison between the commonly used RAID levels, refer to Table 1-2.

Table 1-2 I/O Performance Comparison Between Commonly Used RAID Levels

RAID Level	I/O Performance Description
RAID 0	RAID 0 divides data into smaller data blocks and writes them into different disks. Because multiple disks can be read and written at the same time, RAID 0 improves the I/O performance.

RAID Level	I/O Performance Description
RAID 1	Because the disks in a RAID group exist in pairs, data must be written in two copies at the same time, occupying longer time and resources. As a result, the performance deteriorates.
RAID 1E	Each piece of data must be written at least twice, which reduces the disk write capability.
RAID 5	Because both common data and parity data is kept on member disks, each member disk can be read and written separately. Therefore, RAID 5 provides high data throughput.
RAID 6	RAID 6 needs to write two sets of parity data across member disks, which causes performance deterioration during write operations. RAID 6 is ideal in scenarios where high reliability, high response rate and high transmission rate are required to provide high data throughput, high data redundancy and high I/O performance.
RAID 10	With the high data transmission rate provided by the RAID 0 subgroup, RAID 10 performs well in data storage. The I/O performance is improved as the number of subgroups increases.
RAID 50	RAID 50 outperforms other RAID levels in the scenarios that require high reliability, high response rate and high transmission rate. The I/O performance is improved as the number of subgroups increases.
RAID 60	Its application scenarios are similar to those of RAID 50. However, because two sets of parity data need to be written into each member disk, its performance is deteriorated during write operations. Therefore, RAID 60 is not applicable to tasks that require writing massive data.

Available Capacity

For a description of the storage capacity comparison between the commonly used RAID levels, refer to Table 1-3.

Table 1-3 Storage Capacity Comparison Between Commonly Used RAID Levels

RAID Level	Storage Capacity Description
RAID 0	In all RAID levels, RAID 0 provides the maximum storage capacity. Available capacity = minimum capacity of a member disk x number of member disks
RAID 1	Data must be written in two copies, causing high storage space consumption. Available capacity = minimum capacity of a member disk x number of member disks ÷ 2
RAID 1E	Two stripe groups are used as one unit, and the next stripe group is the striped mirror of the last stripe group, causing high storage space consumption. Available capacity = minimum capacity of a member disk x number of member disks ÷ 2

RAID Level	Storage Capacity Description
RAID 5	Parity data blocks are isolated from common data blocks. In general, parity data occupies the capacity of one member disk. Available capacity = minimum capacity of a member disk x (number of member disks - 1)
RAID 6	Two independent parity data blocks are isolated from common data blocks. In general, the parity data occupies the capacity of two member disks. Available capacity = minimum capacity of a member disk x (number of member disks - 2)
RAID 10	It consists of multiple RAID 1 volumes and has the same capacity as RAID 1. Available capacity = minimum capacity of a member disk x number of member disks ÷ 2
RAID 50	It consists of multiple RAID 5 volumes. Its capacity is related to the number of RAID 5 volumes. Available capacity = total capacity of all the RAID 5 volumes
RAID 60	It consists of multiple RAID 6 volumes. Its capacity is related to the number of RAID 6 volumes. Available capacity = total capacity of all the RAID 6 volumes

Performance Comparison Summary

Each RAID level has its advantages and disadvantages in terms of fault tolerance, I/O performance and available capacity. For the overall comparison of the performance indicators, refer to Table 1-4.

Table 1-4 Performance Comparison Summary

RAID Level	Fault Tolerance Capability	I/O Performance		Number of Required Disks	Disk Usage
		Read Performance	Write Performance		
RAID 0	Low	High	High	N≥2	100%
RAID 1	High	High	Medium	2N (N≥1)	50%
RAID 1E	High	High	Medium	N≥3	50%
RAID 5	Medium	High	Medium	N≥3	(N-1)/N
RAID 6	High	High	Medium	N≥4	(N-2)/N
RAID 10	High	Medium	Medium	2N (N≥2)	50%
RAID 50	High	High	High	N≥n×M (n≥3, M≥2)	(N-M)/N
RAID 60	High	High	High	N≥n×M (n≥4, M≥2)	(N-2M)/N

1. N is the total number of RAID member disks.
2. M is the number of RAID volumes.
3. n is the number of member disks in each volume.

1.3 RAID-Related Features

1.3.1 Disk Group and Virtual Disk

With the explosive growth of modern data centers, data service traffic is increasing day by day, resulting in higher requirements for data storage capacity and read/write speed of servers. As a result, common disks cannot meet these requirements. In this case, it is necessary to form multiple independent disks into a super-large disk group in a specific way to provide better storage performance, I/O performance and security.

A virtual disk is a continuous data storage unit divided based on an overall disk group. Virtual disks can be configured to have larger capacity, higher security, and more data redundancy than a single physical disk.

A virtual disk can be:

- A complete disk group
- Multiple complete disk groups
- Part of a disk group
- Parts of multiple disk groups (a part is divided from each disk group and they constitute a virtual disk)

Related conventions:

- A disk group (DG for short) is also called a drive group or array.
- A virtual disk is also called virtual drive (VD for short), volume or logical disk (LD for short).

1.3.2 Fault Tolerance

Fault tolerance means that data integrity and data processing capabilities are not affected when a disk error or disk fault occurs in a disk group.

Fault tolerance improves the disk system availability, guaranteeing system operation. Therefore, fault tolerance is a very important feature in the fault recovery procedure.

RAID can use redundant data to restore the data errors occurring during the calculation or transmission to achieve fault tolerance. RAID 1, RAID 1E, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60 provide the fault tolerance capability. For the fault tolerance description at each level, refer to “Fault Tolerance Capabilities”.

Note

RAID 0 does not provide fault tolerance. Once a disk fails, its data is lost.

1.3.3 Consistency Check

For RAID 1, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60 with the redundancy function, the RAID controller card calculates the data stored on the disks, and compares these data with the corresponding redundant data. If any inconsistency is found, an attempt is made to automatically fix the inconsistency and save the error information. It is recommended that a consistency check be performed at least once a month.

Note

Because RAID 0 does not provide the redundancy function, it does not support the consistency check.

1.3.4 Hot Spare

Hot Spare

Hot spare refers to the backup performed when the system is operating properly. The hot spare feature of the RAID controller card is implemented by hot spare and emergency spare.

Hot Spare Disk

A hot spare disk is an independent disk in the disk system. When a disk in a RAID group is faulty, the hot spare disk is automatically added to the RAID group to replace the faulty disk for rebuilding data and providing fault tolerance.

On RAID controller card management screens or CLI, an idle disk whose capacity is greater than or equal to that of a member disk in a RAID group and whose media type and interface are the same as those of the member disk can be specified as the hot spare disk of the RAID group.

The RAID controller card supports the following two types of hot spare disks:

- Global hot spare disk, which is shared by all the configured RAID groups of a RAID controller card. One RAID controller card can be configured with one or more global hot spare disks. If a member disk in any RAID group fails, the global hot spare disk can automatically replace the faulty disk.
- Dedicated hot spare disk, which is exclusive to a specific RAID group of a RAID controller card. Each RAID group can be configured with one or more dedicated hot spare disks. If a member disk in the specified RAID group fails, the dedicated hot spare disk can automatically replace the faulty disk.

A hot spare disk has the following features:

- The hot spare disk is only used for the RAID groups with redundancy, including RAID 1, RAID 1E, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60.
- The hot spare disk is only used to replace a faulty disk managed by the same RAID controller card.

Emergency Spare

The emergency spare function means that if a member disk in any RAID group with the redundancy function is faulty and no hot spare disk is specified, an idle disk managed by the RAID controller card automatically replaces the faulty member disk and rebuilds the data to avoid data loss.

Emergency spare requires that the capacity of the idle disk must be greater than or equal to that of the member disk, and the media type and interface of the idle disk must be the same as that of the member disk.

1.3.5 Data Rebuild

Description

If a faulty disk exists in a RAID group, you can use the data rebuild function of the RAID controller card to rebuild data in the faulty disk for a new disk. The data rebuild function is only used for the RAID arrays with redundancy, including RAID 1, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60.

The RAID controller card supports automatic data rebuild:

- If an available hot spare disk is specified for a RAID group, when a member disk is faulty, the hot spare disk automatically replaces the faulty disk and rebuilds the data.
- If no available hot spare disk is specified for a RAID group, when a member disk is faulty, an idle disk in the RAID group automatically replaces the faulty disk and rebuilds the data.

Note

The hot spare disk and idle disk used for data rebuild must meet the following requirements:
- The capacity is larger than or equal to the RAID member disk.
- The media type and interface are the same as those of the member disk.

If no hot spare disk is specified for a RAID group and no idle disk meets requirements, data can be rebuilt only after the faulty disk is replaced with a new one.

After the hot spare disk begins data rebuild, the faulty member disk is marked as removable. If the system is powered off during the data rebuild process, the RAID controller card continues the data rebuild task after the system is restarted.

Rebuild Rate

The rebuild rate is the proportion of CPU resources occupied by a data rebuild task to the overall CPU resources during the system operation. It can be set to 0%–100%.

- The value **0%** indicates that the data rebuild task is started only when there is no other task running in the system.
- The value **100%** indicates that the data rebuild task occupies all CPU resources.

Note

The rebuild rate can be set as required. It is recommended that you set a proper value based on the actual system conditions.

1.3.6 Disk Status

Physical Disk Status

For a description of the possible status of a physical disk managed by a RAID controller card, refer to Table 1-5.

Table 1-5 Physical Disk Status

Status	Description
Online	Indicates that the disk is a member disk of a virtual disk and is online and can be used properly.
Unconfigured Good	Indicates that the disk is in normal status, but is not a member disk or a hot spare disk of a virtual disk.
Hot Spare	Indicates that the disk is set as a hot spare disk.
Failed	This status appears when an unrecoverable error occurs on a disk in **Online** or **Hot Spare** status.
Rebuild/Rebuilding	Indicates that the disk is rebuilding data to ensure the data redundancy and integrity of the virtual disk. In this case, the performance of the virtual disk is affected.
Unconfigured Bad	This status appears when an uninitialized disk or a disk in **Unconfigured Good** status has an unrecoverable error.
Missing	This status appears after the disk in the **Online** status is removed.
Offline	Indicates that the disk is a member disk of a virtual disk and is offline and can be used properly.
Shield State	Indicates the temporary status in which the physical disk is performing a diagnosis operation.
Copyback	Indicates that a new disk is replacing the hot spare disk.

Status	Description
Unsupport	Indicates that the specification of the disk exceeds the current specification of the RAID controller card.
Raw (Pass Through)	Indicates that the disk is a pass-through disk in HBA mode. For the concept of pass-through, refer to "1.3.13 Disk Pass-Through".
Ready	Indicates that the disk in this status can be used to configure a RAID volume, which is applicable to the **RAID** mode and **Mixed** mode of the RAID controller card. ● In **RAID** mode, the connected disks can be used in the system only after they form a RAID volume, and the disks in **Ready** status are not reported to the OS. ● In **Mixed** mode, the connected disks can be used directly or form a RAID volume, and the disks in **Ready** status are reported to the OS.
Predictive Failure	Indicates that the disk is about to change to the **Failed** status. The data on the disk needs to be backed up and the disk needs to be replaced.

Virtual Disk Status

For a description of the possible status of a virtual disk created under a RAID controller card, refer to Table 1-6.

Table 1-6 Virtual Disk Status

Status	Description
Optimal	Indicates that the virtual disk is in good condition and all member disks are online.
Degraded	Indicates that the virtual disk is available but abnormal, and some member disks are faulty or offline.
Failed	Indicates that the virtual disk is faulty.
Partial Degraded	This status appears when the number of faulty or offline disks in the RAID group does not exceed the maximum number of faulty disks supported by the RAID array.
Offline	This status appears when the number of faulty or offline disks in the RAID group exceeds the maximum number of faulty disks supported by the RAID array.

1.3.7 Read/Write Policy for Virtual Disks

Overview

When creating a virtual disk, you need to define the data read/write policies to standardize subsequent data read and write operations on it.

Data Read Policy

"**Read Policy**" is generally displayed on the configuration screen. The RAID controller card supports the following two types of data read policies:

- Read-ahead mode: When a RAID controller card reads the required data from the virtual disk, it simultaneously reads the subsequent data and writes it into the cache. If a user needs to access the data, the data can be directly read from the cache to reduce the track seeking operation, save the response time, and improve the data read speed.

  Notice

  To use this policy, the RAID controller card must support data protection against a power supply failure. However, if the super capacitor is abnormal at this time, the data may be lost.

- Non-read-ahead mode: The RAID controller card reads data from the virtual disk only after receiving a data read command. It does not perform the read-ahead operation.

Data Write Policy

"**Write Policy**" is generally displayed on the configuration screen. The RAID controller card supports the following three types of data write policies:

- Write back: If this policy is used, when data needs to be written to a virtual disk, it is directly written into the cache. The RAID controller card refreshes the data to the virtual disk only when the written data is accumulated to a certain extent, achieving batch data write and improving the data write speed. After receiving all the transmitted data, the cache returns the data transmission completion signal to the host.

  Notice

  To use this policy, the RAID controller card must support data protection against a power supply failure. However, if the super capacitor is abnormal at this time, the data may be lost.

- Direct write: The RAID controller card directly writes data to the virtual disk without being cached. However, the write speed is low. After the virtual receives all the transmitted data, the RAID controller card returns the data transmission completion signal to the host.

  Note

  This policy does not require the RAID controller card to support data protection against a power supply failure. Even if the super capacitor fails, there is no impact.

- BBU-related write-back:

→ When the BBU of the RAID controller card is present and operates properly, the write operation from the RAID controller card to the virtual disk is transited through the cache (write-back mode).

→ When the BBU of the RAID controller card is not present or the BBU is faulty, the write operation from the RAID controller card to the virtual disk is automatically switched to the direct write mode without being cached.

1.3.8 Data Protection Against a Power Supply Failure

Protection Principle

The speed of writing data into the high-speed cache of a RAID controller card is greater than that of writing data into a disk. When the server writes a large amount of data, the high-speed cache of the RAID controller card can be used to improve system performance.

After the high-speed cache is enabled, the advantages and disadvantages are as follows:

- The write performance of the server is improved. When the write pressure on the server is reduced or the high-speed cache of the RAID controller card is to be full, data is written into a disk from the high-speed cache.
- The risk of data loss increases. When the server is powered off accidentally, the data in the high-speed cache of the RAID controller card will be lost.

To improve the high read/write performance of the server and ensure the data security in the high-speed cache of the RAID controller card, you can configure a super capacitor for the RAID controller card. In case of unexpected power failure of the server, the super capacitor is used to supply power to the RAID controller card, and data in the high-speed cache is written into the NAND flash in the super capacitor for storage.

Super Capacitor Power Calibration

Because protection against a power supply failure requires a super capacitor, the RAID controller card needs to record the discharge curve and the maximum power of the super capacitor to learn about the status of the super capacitor. In addition, to extend the lifespan of the super capacitor, the automatic calibration mode of the super capacitor is enabled on the RAID controller card by default.

The RAID controller card needs to keep the battery level of the super capacitor at a relatively stable value. Therefore, the battery level of the super capacitor is calibrated through the three-phase charging/discharging operations. The three-phase charging/discharging operations are as follows:

1. Charges the super capacitor with maximum power.
2. Automatically starts the calibration process to completely discharge the super capacitor.
3. Recharges the super capacitor to the maximum power.

The super capacitor has the following characteristics during the calibration process:

- The write policy of the RAID controller card is automatically adjusted to direct write mode to ensure data security.
- The write performance of the RAID controller card is low.
- The power calibration period depends on the charging and discharging speeds of the super capacitor.

1.3.9 Disk Striping

Basic Concepts

Disk striping means that the disk space is divided into multiple stripes in accordance with the specified size and data blocks are also divided in accordance with the stripe size when data is written. The specific concepts include the following:

- Stripe width: number of disks used in a disk group for striping.
- Stripe size of a disk group: total amount of data written at the same time on all member disks in the disk group by the RAID controller card.
- Stripe size of a disk: amount of data written into a single disk.

For example, a RAID 0 group composed of four member disks divides the sequential data into 12 data blocks and writes them to each member disk in the way shown in Figure 1-9.

Figure 1-9 Disk Striping Example

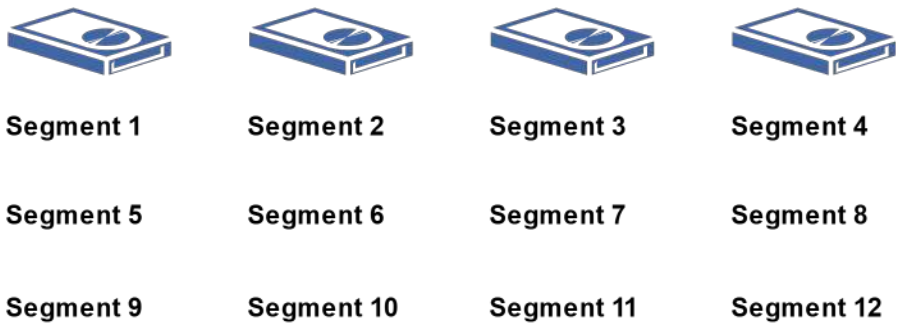

In the example RAID 0, suppose the size of each data block is 2 KB:

- The stripe width is the number of member disks, that is, 4.
- The stripe size of the disk group is the sum of 12 data blocks, that is, 24 KB.
- The stripe size of each disk is the sum of three data blocks, that is, 6 KB.

Feature

Because most disks have restrictions on the number of accesses (I/O operation per second) and data transmission rate (data volume per second), when multiple processes access a disk at the same time, the disk restriction may be triggered. Subsequent processes are suspended.

Striping is a technology used to automatically balance the I/O load across multiple physical disks. By dividing a piece of sequential data into multiple data blocks and writing them into different disks, multiple processes can access different parts of the data at the same time. In addition, when a disk needs sequential access to data, it can obtain the maximum I/O parallel capability.

Note

This feature does not ensure data redundancy.

1.3.10 Disk Mirroring

Disk mirroring means that the same data is written into two disks at the same time when a data write task is executed, thus achieving 100% data redundancy. Because the data on the two disks is exactly the same, when one disk is faulty, the data will not be lost, and the system automatically switches from the faulty disk to the mirrored disk for data read and write.

Disk mirroring is applicable to RAID 1 and RAID 10. It brings a complete redundancy of 100%, but it is costly and the actual disk usage is only 50%, because each disk needs a backup disk during the mirroring process, see Figure 1-10.

Figure 1-10 Disk Mirroring Example

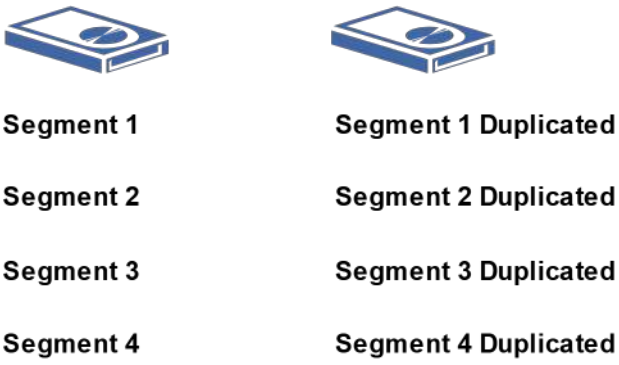

1.3.11 Foreign Configuration

The foreign configuration is different from the configuration of the current **RAID** controller card, and is usually displayed as **Foreign Configuration** on the configuration screen.

In the following cases, the foreign configuration exists:

- The RAID configuration information exists in a physical disk newly installed on a server, and the RAID controller card identifies such information as a foreign configuration.
- After the RAID controller card of a server is replaced, the new RAID controller card identifies the existing RAID information as foreign configuration.
- After a member disk of a RAID group is hot swapped, the member disk is marked as carrying foreign configuration.

The detected foreign configuration can be processed in accordance with the actual server conditions:

- If the RAID information carried by the newly inserted disk does not meet the requirements of the current scenario, the configuration can be deleted.
- After the RAID controller card is replaced, if you want to use the previous configuration, you can import that configuration to apply it on the new RAID controller card.

1.3.12 Disk Energy Saving

A RAID controller card has the disk energy saving function, and can control disk rotation in accordance with the disk configurations and I/O activities. All rotary SAS and SATA disks support this function.

When the disk energy saving function is enabled, both idle disks and idle hot spare disks mounted under the RAID controller card are in energy saving status. When related operations (for example, creating a RAID volume, creating a hot spare disk, expanding a disk dynamically, and rebuilding a hot spare disk) are performed, the disk in energy saving status can be woken up.

1.3.13 Disk Pass-Through

Disk pass-through, namely, the JBOD function, also called transparent command transmission, is a data transmission mode that only ensures transmission quality without processing the command by the transmission device.

After the disk pass-through function is enabled, the RAID controller card can transparently transmit commands to the connected disk. If no RAID group is configured, user commands can be transparently transmitted to a disk, so that upper-layer service software or management software can access the disk.

For example, during the installation of the server operating system:

- For a RAID controller card that supports the disk pass-through function, you can use a disk mounted under the RAID controller card as the installation disk.
- For a RAID controller card that does not support the disk pass-through function, you can use only the virtual disk configured under the RAID controller card as the installation disk.

Chapter 2
Applicable Server Models

This document is applicable to VANTAGEO rack servers based on the **Eagle Stream** platform, including:

- 1240-RE
- 2240-RE
- 4440-RE

Chapter 3
VT SmartIOC 2100 RAID Controller Card

Table of Contents

A VT SmartROC 2100 RAID controller card is integrated with the BIOS configuration utility. By using the program, you can configure the disks that are supported by the VT SmartROC 2100 RAID controller card to a RAID volume of a specific level.

Note

The operation screens of the BIOS configuration utility displayed in this chapter are for reference only, and may not be the same as the actual ones.

3.1 Capability Features

For a description of the capabilities of a VT SmartIOC 2100 RAID controller card, refer to Table 3-1.

Table 3-1 Descriptions of the Capabilities of a VT SmartIOC 2100 RAID Controller Card

Capability Item	Capability Parameter
Product form	Mezz card
Controller chip	PMC PM8238
Host interface	PCIe 3.0x8
SAS interface	12 Gb SAS

Capability Item	Capability Parameter
Number of ports	16+2
Drive interface	SAS and SATA
Drive type	HDD and SSD
Whether drives are hot swappable	Supported
Maximum number of RAID groups	8
Number of drives	238
RAID level	RAID 0, RAID 1, RAID 10, RAID 5
JBOD mode	Supported
Cache	None
Cache protection	None
Out-of-band management	Supported
Consistency check/verification and fix	Supported
Online capacity expansion	Supported
Online RAID level migration	Supported
Automatic rebuild	Supported
Manufacturer tool support	arcconf

3.2 Initial Configuration (Legacy Mode)

Figure 3-1 shows the initial configuration flow of a VT SmartIOC 2100 RAID controller card.

Figure 3-1 Initial Configuration Flow of a VT SmartIOC 2100 RAID Controller Card

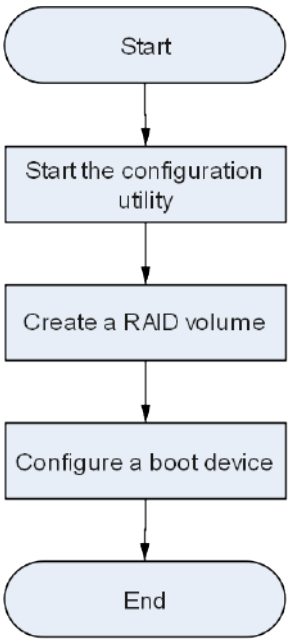

3.2.1 Starting the Configuration Utility

Abstract

This procedure describes how to start the BIOS configuration utility of a VT SmartIOC 2100 RAID controller card to log in to the management screen and complete the subsequent initial and common configurations.

Prerequisite

The boot mode is already set to **Legacy** in BIOS. For details, refer to "5.3.1 Setting the Boot Mode to Legacy".

Steps

1. Start the server system.
2. During the POST process, press **Ctrl+A** to start the BIOS configuration utility of the VT SmartIOC 2100 RAID controller card. The screen as shown in Figure 3-2 is displayed.

Figure 3-2 BIOS Configuration Utility Screen

3. In the **Options** area, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Configuration Menu** screen is displayed, see Figure 3-3.

Figure 3-3 Configuration Menu Screen

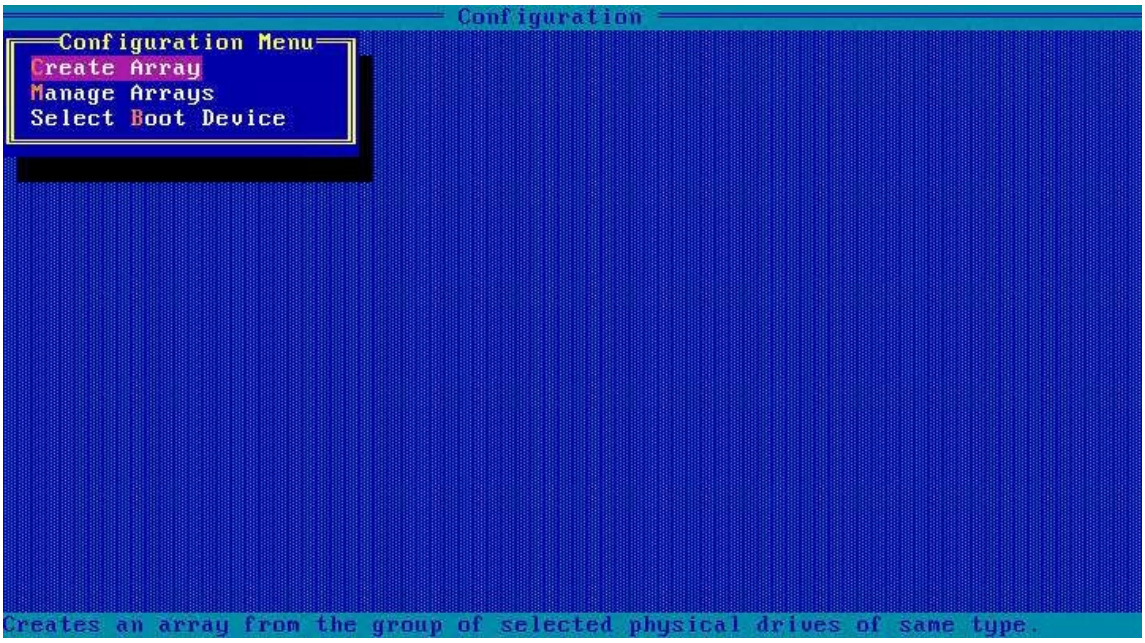

3.2.2 Creating a RAID Volume

Abstract

You can create RAID volumes at different levels as required.

The procedures for creating RAID volumes at different levels are similar. This procedure uses a RAID 0 volume in legacy mode as an example.

Prerequisite

Sufficient SATA and SAS disks are installed on the server.

Context

For a description of the number of disks required to create a RAID volume, refer to Table 3-2.

Table 3-2 Number of Disks Required for Creating a RAID Volume

RAID Level	Description
RAID 0	RAID 0 requires at least one disks.
RAID 1	RAID 1 requires at least two disks. Disks with different capacities can be used in a RAID 1 volume, but the logical capacity of each member disk depends on the space of the disk with the smallest capacity.
RAID 5	RAID 5 requires at least three disks.
RAID 6	RAID 6 requires at least four disks.
RAID 10	RAID 10 requires at least four disks. A RAID 10 volume consists of at least two RAID 1 volumes. For example, if there are four disks to be used in RAID 10 mode, you need to add them to two drive groups, each of which is mounted with two disks in RAID 1 mode.
RAID 50	RAID 50 requires at least six disks. A RAID 50 volume consists of at least two RAID 5 volumes. For example, if there are six disks to be used in RAID 50 mode, you need to add them to two drive groups, each of which is mounted with three disks in RAID 5 mode.
RAID 60	RAID 60 requires at least eight disks. A RAID 60 volume consists of at least two RAID 6 volumes. For example, if there are eight disks to be used in RAID 60 mode, you need to add them to two drive groups, each of which is mounted with four disks in RAID 6 mode.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Create Array**, and then press **Enter**. On the displayed **Select drives to create Array** screen, all the disks that can be used to create a RAID volume are displayed, see Figure 3-4.

Figure 3-4 Select Drives to Create Array Screen

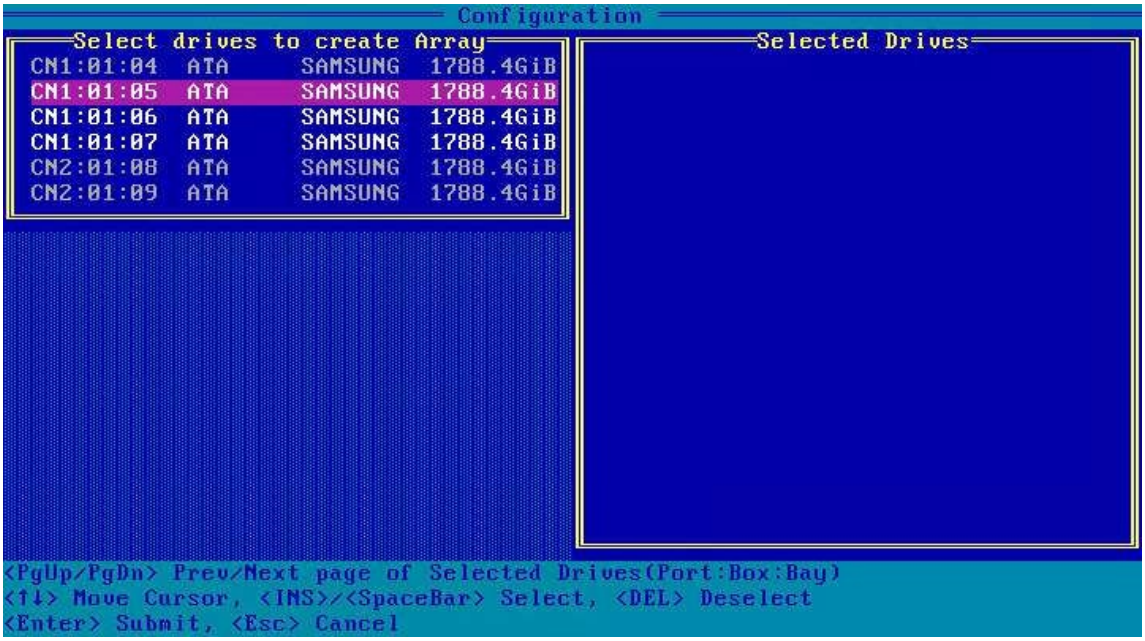

2. Use the arrow keys to select the disks to be used to create a RAID volume, and then press **Insert** to add these disks to the **Selected Drives** list, see Figure 3-5.

Figure 3-5 Selected Drives List

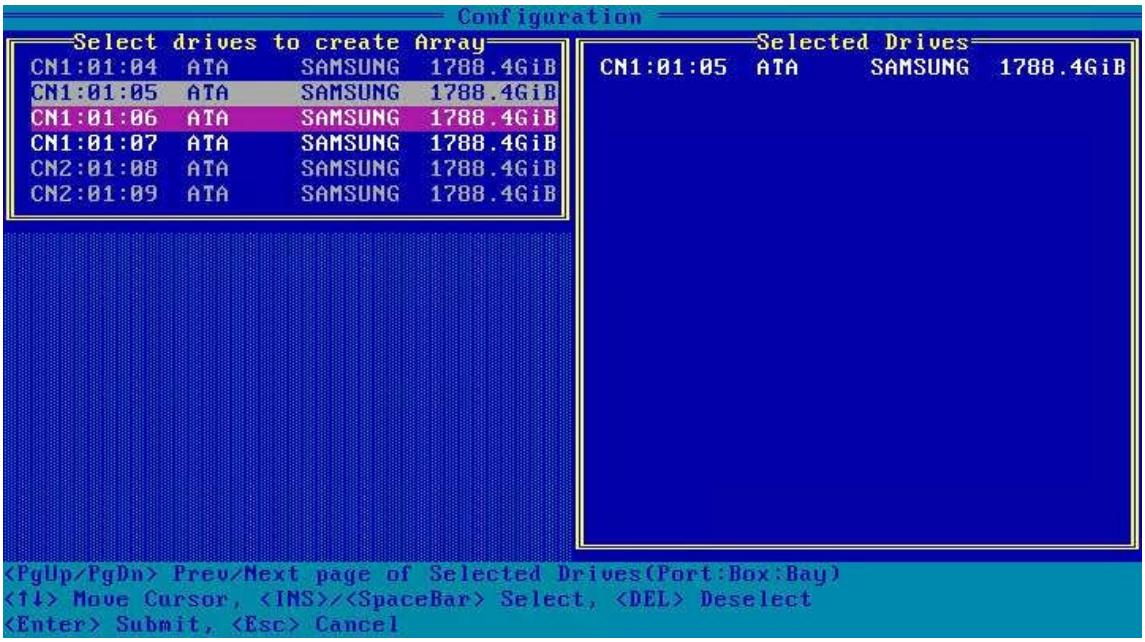

Note

- The disks for creating a RAID volume must be of the same type. It is forbidden to select disks with interface types such as SATA and SAS at the same time.
- Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

3. Press **Enter**. The **Create Logical Drive** screen is displayed, see Figure 3-6.

Figure 3-6 Create Logical Drive Screen

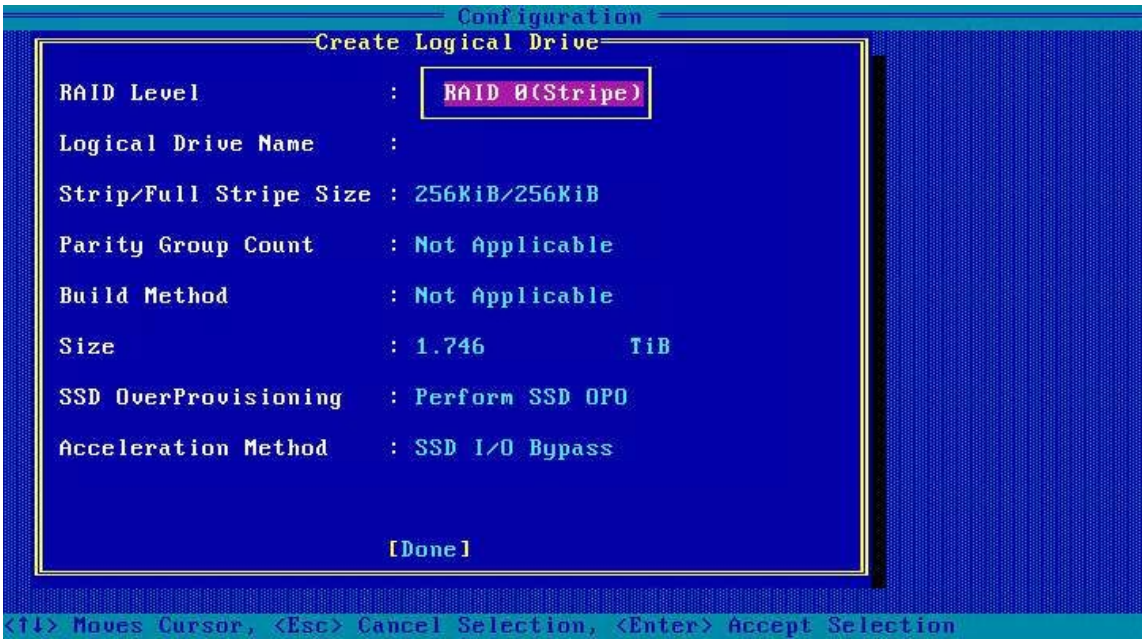

4. Use **Tab**/**Tab**+**Shift** to select the parameters that you want to modify. In the displayed operation box, use the arrow keys to select the related parameters, and then press **Enter** for confirmation. For a description of the parameters on the **Create Logical Drive** screen, refer to Table 3-3.

Table 3-3 Descriptions of Parameters on the Create Logical Drive Screen

Parameter	Description
RAID Level	Sets a RAID standard, for example, **RAID 0(Stripe)**.
Logical Drive Name	Sets the RAID volume name, for example, "0".
Strip/Full Stripe Size	The stripe size should be equal to the size of average disk IO requests generated by server applications. In the optimum status, only one IO operation is executed for each IO request. The recommended stripe size configurations are as follows: ● For a Web server, 8 KB is recommended. ● For a groupware server (such as an email server), 16 KB is recommended. ● For a database server, 16 KB or 32 KB is recommended. ● For a file server, 32 KB or 64 KB is recommended ● For a video file server, 64 KB, 128 KB, or 256 KB is recommended.
Parity Group Count	Configures logical-device parity groups in accordance with the number of physical devices in the array. It is not applicable to all RAID levels.

Parameter	Description
Build Method	Sets the RAID initialization method, which is used to determine how the logical devices prepare for read and write, and how long the initialization takes. ● default: When the logical devices can be accessed by the operating system, parity blocks are initialized at the back end. A lower RAID level can achieve faster parity initialization. ● RPI: The data and parity blocks at the front end are overwritten. Before the parity initialization procedure is completed, logical devices remain invisible and unavailable to the operating system. All parity groups are initialized in parallel, but the initialization of a single parity group (RAID 5) is faster. The RAID level does not affect the performance during the RAID initialization. ● Not Applicable: unavailable. Keep the default configuration unless otherwise specified.
Size	Displays the storage size of the disk array in accordance with the total storage space of the disks added to the disk array. By default, the RAID created uses all the available disk space.
SSD Over Provisioning Optimization	Sets whether to enable the SSD over-provisioning function to perform excessive-configuration optimization on the RAID group when creating a RAID volume composed of SSDs.
Acceleration Method	Sets the caching mode for the RAID volume. ● IO Bypass: This option is valid only when the RAID logical volume is formed by SSDs. ● Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time. ● None: disables the controller cache. Neither **IO Bypass** nor **Controller Cache** is used. Keep the default configuration unless otherwise specified.

5. Use **Tab** to select **Done**, and then press **Enter** to create the RAID volume, see Figure 3-7.

Figure 3-7 Creating a RAID Volume

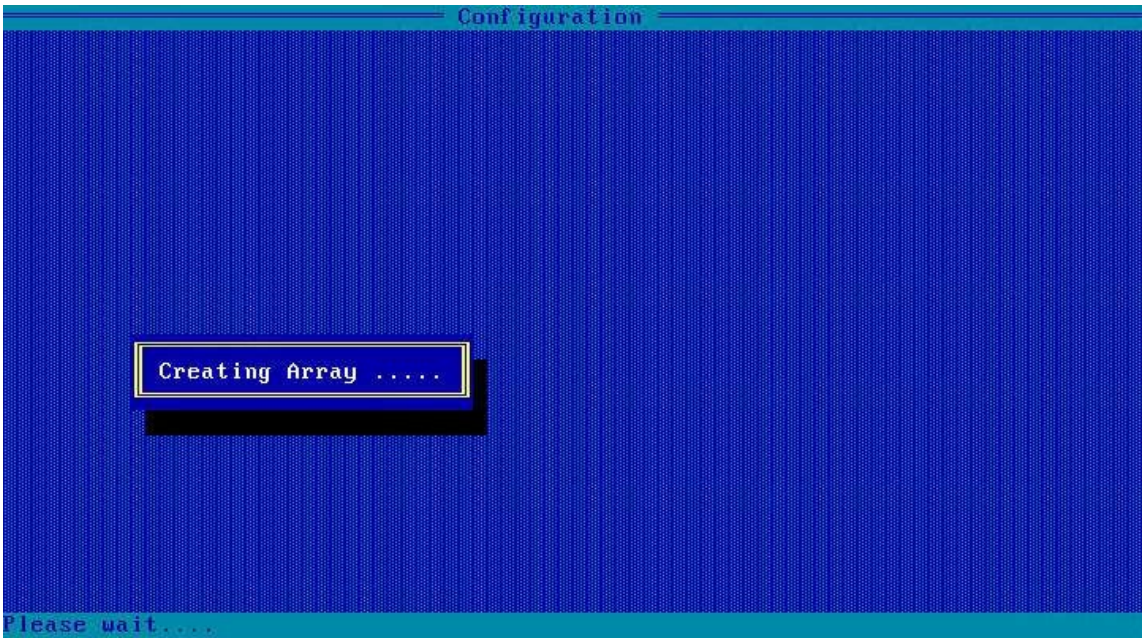

6. The RAID volume is created successfully, see Figure 3-8. Press any key to return to the **Configuration Menu** screen.

Figure 3-8 RAID Volume Created Successfully

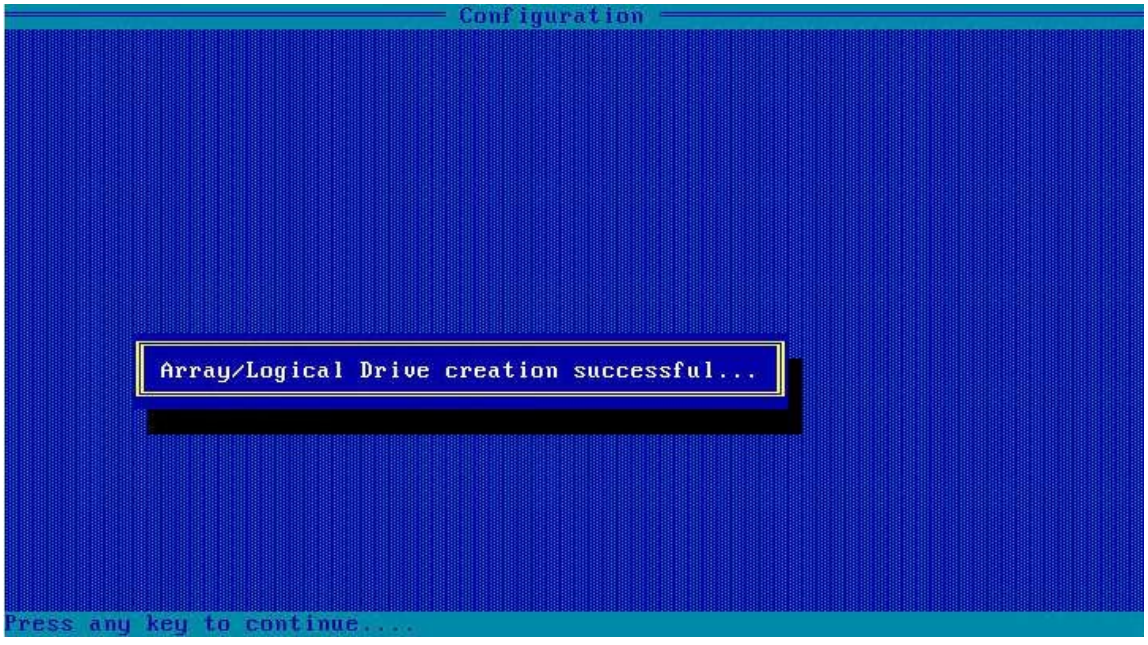

3.2.3 Configuring a Boot Device

Abstract

After a RAID volume is created, if you need to install an operating system on the RAID volume, and there are multiple RAID volumes on the RAID controller card, you must set the RAID volume as a boot device.

Prerequisite

A RAID volume is created successfully. For details, refer to “3.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** page is displayed, see Figure 3-9.

Figure 3-9 List of Arrays Screen

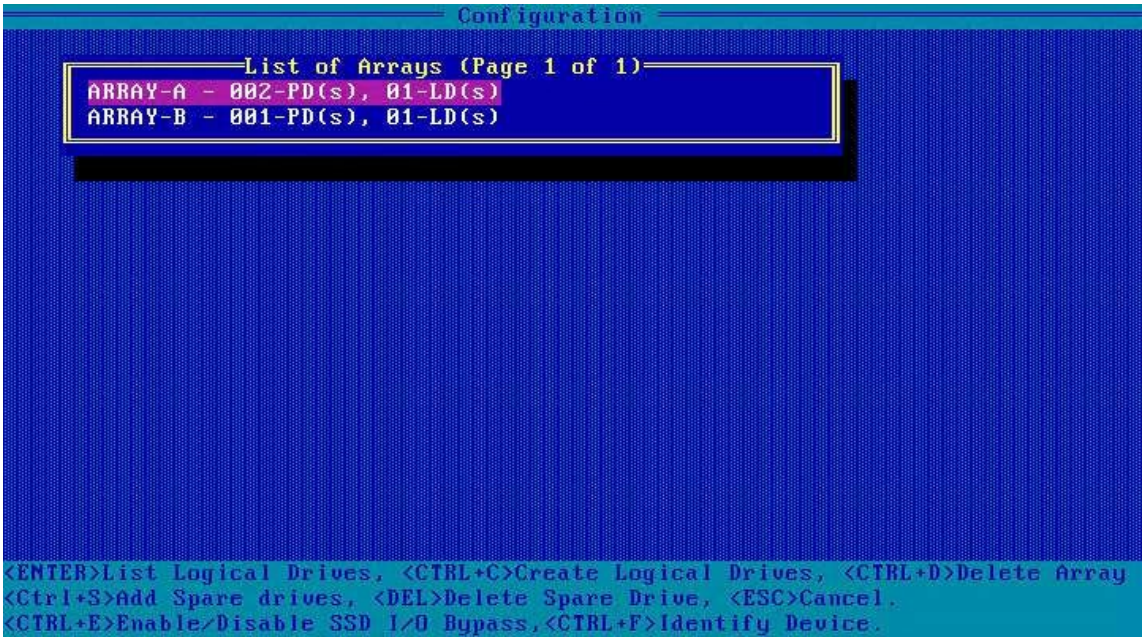

2. Use the arrow keys to select the array where the RAID volume to be set as a boot device is located, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 3-10.

Figure 3-10 List of Logical Drives Screen

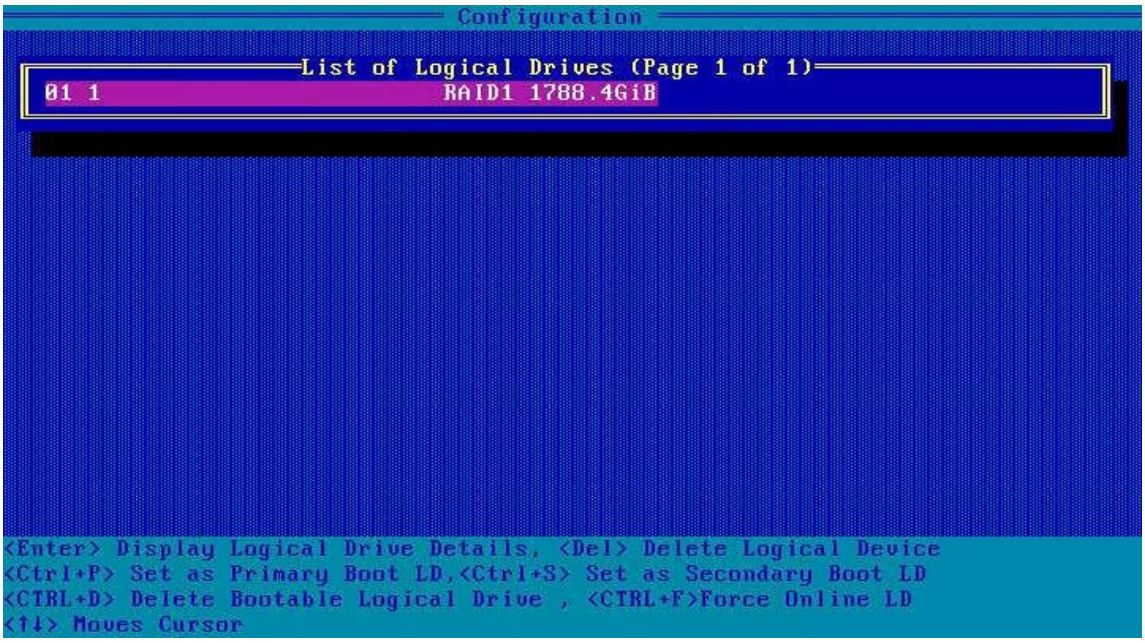

3. Select the RAID volume to be booted first, and then press **Ctrl+P** to configure it as the first boot device, see Figure 3-11.

Figure 3-11 Configuring a Boot Device

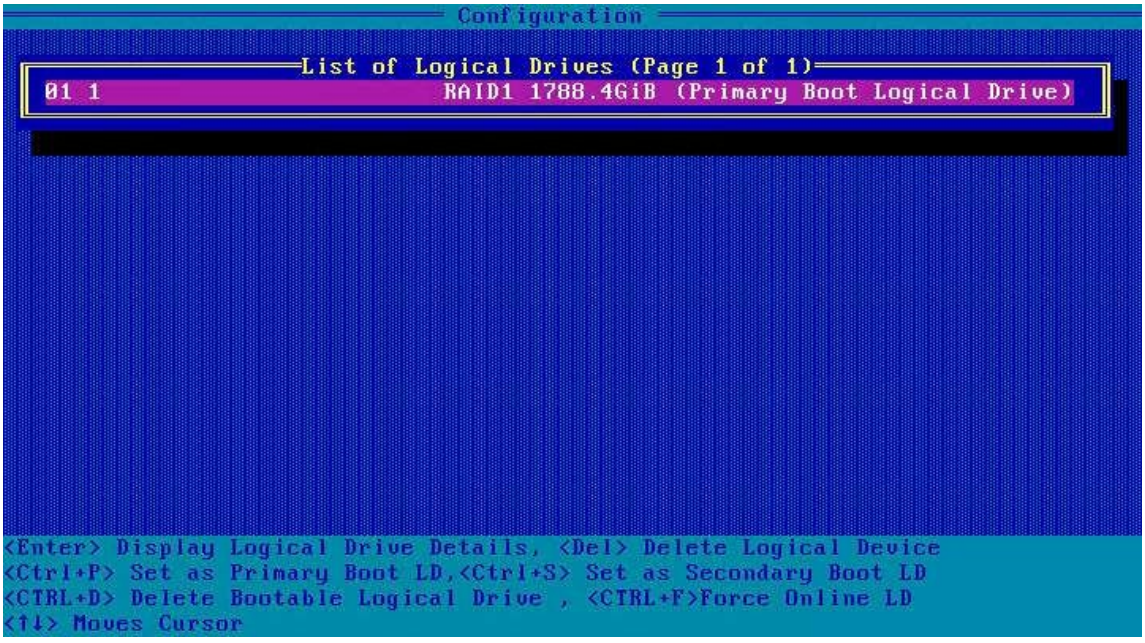

3.3 Initial Configuration (UEFI Mode)

Figure 3-12 shows the initial configuration flow of a VT SmartIOC 2100 RAID controller card.

Figure 3-12 Initial Configuration Flow of a VT SmartIOC 2100 RAID Controller Card

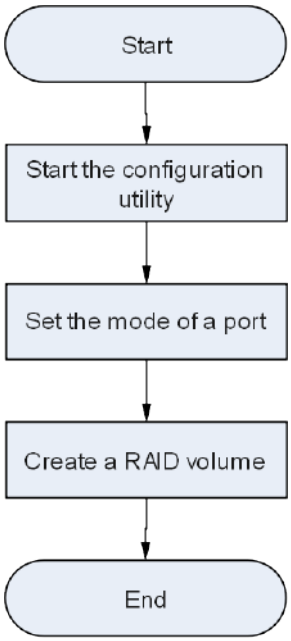

Note

- The SmartIOC 2100 RAID controller card does not support the boot disk configuration in UEFI mode.
- The SmartIOC 2100 RAID controller card does not support the co-existence of configurations in both UEFI and legacy modes. If the mode is switched from UEFI to legacy, the configuration in UEFI mode must be cleared. Otherwise, the normal operation of the RAID controller card is affected.

3.3.1 Starting the Configuration Utility

Abstract

This procedure describes how to start the BIOS configuration utility of a VT SmartIOC 2100 RAID controller card to log in to the management screen and complete the subsequent initial and common configurations.

Prerequisite

The boot mode is already set to UEFI in BIOS. For details, refer to "5.3.2 Setting the Boot Mode to UEFI".

Steps

1. Start the server system.
2. During the POST process, press **F2**/**DEL**. The **Aptio Setup** screen is displayed, see Figure 3-13.

Figure 3-13 Aptio Setup Screen

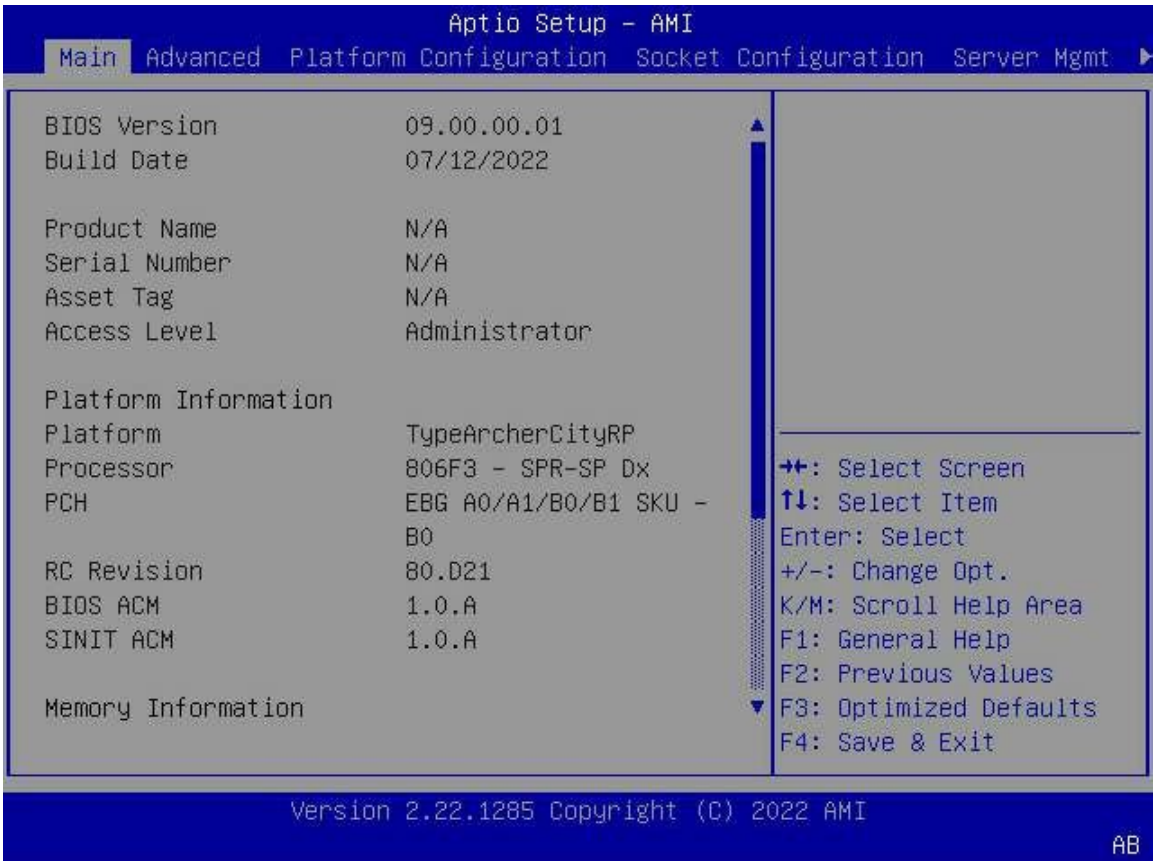

3. Use the arrow keys to select **Advanced**, and then press **Enter**. The **Advanced** screen is displayed, see Figure 3-14.

Figure 3-14 Advanced Screen

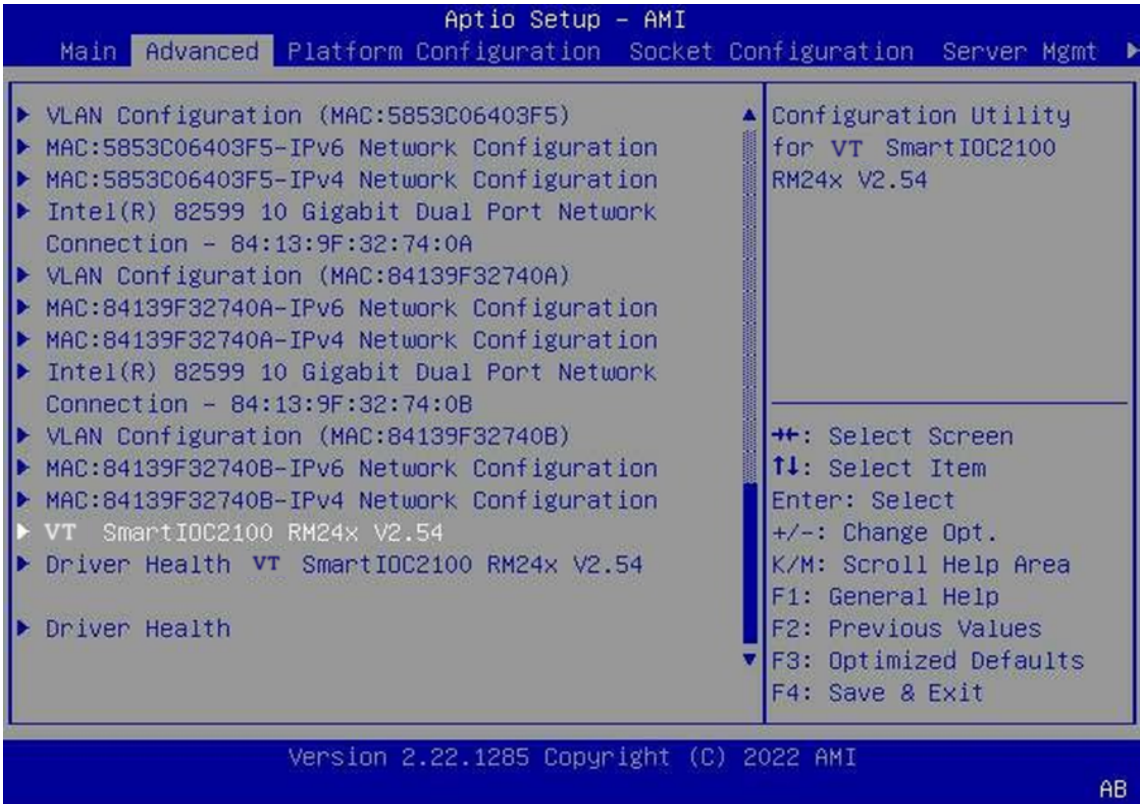

4. Use the arrow keys to select **VT SmartIOC2100 RM24x V2.54**, and press **Enter**. The controller management screen is displayed, see Figure 3-15.

Figure 3-15 Controller Management Screen

For a description of the functions of the menus on the controller management screen, refer to Table 3-4.

Table 3-4 Functions of Menus on the Controller Management Screen

Menu	Function Description
Controller Information	Displays the basic information, firmware, current temperature, and port configuration of the controller.
Configure Controller Settings	Provides advanced configuration options for the controller.
Array Configuration	Creates an array or RAID.
Disk Utilities	Displays the list of disk devices mounted under the controller as well as the basic disk information. It allows you to turn on the disk location indicator, erase disk data and upgrade the firmware.
Set Bootable Device(s) for Legacy Boot Mode	Configures, or clears the primary and secondary boot disks.
Administration	Allows the controller administrator to perform operations, such as upgrading the firmware and restoring factory defaults.

3.3.2 Creating a RAID Volume

Abstract

You can create RAID volumes at different levels as required.

The procedures for creating RAID volumes at different levels are similar. This procedure uses a RAID 1 volume in UEFI mode as an example.

Prerequisite

- Sufficient SATA and SAS disks are installed on the server.
- The port mode for the disks to be connected is already set. For details, refer to "3.5.1 Setting the Mode of a Port".

Context

For a description of the number of disks required to create a RAID volume, refer to Table 3-5.

Table 3-5 Number of Disks Required for Creating a RAID Volume

RAID Level	Description
RAID 0	RAID 0 requires at least one disks.
RAID 1	RAID 1 requires at least two disks. Disks with different capacities can be used in a RAID 1 volume, but the logical capacity of each member disk depends on the space of the disk with the smallest capacity.
RAID 5	RAID 5 requires at least three disks.
RAID 6	RAID 6 requires at least four disks.
RAID 1+0	RAID 1+0 requires at least four disks. A RAID 1+0 volume consists of at least two RAID 1 volumes. For example, if there are four disks to be used in RAID 1 +0 mode, you need to add them to two drive groups, each of which is mounted with two disks in RAID 1 mode.
RAID 50	RAID 50 requires at least six disks. A RAID 50 volume consists of at least two RAID 5 volumes. For example, if there are six disks to be used in RAID 50 mode, you need to add them to two drive groups, each of which is mounted with three disks in RAID 5 mode.
RAID 60	RAID 60 requires at least eight disks. A RAID 60 volume consists of at least two RAID 6 volumes. For example, if there are eight disks to be used in RAID 60 mode, you need to add them to two drive groups, each of which is mounted with four disks in RAID 6 mode.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-16.

Figure 3-16 Array Configuration Screen

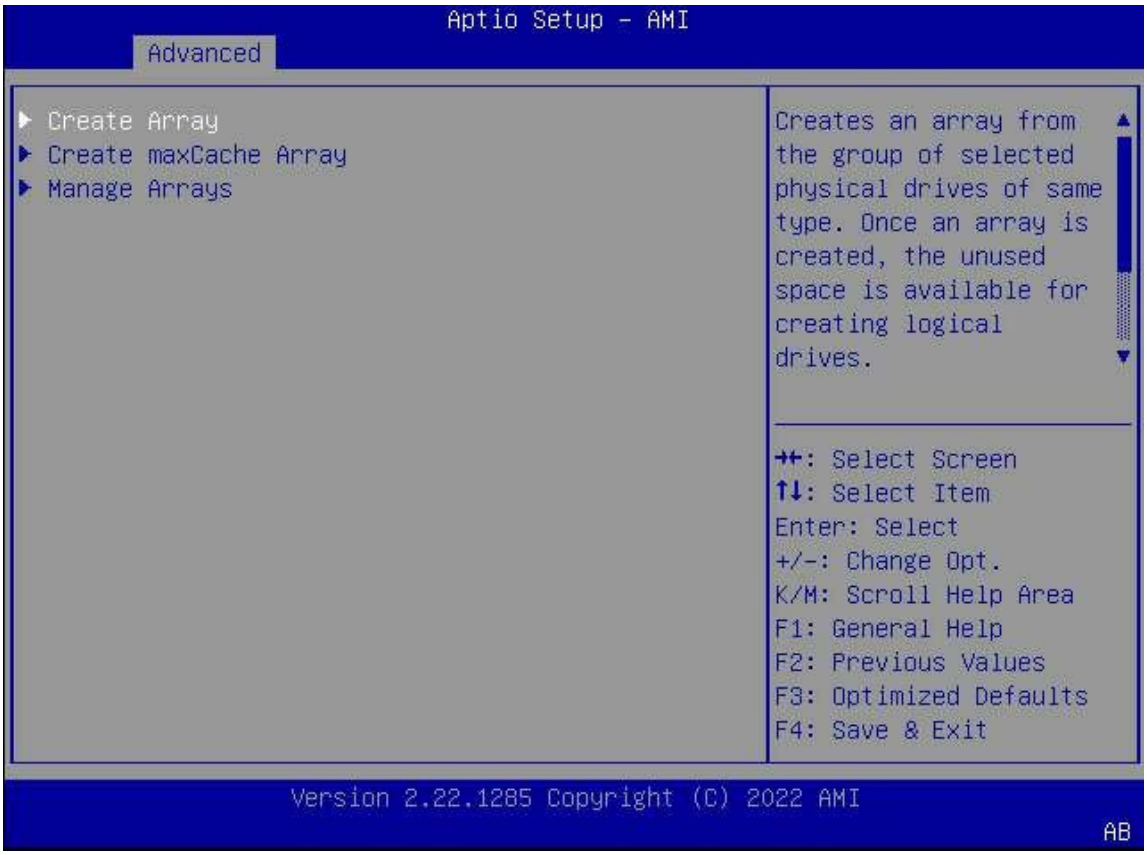

2. Use the arrow keys to select **Create Array**, and then press **Enter**. In the displayed disk list, all the disks that can be used to create a RAID volume are displayed, see Figure 3-17.

Figure 3-17 Viewing the Disk List

3. Select the disk to be added to the Array disk group, and then press **Enter** to set the disk port to **Enabled** status, see Figure 3-18.

Figure 3-18 Confirming the Configuration

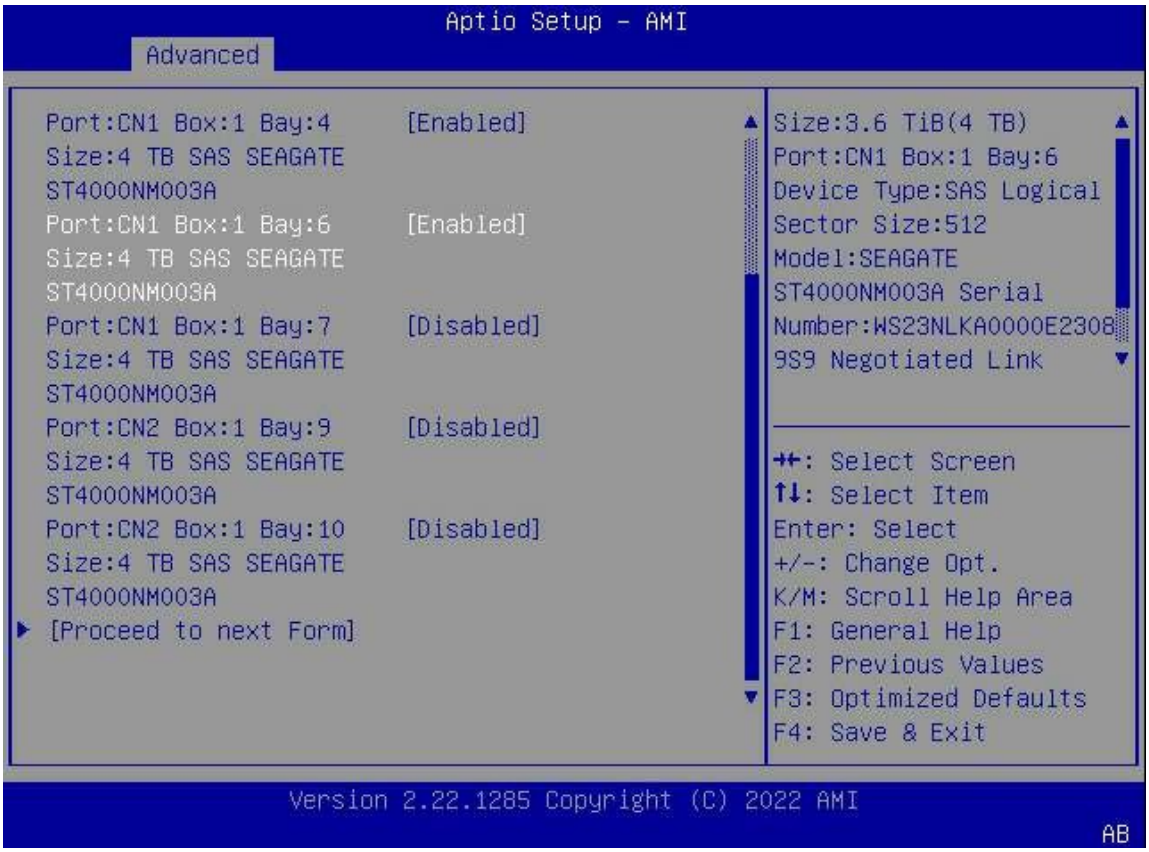

Note

The disks for creating a RAID volume must be of the same type. It is forbidden to select disks with interface types such as SATA and SAS at the same time.

4. Use the arrow keys to select **Proceed to Next Form**, and then press **Enter**. The screen for creating RAID is displayed, see Figure 3-19.

Figure 3-19 Creating RAID

5. Use the arrow keys to select **RAID Level**, and then press **Enter**. From the displayed short-cut menu, select the desired RAID level, see Figure 3-20.

Figure 3-20 Selecting a RAID Level

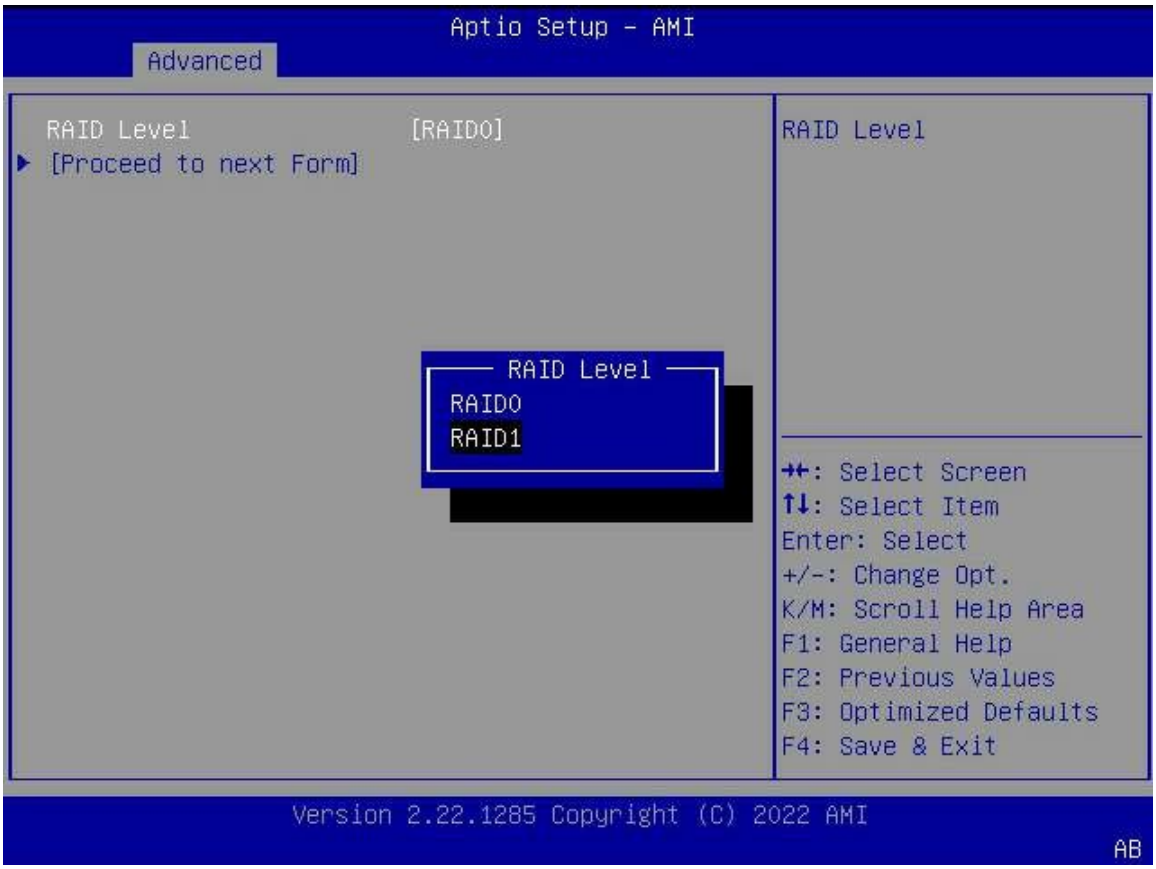

6. Use the arrow keys to select **Proceed to Next Form**, and press **Enter**. The screen for setting RAID logical volumes is displayed, see Figure 3-21.

Figure 3-21 Setting RAID Logical Volumes

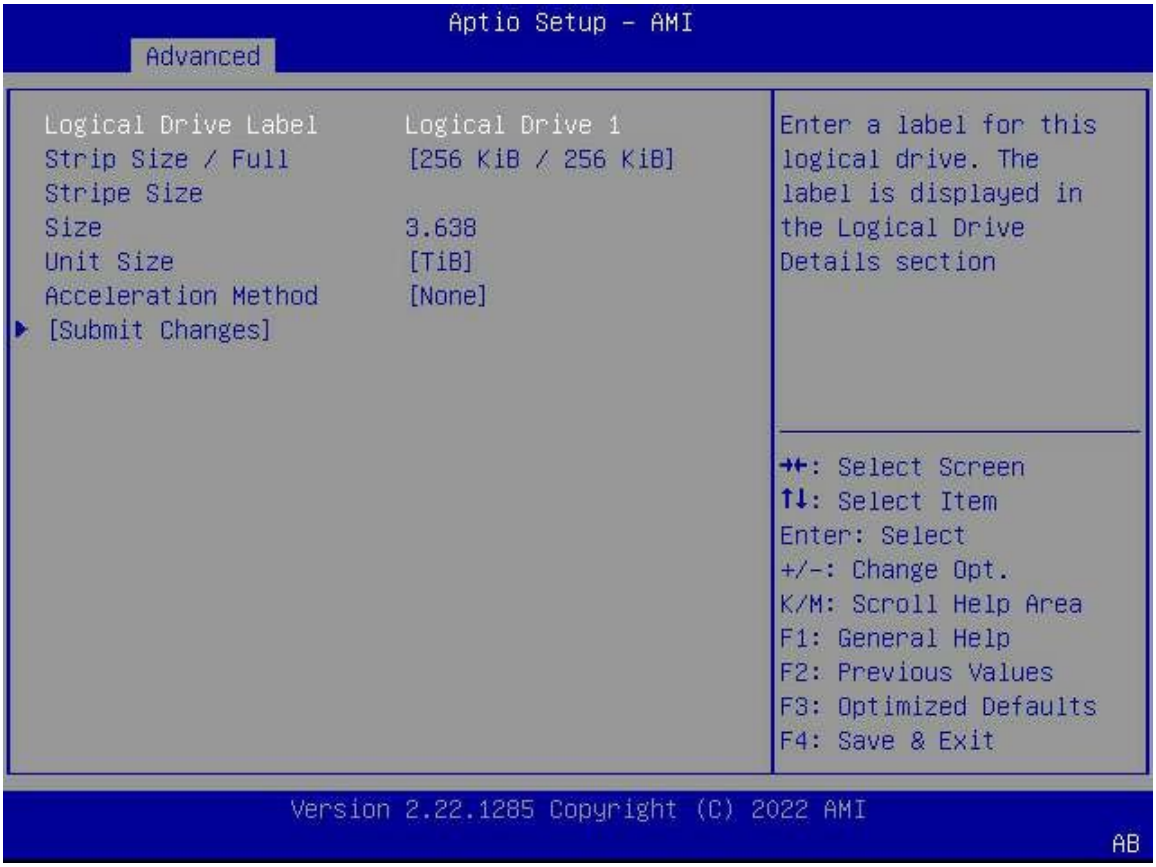

7. Use the arrow keys to select the parameters that you want to modify. In the displayed operation boxes, configure the related parameters. For a description of the parameters for setting a logical volume, refer to Table 3-6.

Table 3-6 Parameter Descriptions for Logical Volume Configuration

Parameter	Description
Logical Drive Label	Sets the name of the RAID logical volume, for example, "Logical Drive 1".
Strip Size/Full Strip Size	● Strip Size indicates the size of the current stripe. The stripe size should be equal to the size of average disk IO requests generated by server applications. In the optimum status, only one IO operation is executed for each IO request. The size can be 16 KiB, 32 KiB, 64 KiB, 128KiB, 512 KiB, or 1024 KiB. The default value is 256 KiB. The recommended stripe size configurations are as follows: → For a Web server, 8 KiB is recommended. → For a groupware server (such as an email server), 16 KiB is recommended. → For a database server, 16 KiB or 32 KiB is recommended. → For a file server, 32 KiB or 64 KiB is recommended. → For a video file server, 64 KiB, 128 KiB, or 256 KiB is recommended.

Parameter	Description
	● Full Strip Size indicates the total size of all stripes. When you set **Strip Size**, the system automatically calculates **Full Stripe Size**.
Size	Displays the storage size of the RAID volume in accordance with the total storage space of the disks added to the RAID volume. By default, all available space is used to create a RAID logical volume. To create multiple RAID logical volumes, you can define the size of the volumes.
Unit Size	Select the unit (MiB/GiB/TiB) of the logical drive.
Acceleration Method	Sets the caching mode for the RAID volume. ● IO Bypass: This option is valid only when the RAID logical volume is formed by SSDs. ● Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time. ● None: disables the controller cache. Neither **IO Bypass** nor **Controller Cache** is used. Keep the default configuration unless otherwise specified.

8. Use the arrow keys to select **Submit Changes**, and press **Enter**. The RAID volume is created successfully, see Figure 3-22.

Figure 3-22 RAID Volume Created Successfully

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.4 Common Configurations (Legacy Mode)

By using the BIOS configuration utility, you can configure and maintain a created RAID volume. For a description of the common operations on a SmartIOC 2100 RAID controller card in legacy mode, refer to Table 3-7.

Table 3-7 Common Operations on a SmartIOC 2100 RAID Controller Card

Common Operation	Description
Querying RAID volume information	Refer to "3.4.1 Querying RAID Volume Information".
Creating a hot spare disk	Refer to "3.4.2 Creating a Hot Spare Disk".
Deleting a RAID volume	Refer to "3.4.3 Deleting a RAID Volume".
Deleting a hot spare disk	Refer to "3.4.4 Deleting a Hot Spare Disk".
Locating a disk	Refer to "3.4.5 Locating a Disk".

Common Operation	Description
Configuring a pass-through disk	Refer to "3.4.6 Configuring a Pass-Through Disk".

3.4.1 Querying RAID Volume Information

Abstract

This procedure describes how to query the RAID information created on a SmartIOC 2100 RAID controller card, such as the RAID volume status and member disk status.

Prerequisite

A RAID volume is created successfully. For details, refer to “3.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 3-23.

Figure 3-23 List of Arrays Screen

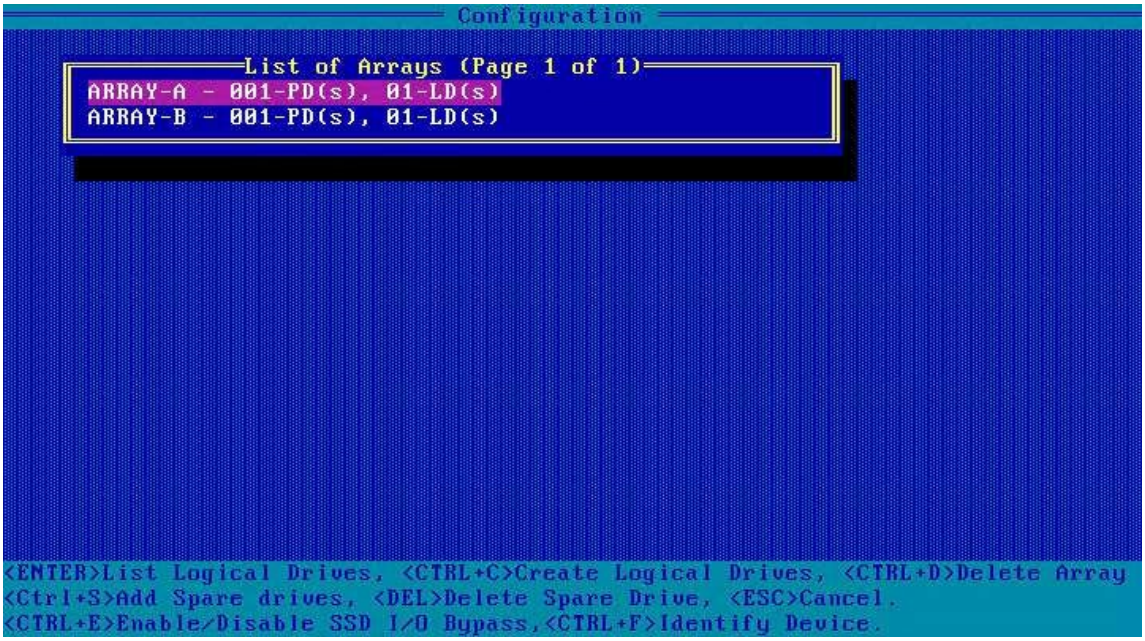

2. Select the RAID volume whose properties you want to view, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 3-24.

Figure 3-24 List of Logical Drives Screen

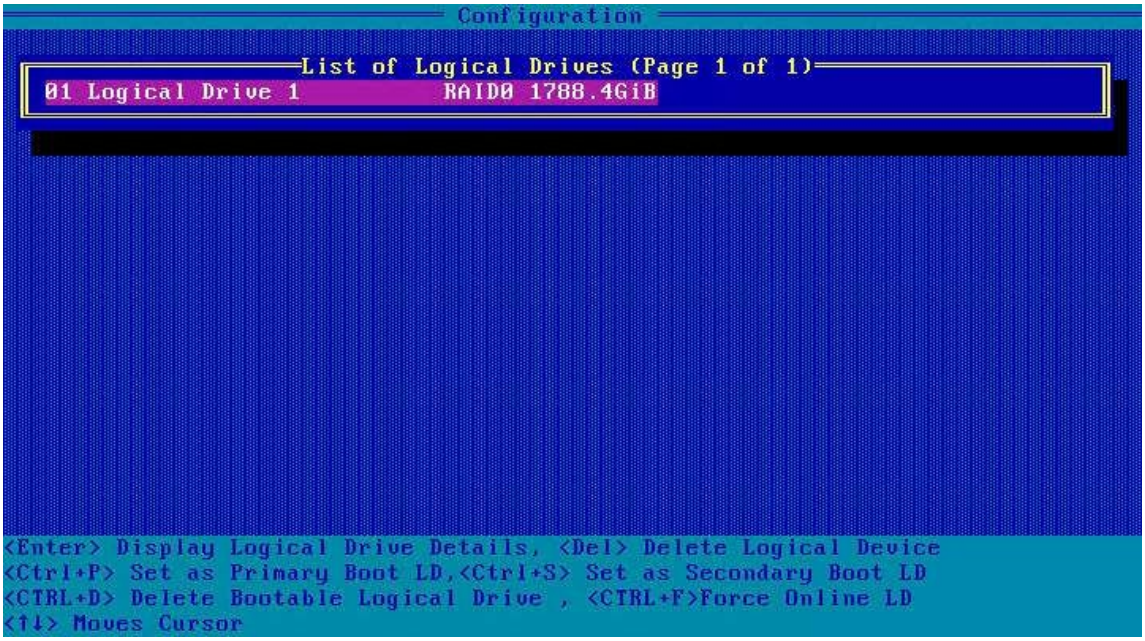

3. Press **Enter**. In the displayed **Logical Drive Details** dialog box, view the property information about the RAID volume, see Figure 3-25.

Figure 3-25 Logical Drive Details Dialog Box

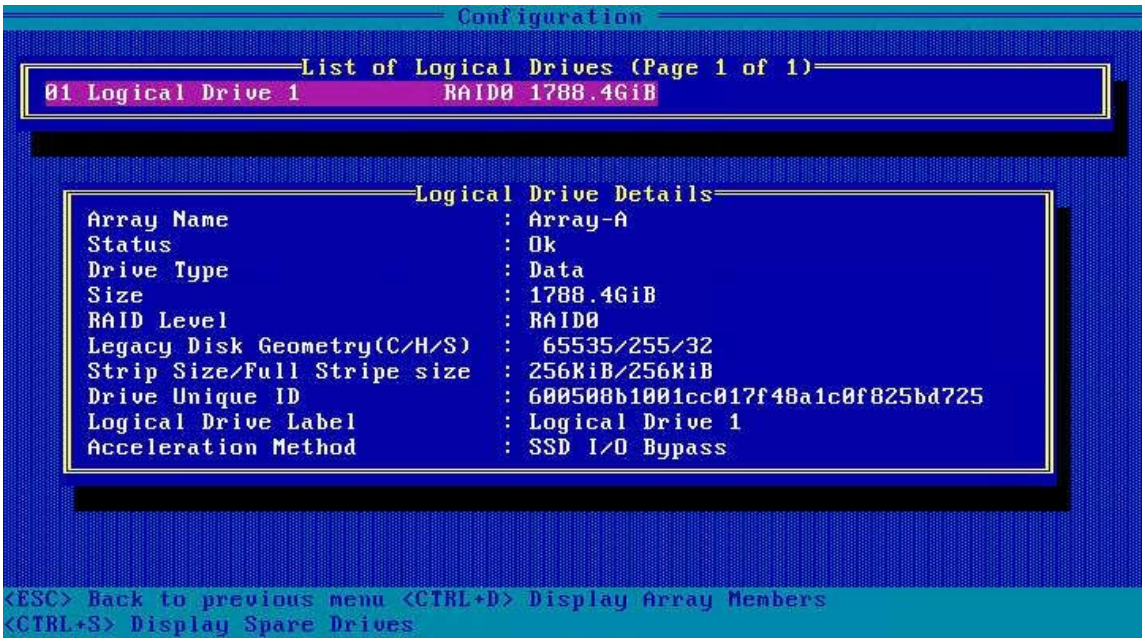

4. (Optional) To view the member disk information, press **Ctrl+D**. The **Array Member Drives** dialog box is displayed, see Figure 3-26.

Figure 3-26 Array Member Drives Dialog Box

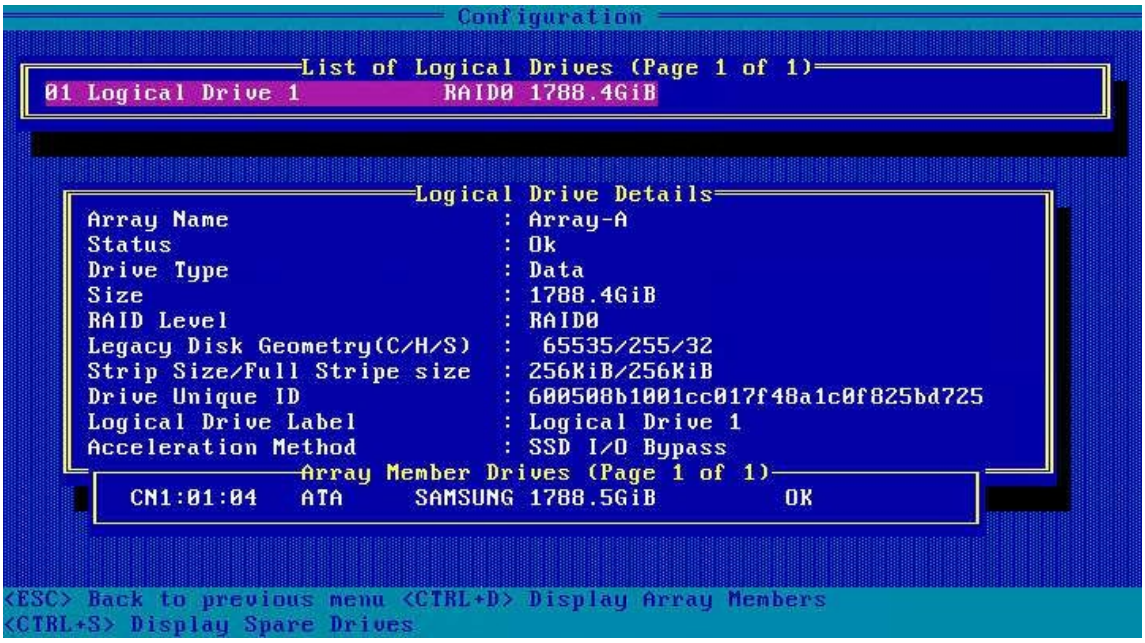

3.4.2 Creating a Hot Spare Disk

Abstract

A hot spare disk improves the data security of a RAID array. For a description of the hot spare disk types supported by a SmartIOC 2100 RAID controller card, refer to Table 3-8.

Table 3-8 Hot Spare Disk Types

Type	Description
Dedicated	● This type of hot spare disks is exclusive to the specified one or more disk groups of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a dedicated hot spare disk temporarily takes over the faulty disk.
Auto Replace	● This type of hot spare disks provides the hot standby function for a disk group of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a hot spare disk of this type automatically replaces the faulty disk.

Prerequisite

There are sufficient idle disks on the server.

Context

When creating a hot spare disk, pay attention to the following points:

- Multiple hot spare disks can be created for a disk group, but only one type of hot spare disk can be set at a time. That is, either **Dedicated** or **Auto Replace** is specified.
- An idle disk can be set as a hot spare disk. The disk that has been used to create a RAID volume cannot be set as a hot spare disk.
- The hot spare disk must be of the same type as that of any member disk in the corresponding disk group. That is, all of them are SATA disks or SAS disks, and the hot spare disk's capacity must not be less than the maximum capacity of the member disks.
- Disk groups at all levels except RAID 0 support hot spare disks.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 3-27.

Figure 3-27 List of Arrays Screen

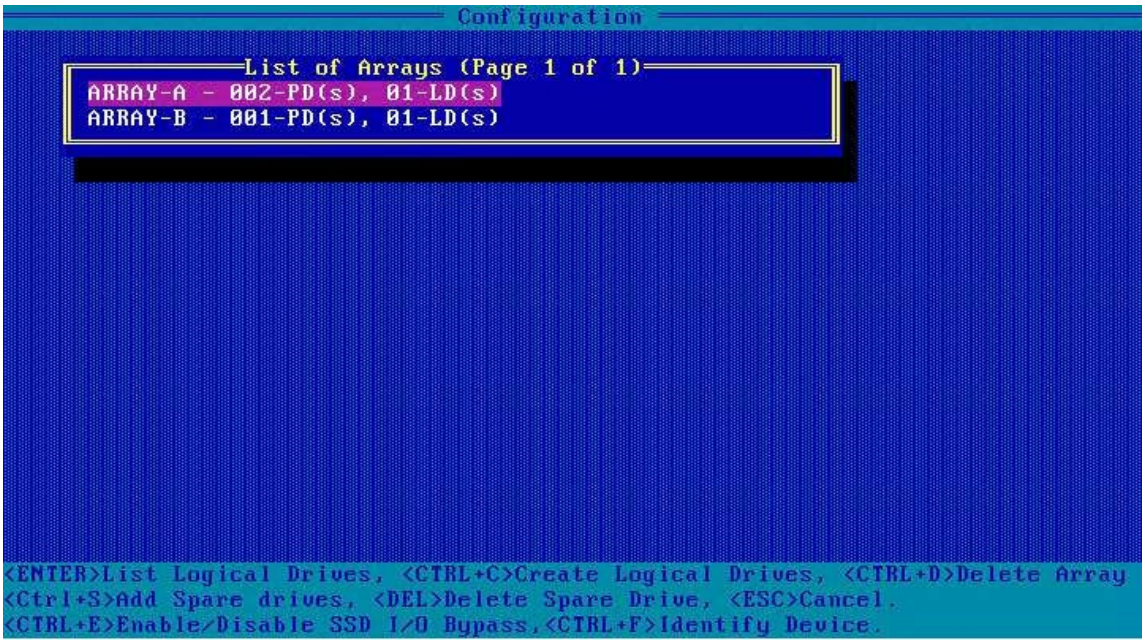

2. Use the arrow keys to select the RAID volume for which you want to configure a hot spare disk, and press **Ctrl+S**. The **Select Hotspare Drives** screen is displayed, see Figure 3-28.

Figure 3-28 Select Hotspare Drives Screen

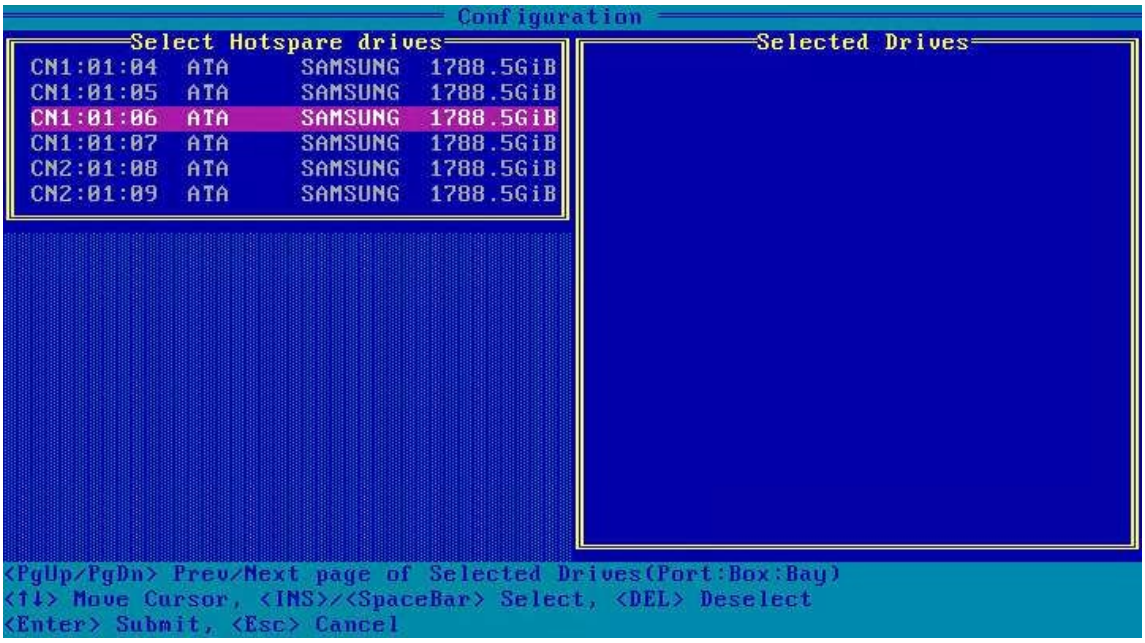

3. Use the arrow keys to select the idle disk to be set as a hot spare disk, and then press **Insert** to add the disk to the **Selected Drives** list, see Figure 3-29.

Figure 3-29 Selected Drives Screen

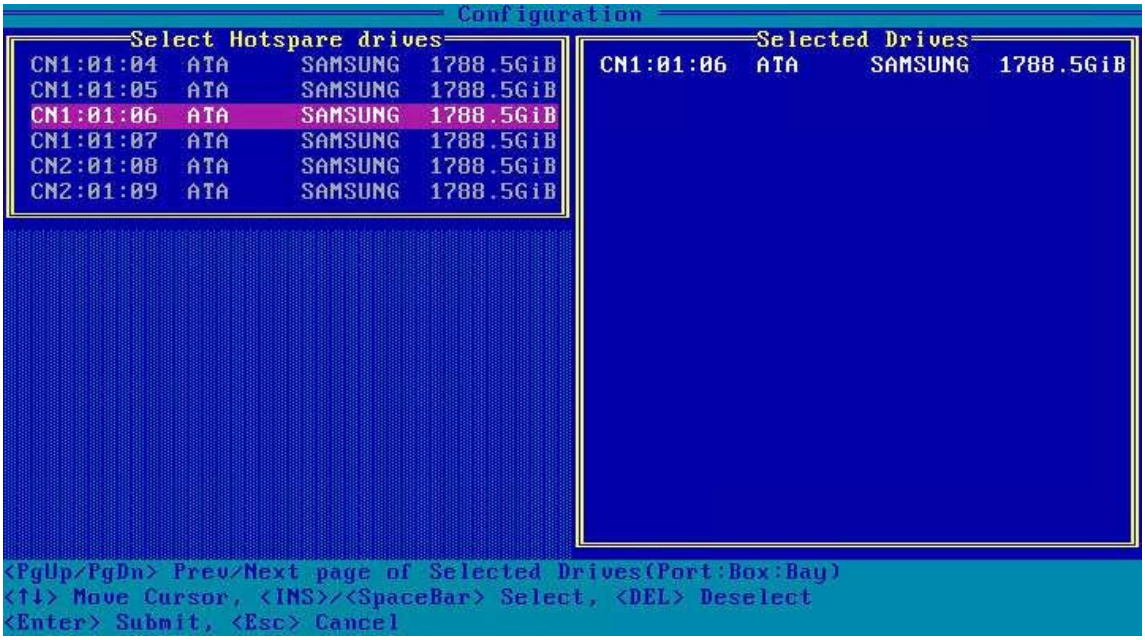

Note

Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

4. Press **Enter**. A confirmation dialog box is displayed, see Figure 3-30.

Figure 3-30 Confirmation Dialog Box

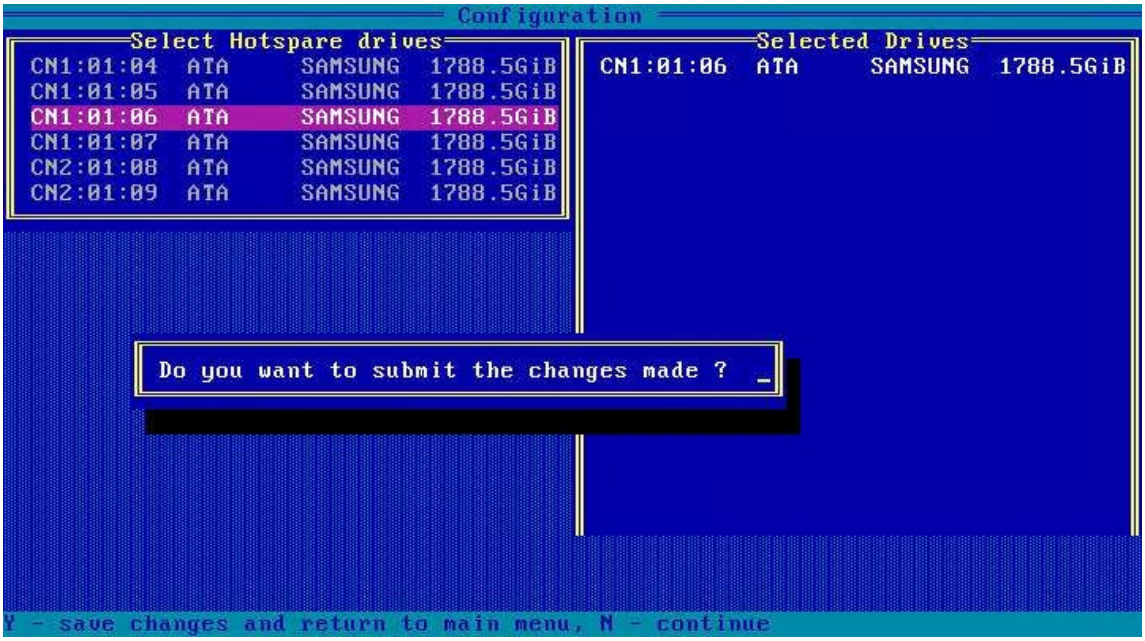

5. At the cursor in the confirmation dialog box, enter **Y**. The **Select Spare Type** screen is displayed, see Figure 3-31.

Figure 3-31 Select Spare Type Screen

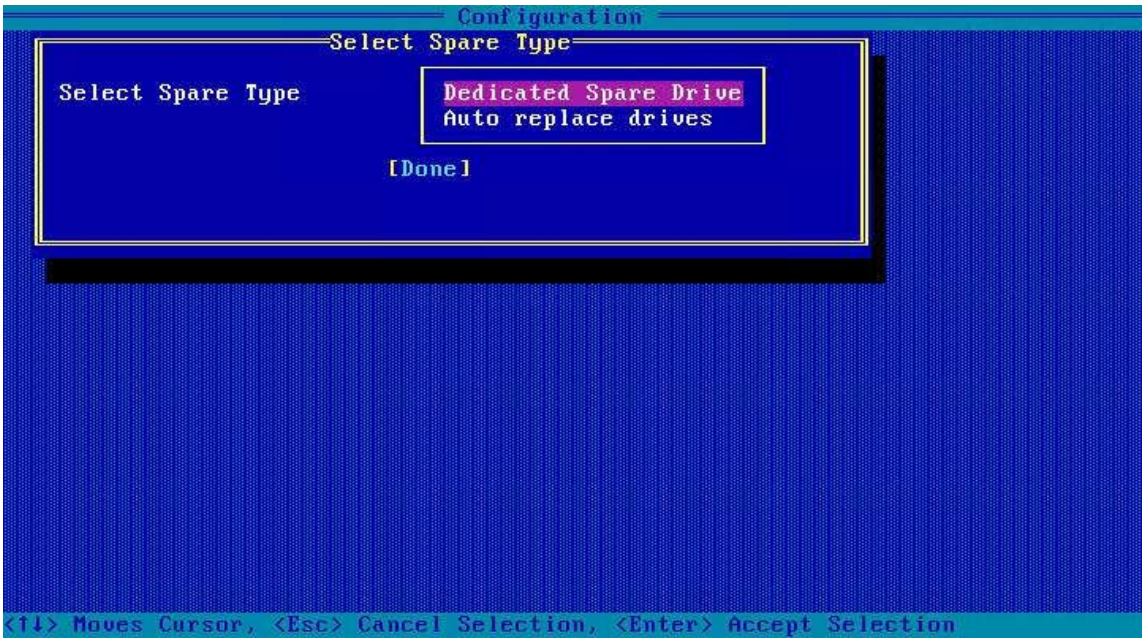

6. Use the arrow keys to select the type of hot spare disk to be created, and then press **Enter** for confirmation, see Figure 3-32.

Figure 3-32 Select Spare Type Screen

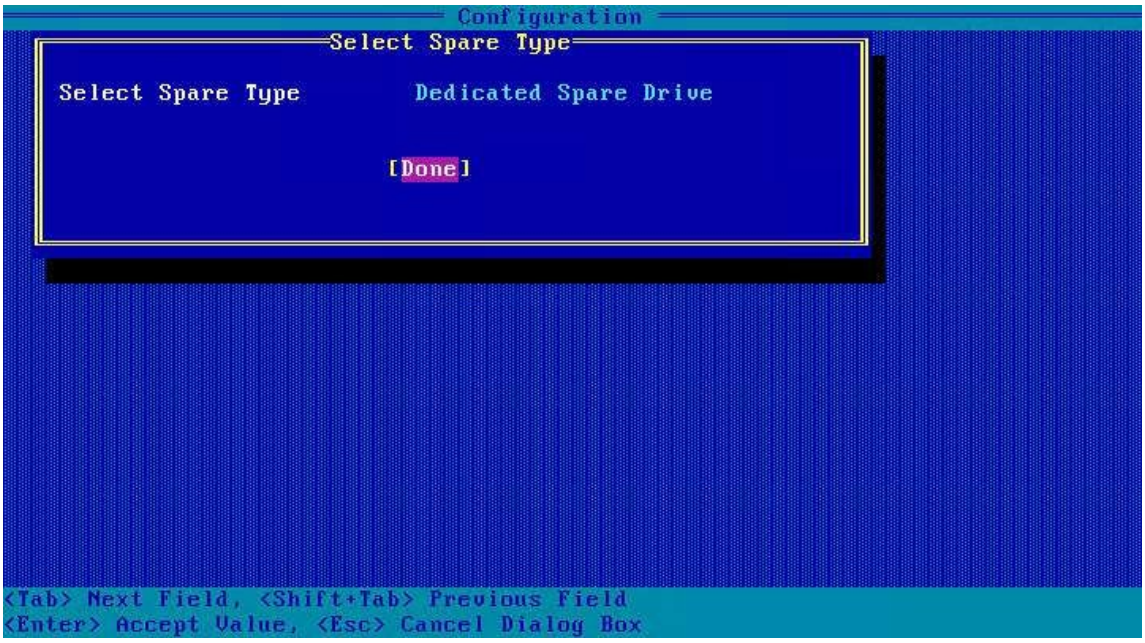

7. Use **Tab** to select **Done**, and then press **Enter** to create the hot spare disk, see Figure 3-33.

Figure 3-33 Creating a Hot Spare Disk

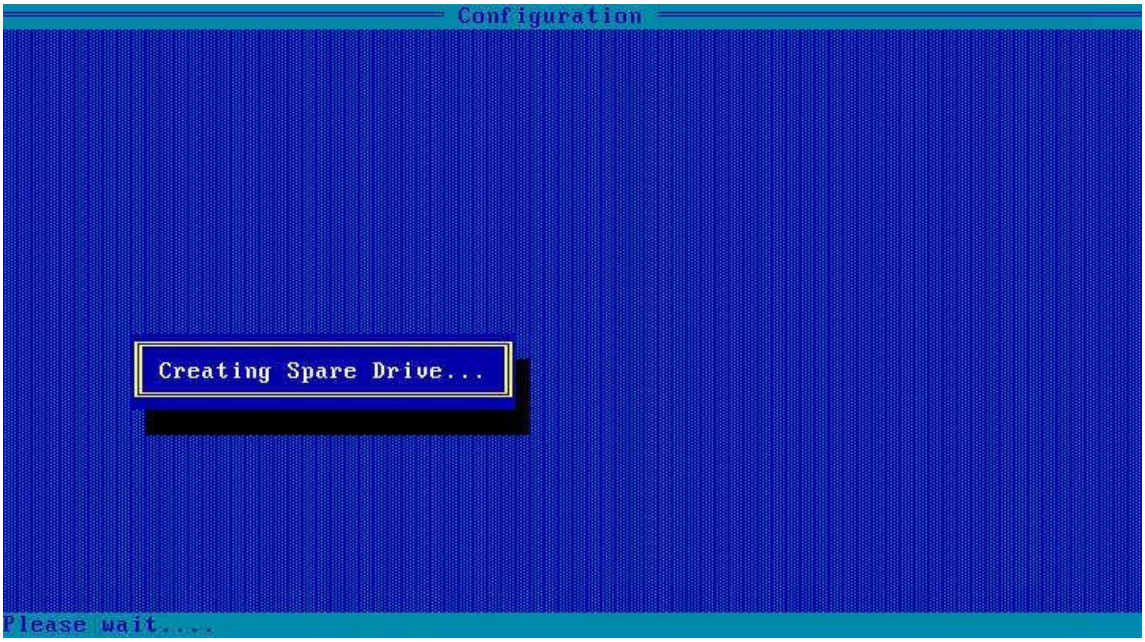

3.4.3 Deleting a RAID Volume

Abstract

When a server no longer needs a RAID volume, you can delete the RAID volume to release the disk space.

Notice

The data that is lost during deletion of the RAID volume cannot be restored. Therefore, you must make sure that you have backed up important data before deleting the volume.

Prerequisite

A RAID volume is created successfully. For details, refer to “3.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 3-34.

Figure 3-34 List of Arrays Screen

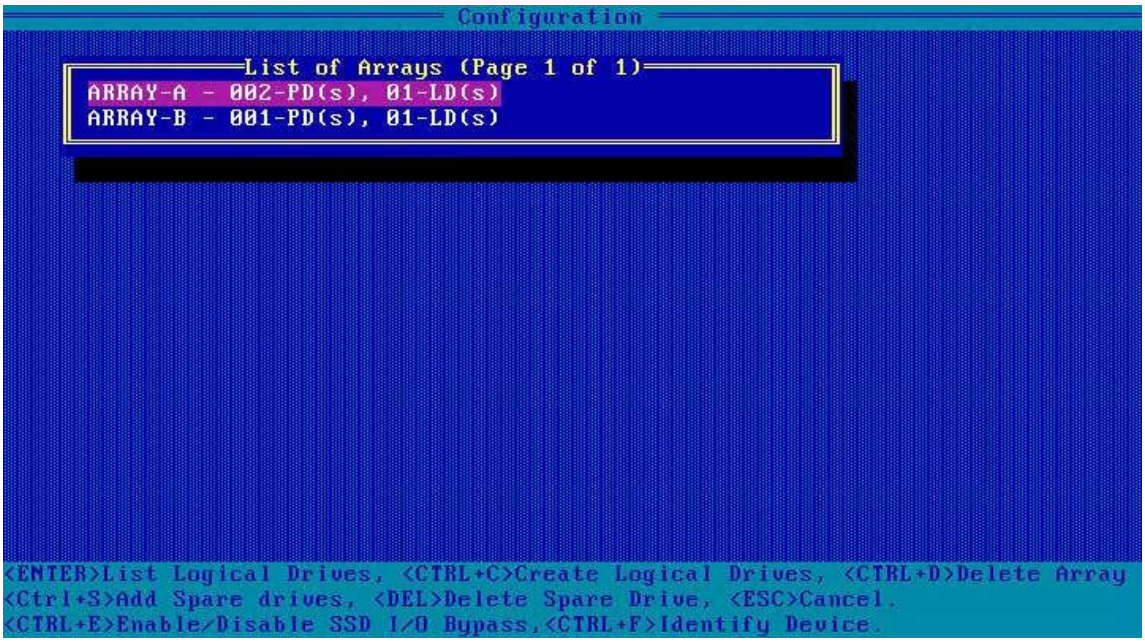

2. Use the arrow keys to select the array where the RAID volume to be deleted is located, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 3-35.

Figure 3-35 List of Logical Drives Screen

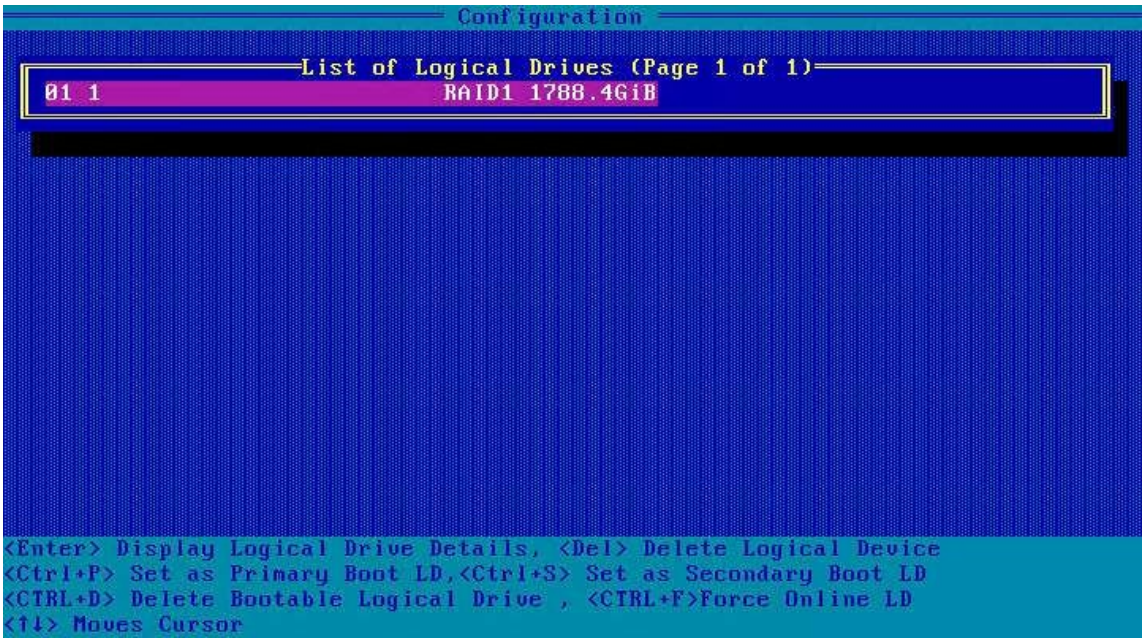

3. Use the arrow keys to select the RAID volume to be deleted, and press **Delete**. A warning dialog box is displayed, see Figure 3-36.

Figure 3-36 Warning Dialog Box

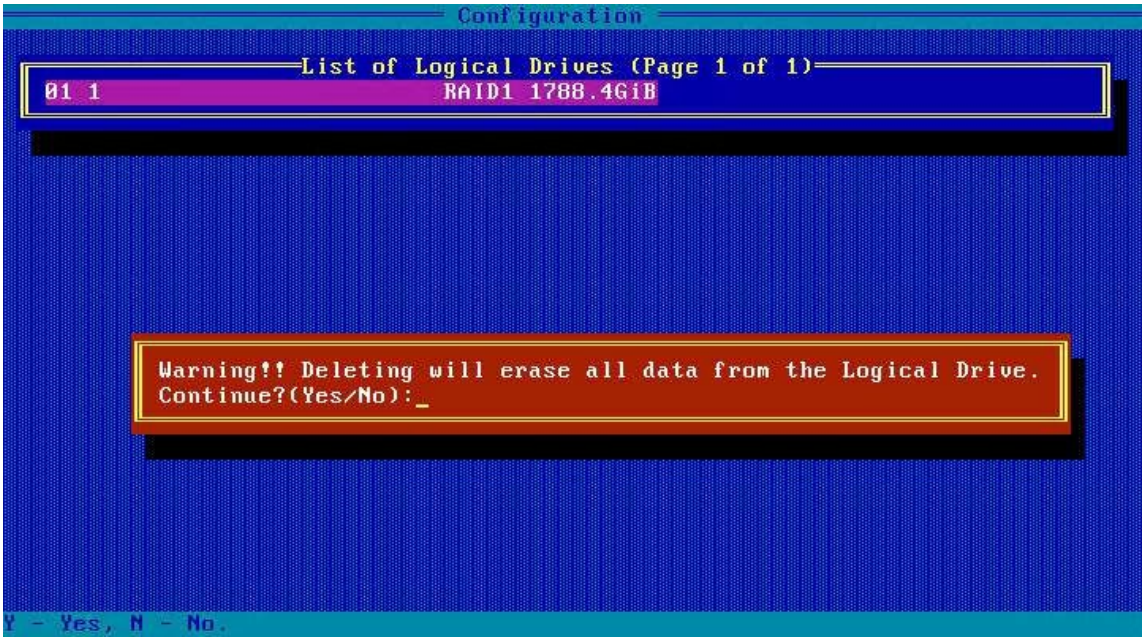

4. At the cursor in the warning dialog box, enter **Y** to delete the selected hot spare disk, see Figure 3-37.

Figure 3-37 Deleting a RAID Volume

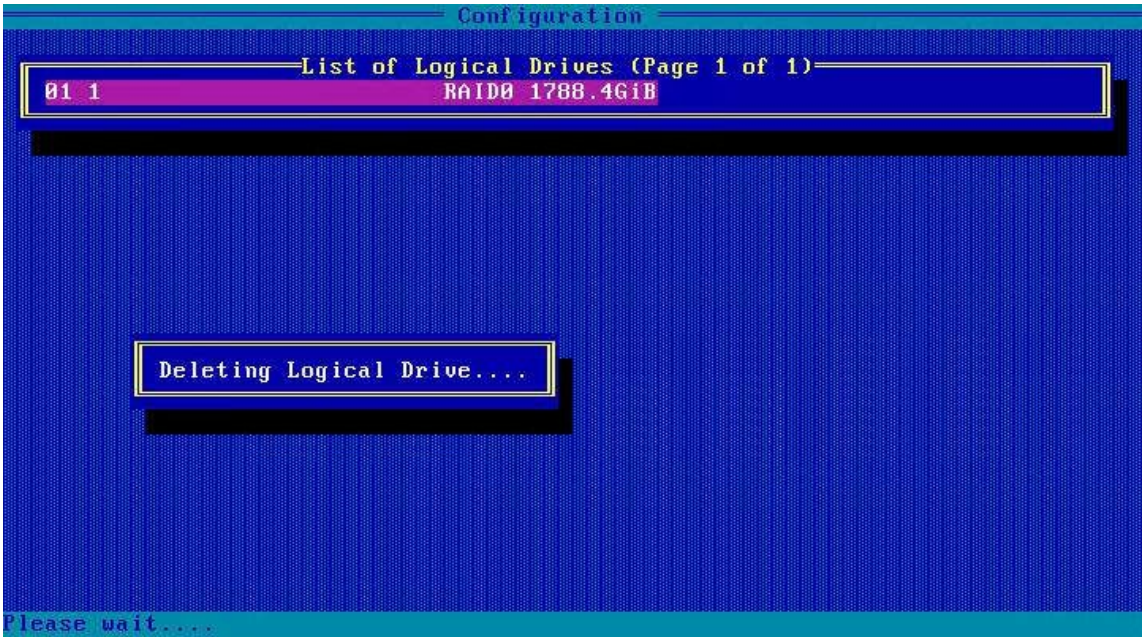

3.4.4 Deleting a Hot Spare Disk

Abstract

When the number of disks of a server cannot meet the requirements, you can delete an existing hot spare disk and restore it to a common disk.

Prerequisite

A hot spare disk is already created. For details, refer to “3.4.2 Creating a Hot Spare Disk”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 3-38.

Figure 3-38 List of Arrays Screen

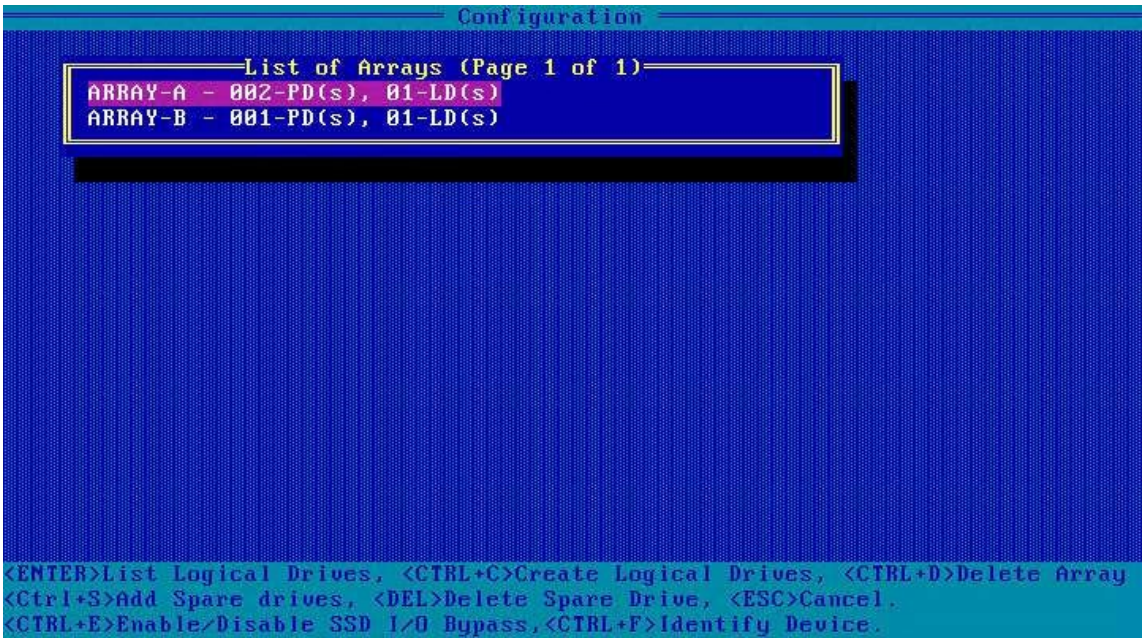

2. Use the arrow keys to select the array where the hot spare disk to be deleted is located, and press **Delete**. The **Select Hotspare drives** screen is displayed, see Figure 3-39.

Figure 3-39 Select Hotspare Drives Screen

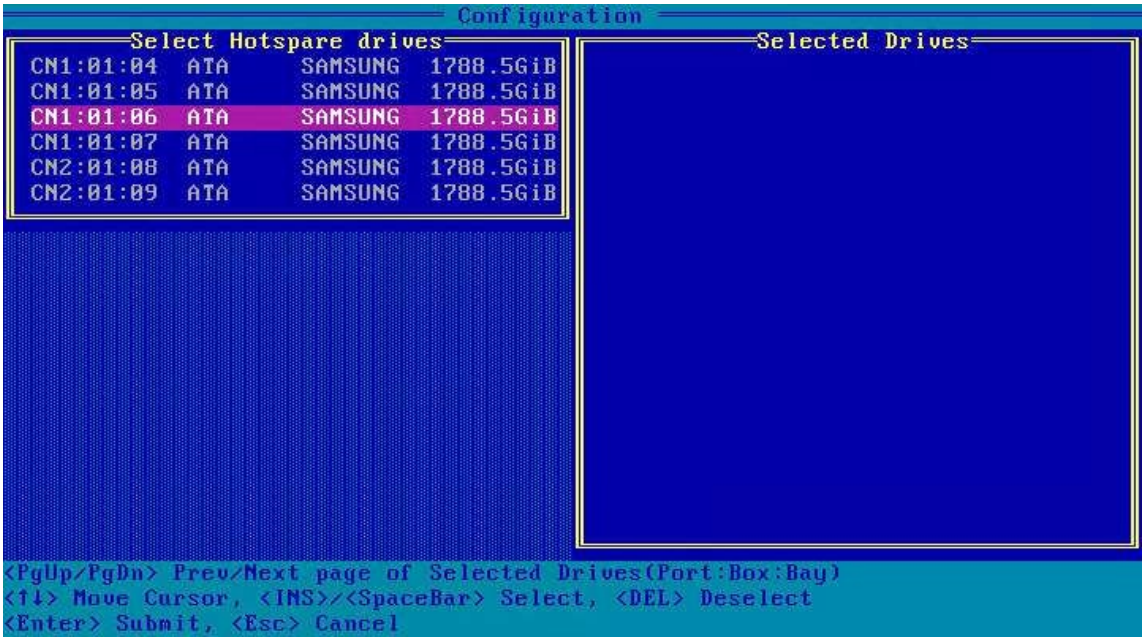

3. Use the arrow keys to select the hot spare disk to be deleted, and then press **Insert** to add the disk to the **Selected Drives** list, see Figure 3-40.

Figure 3-40 Selected Drives List

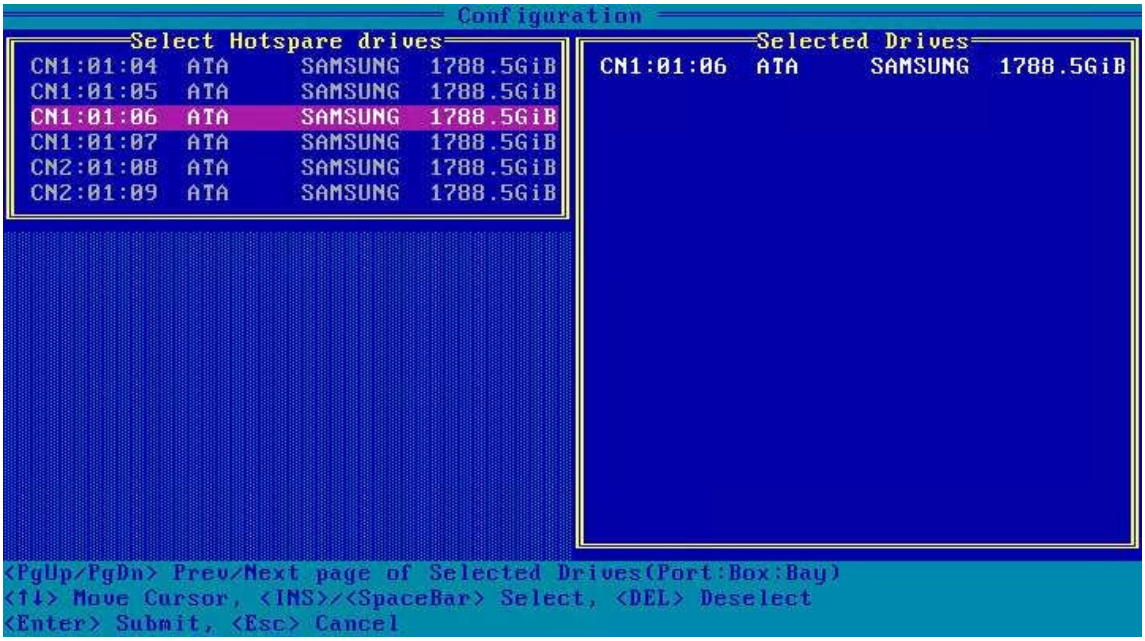

Note

Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

4. Press **Enter**. A confirmation dialog box is displayed, see Figure 3-41.

Figure 3-41 Confirmation Dialog Box

5. At the cursor in the confirmation dialog box, enter **Y** to delete the selected hot spare disk, see Figure 3-42.

Figure 3-42 Deleting a Hot Spare Disk

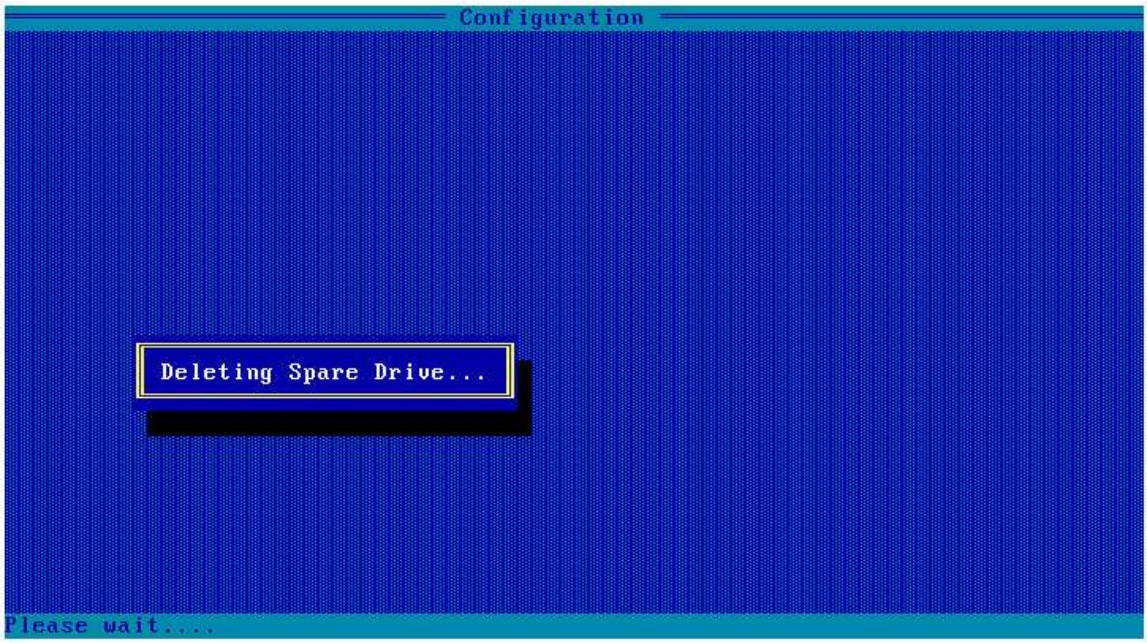

3.4.5 Locating a Disk

Abstract

After the indicator of a disk is lit, you can locate the disk so that you can easily replace or maintain it.

Steps

1. In the Options area on the BIOS configuration utility screen, use the arrow keys to select **Disk Utilities**, and then press **Enter**. The **Select Disks and press <Enter>** screen is displayed, see Figure 3-43.

Figure 3-43 Select Disks and Press <Enter> Screen

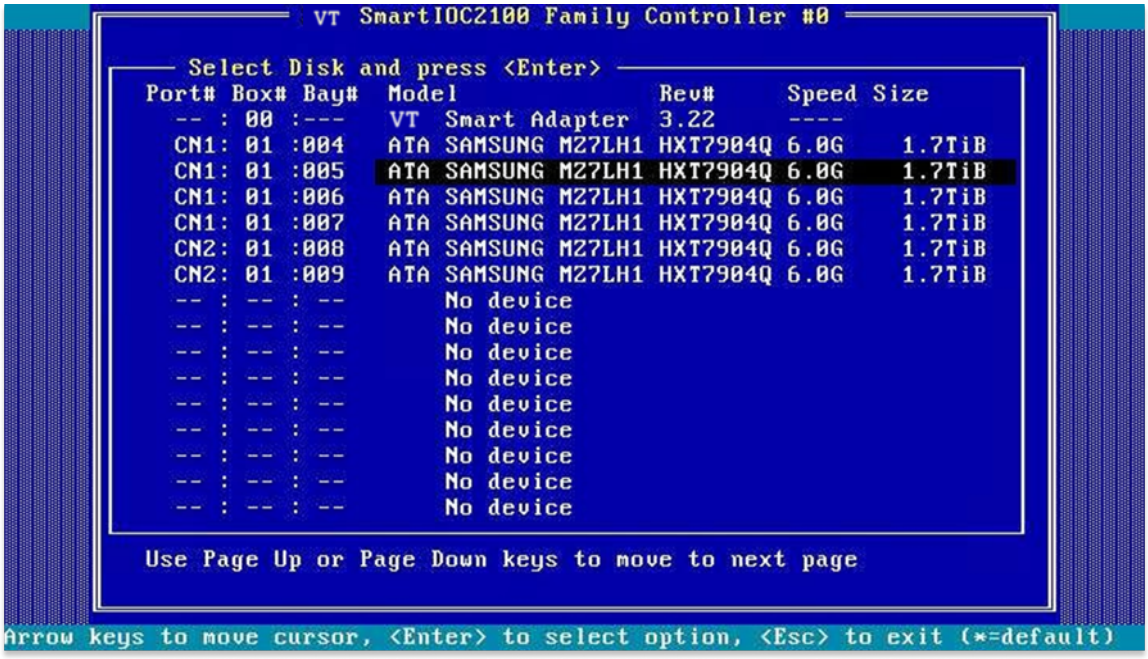

2. Use the arrow keys to select the disk to be located, and then press **Enter**. A function menu is displayed, see Figure 3-44.

Figure 3-44 Function Menu

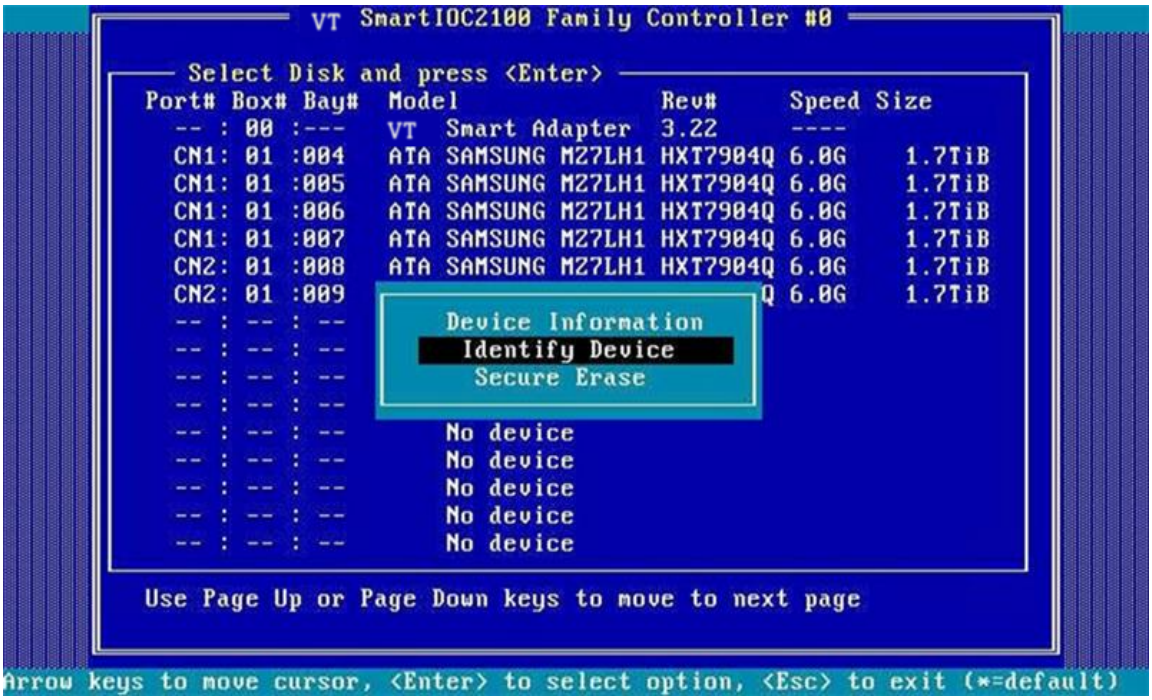

3. Use the arrow keys to select **Identify Device**, and then press **Enter**. A prompt message is displayed, see Figure 3-45. At this time, the red indicator of the corresponding disk is lit and flashes continuously.

Figure 3-45 Disk Indicator Lit

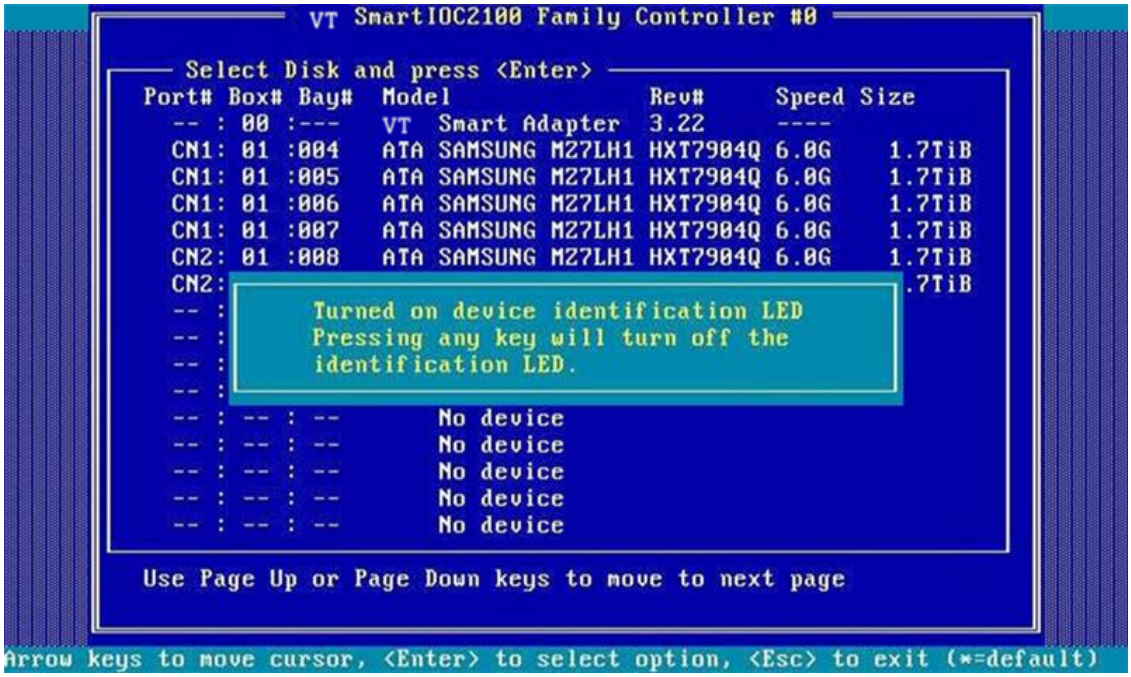

4. Press any key to go out the disk indicator. The disk locating ends.

3.4.6 Configuring a Pass-Through Disk

Abstract

In accordance with the actual RAID configuration, you can configure a pass-through disk in either of the following ways:

- Setting the mode of all ports in batches:
  If the disks connected to all ports are not used to create a RAID volume, the mode of these ports can be set to HBA in a unified manner.
- Setting the mode of ports separately:
  If the disks connected to some ports have been used to create a RAID volume, the mode of these ports can be set to Mixed and that of other ports can be set to HBA.

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.
- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

1. In the Options area on the BIOS configuration utility screen, use the arrow keys to select **Controller Details**, and then press **Enter**. The controller details are displayed. The value of **Controller Mode** is **MIXED**, see Figure 3-46.

Figure 3-46 Controller Details

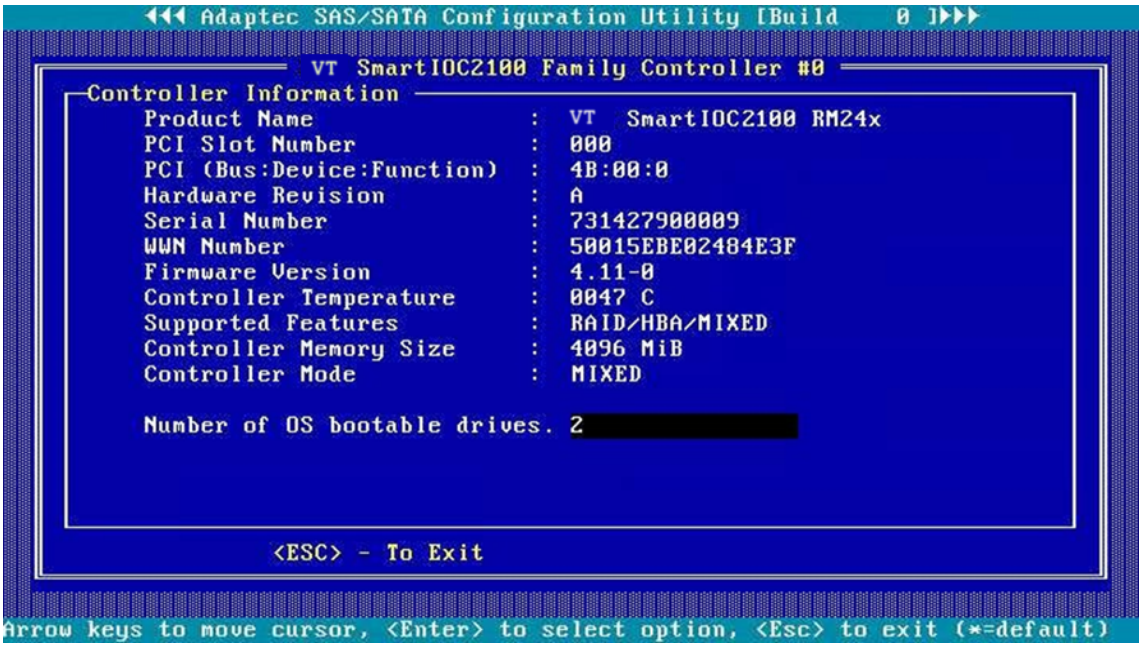

2. Press **ESC** to return to the BIOS configuration screen. Use the arrow keys to select **Configure Controller Settings** and press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-48.

Figure 3-47 Configure Controller Settings Screen

3. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**. The **Configure Controller Port Mode** screen is displayed, see Figure 3-48.

Figure 3-48 Configure Controller Port Mode Screen

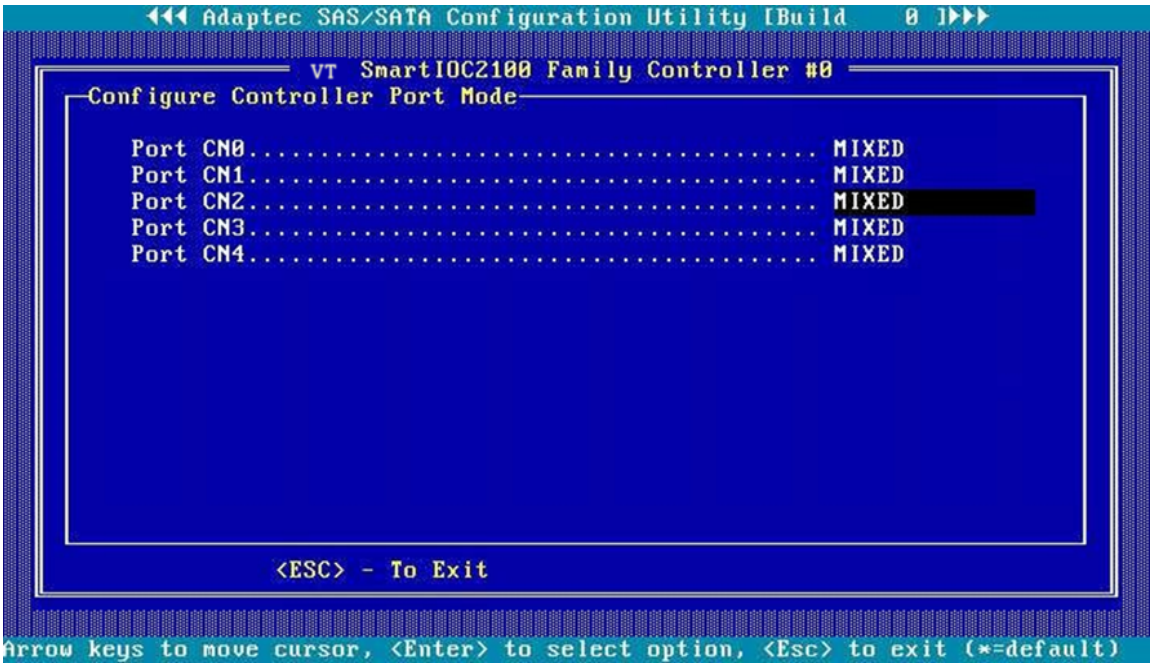

4. Use the arrow keys to select the port whose connected disk is not used to create a RAID volume, and press **Enter**. In the displayed dialog box, set the port mode to **HBA**, see Figure 3-50.

Figure 3-49 Setting the Port Mode to HBA

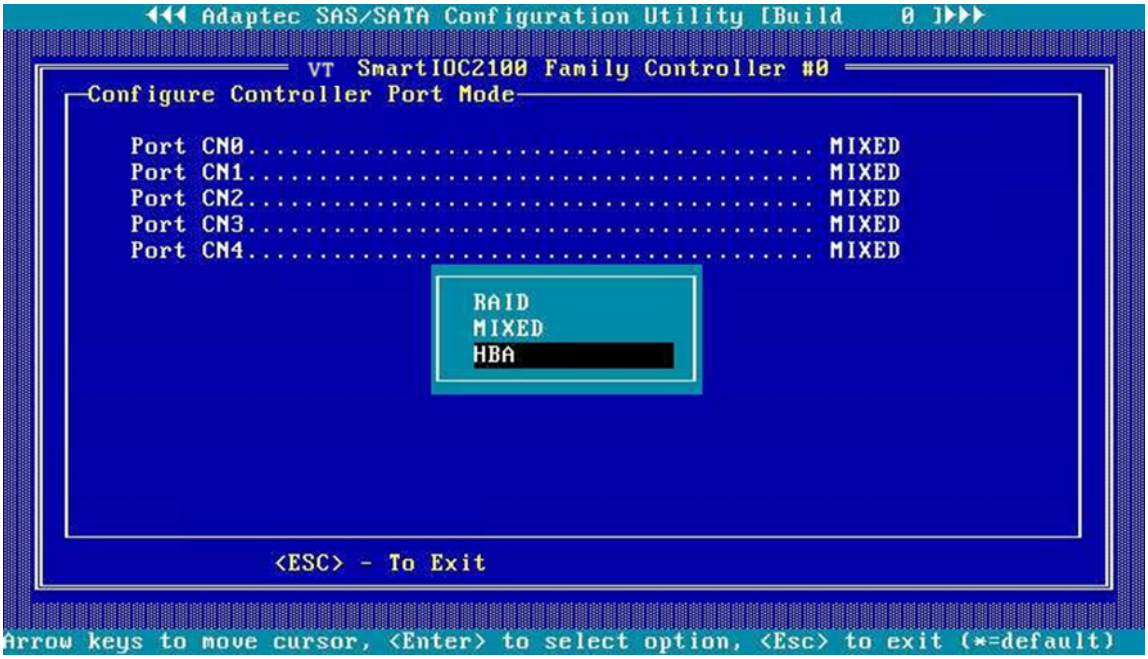

5. (Optional) If the connected disk is already used to create a RAID volume, set the port mode to **MIXED**, see Figure 3-50.

Figure 3-50 Setting the Port Mode to MIXED

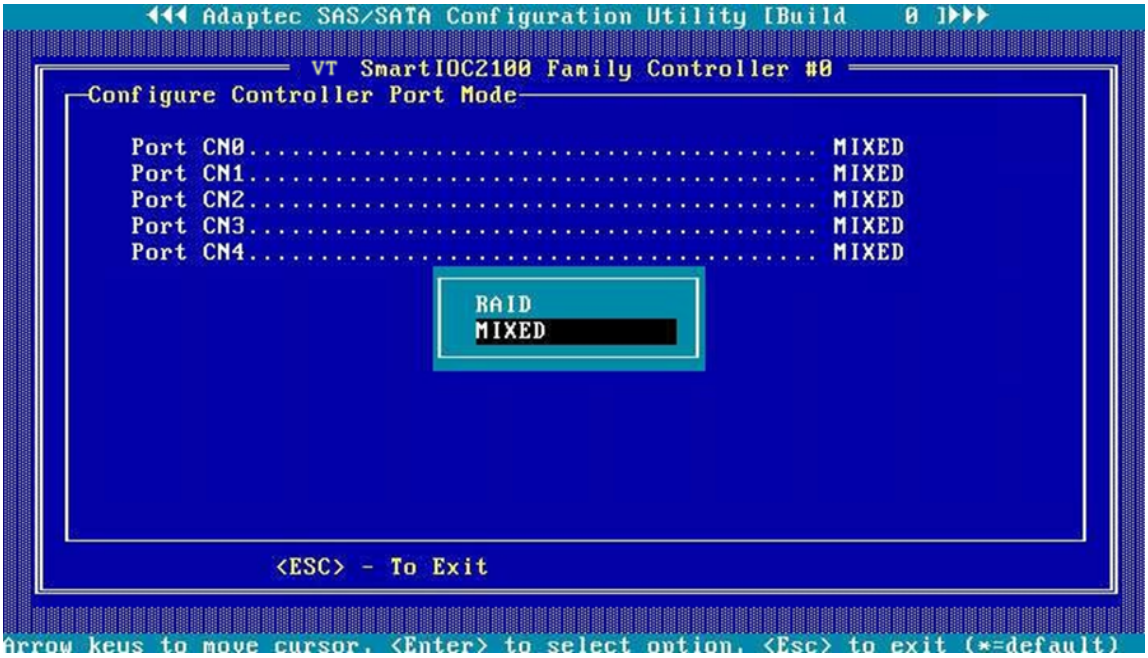

Note

The mode of the ports whose connected disks are already used to create a RAID volume cannot be set to HBA.

6. Set the port mode as required, and then press **Esc** to exit. In the displayed dialog box, select **Yes** to save the configuration, see Figure 3-51.

Figure 3-51 Saving the Configuration

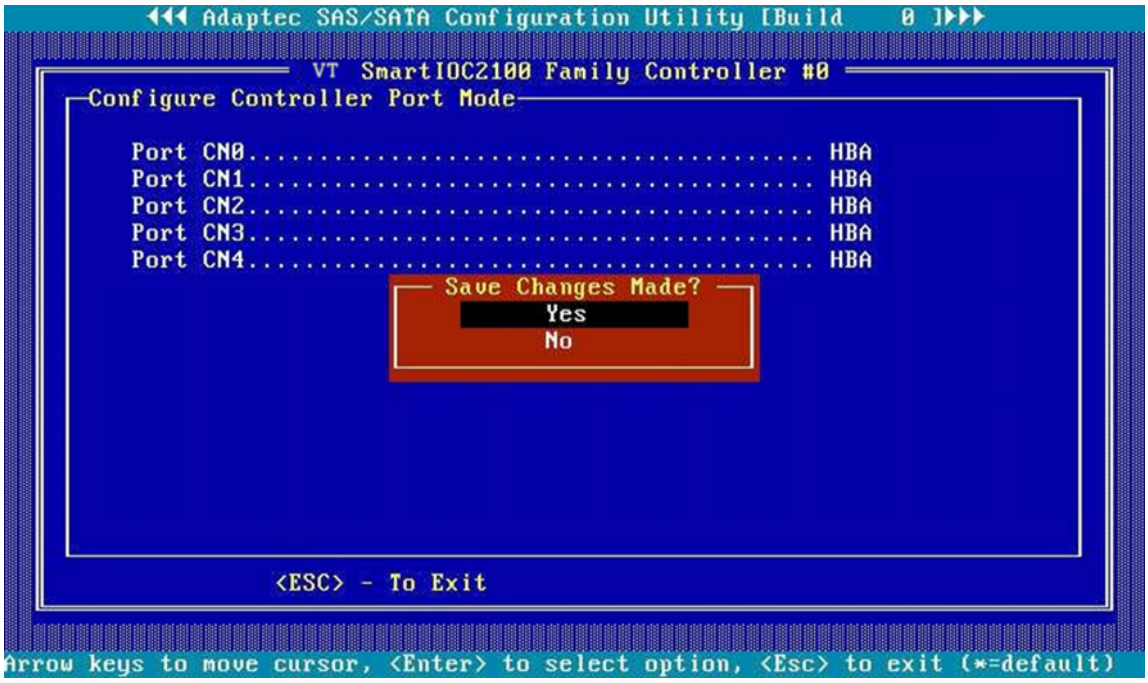

7. Wait until a configuration success message is displayed, see Figure 3-52.

Figure 3-52 Successful Configuration

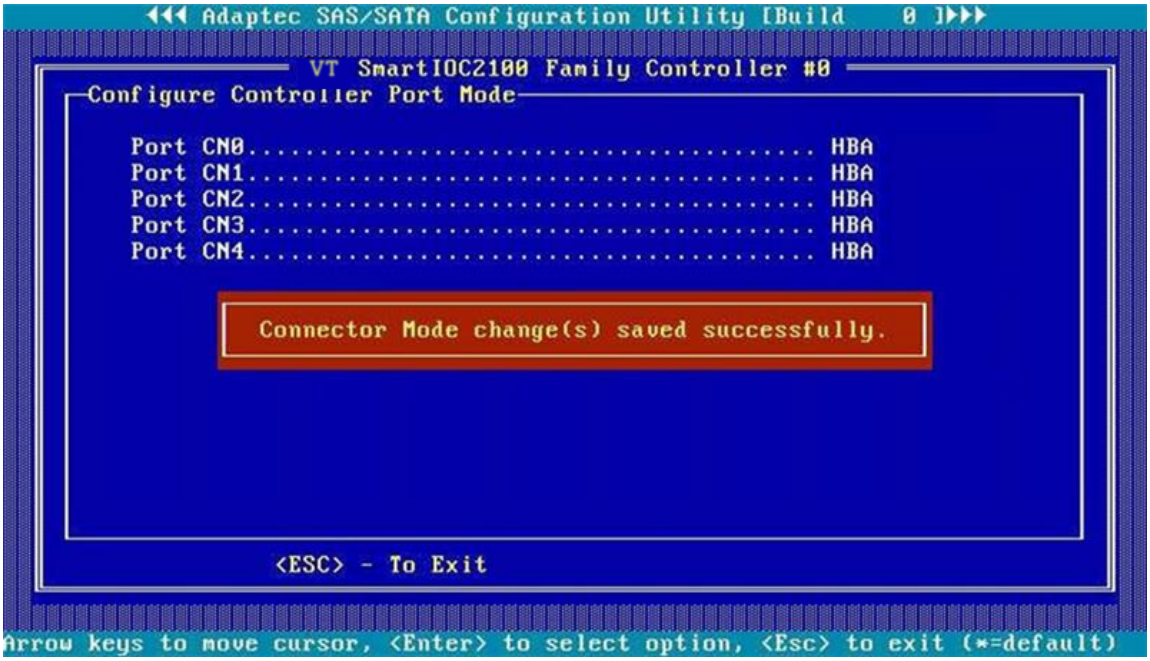

8. Press **Esc** multiple times until the **Exit Utility** dialog box is displayed, see Figure 3-53.

Figure 3-53 Exit Utility Dialog Box

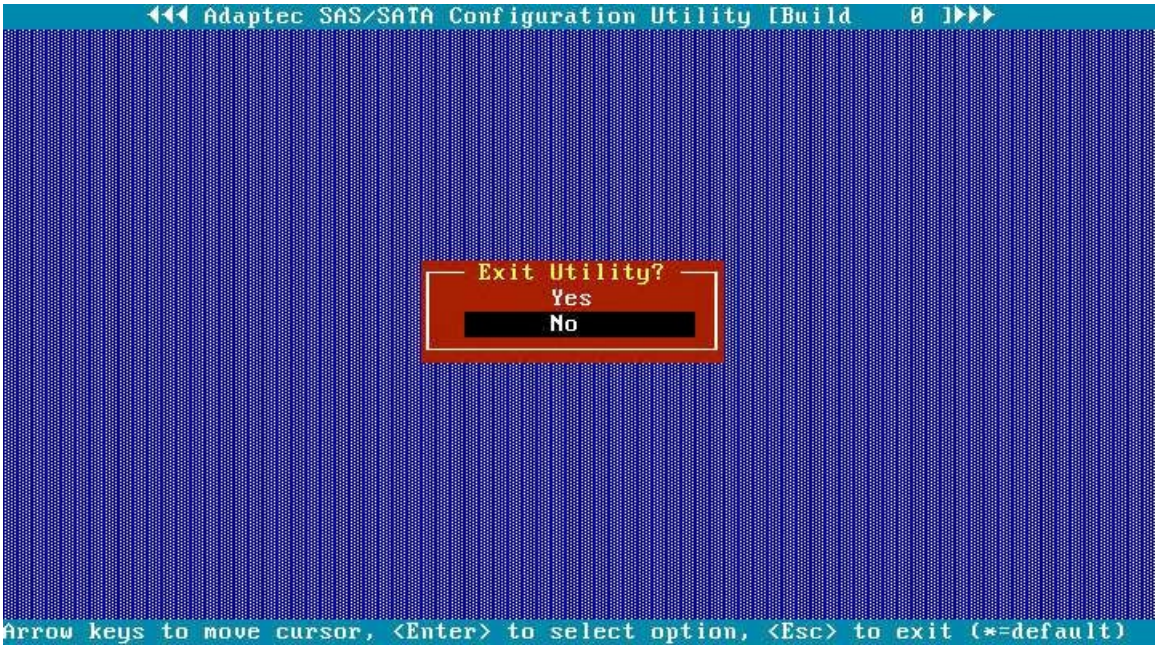

9. In the **Exit Utility** dialog box, select **NO**, and then press **Enter**. The BIOS configuration utility screen is displayed, see Figure 3-54.

Figure 3-54 BIOS Configuration Utility Screen

10. Use the arrow keys to select **Controller Details**, and then press **Enter**. The controller details are displayed, see Figure 3-55. The value of **Controller Mode** is **HBA**.

Figure 3-55 Controller Details

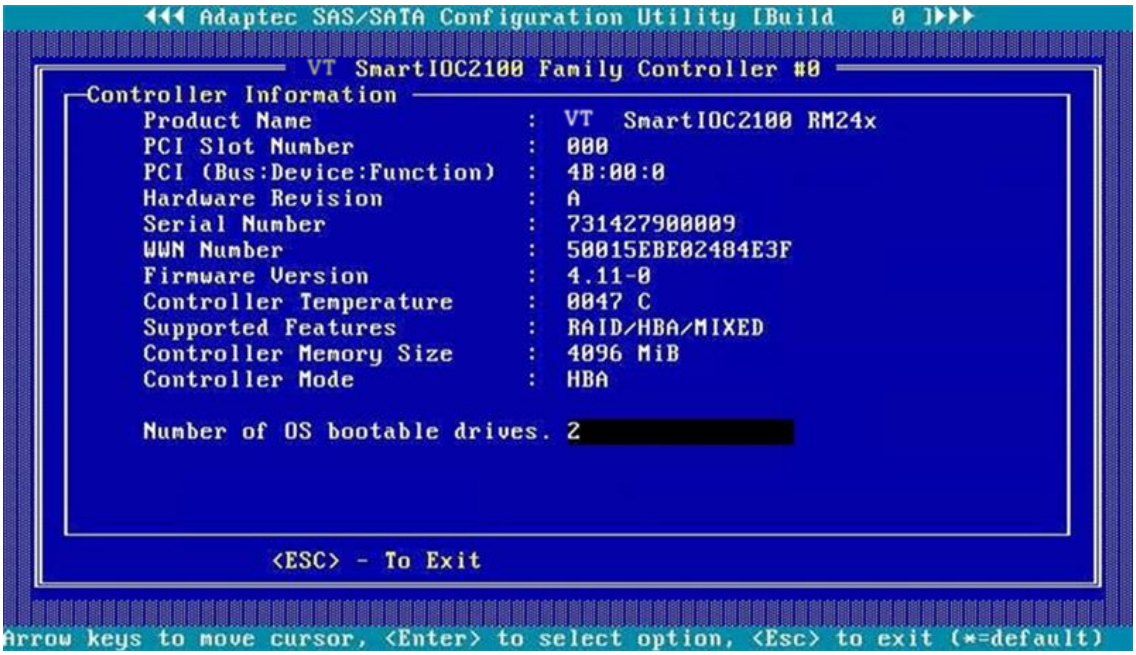

3.5 Common Configurations (UEFI Mode)

By using the BIOS configuration utility, you can configure and maintain a created RAID volume. For a description of the common operations on a SmartIOC 2100 RAID controller card in UEFI mode, refer to Table 3-9.

Table 3-9 Common Operations on a SmartIOC 2100 RAID Controller Card

Common Operation	Description
Setting a port mode	Refer to "3.5.1 Setting the Mode of a Port".
Locating a disk	Refer to "3.5.2 Locating a Disk".
Creating a hot spare disk	Refer to "3.5.3 Creating a Hot Spare Disk".
Changing a hot spare disk	Refer to "3.5.4 Changing a Hot Spare Disk".
Deleting a hot spare disk	Refer to "3.5.5 Deleting a Hot Spare Disk".
Configuring the power mode	Refer to "3.5.6 Configuring the Performance or Power Mode".
Deleting a RAID volume	Refer to "3.5.7 Deleting a RAID volume".
Deleting a disk group	Refer to "3.5.8 Deleting a Disk Group".
Clearing RAID configuration information	Refer to "3.5.9 Clearing RAID Configuration Information".
Configuring a pass-through disk	Refer to "3.5.10 Configuring a Pass-Through Disk".

3.5.1 Setting the Mode of a Port

Abstract

The ports of a SmartIOC 2100 RAID controller card, namely, the ports connected with a disk backplane and disk cables, support three modes: RAID, HBA and Mixed. Before adding the disk corresponding to a port to a RAID logical volume, you must check the port mode.
The SmartIOC 2100 RAID controller card supports setting the port mode in the following two ways:

- Setting the mode of all ports in batches
- Setting the mode of a port separately

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.
- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

- Setting the Mode of All Ports in Batches
  1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-56.

Figure 3-56 Configure Controller Settings Screen

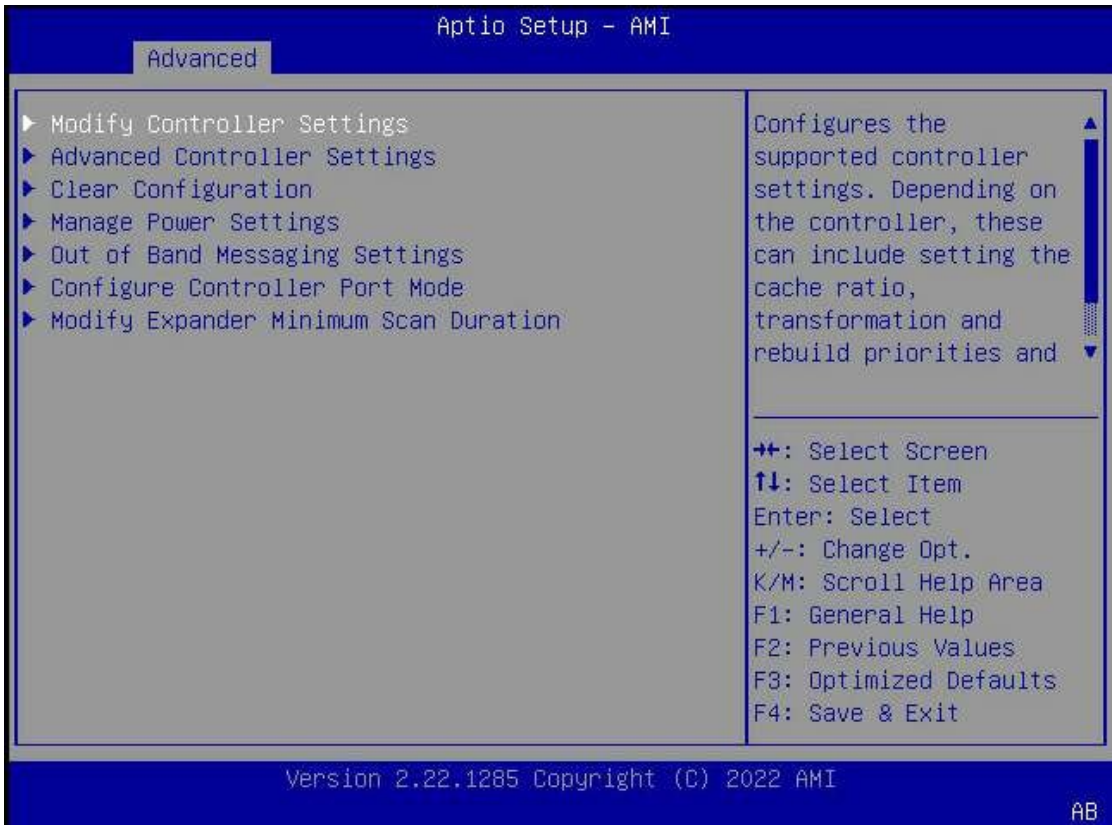

2. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**.

   The **Configure Controller Port Mode** screen is displayed, see Figure 3-57.

Figure 3-57 Configure Controller Port Mode Screen

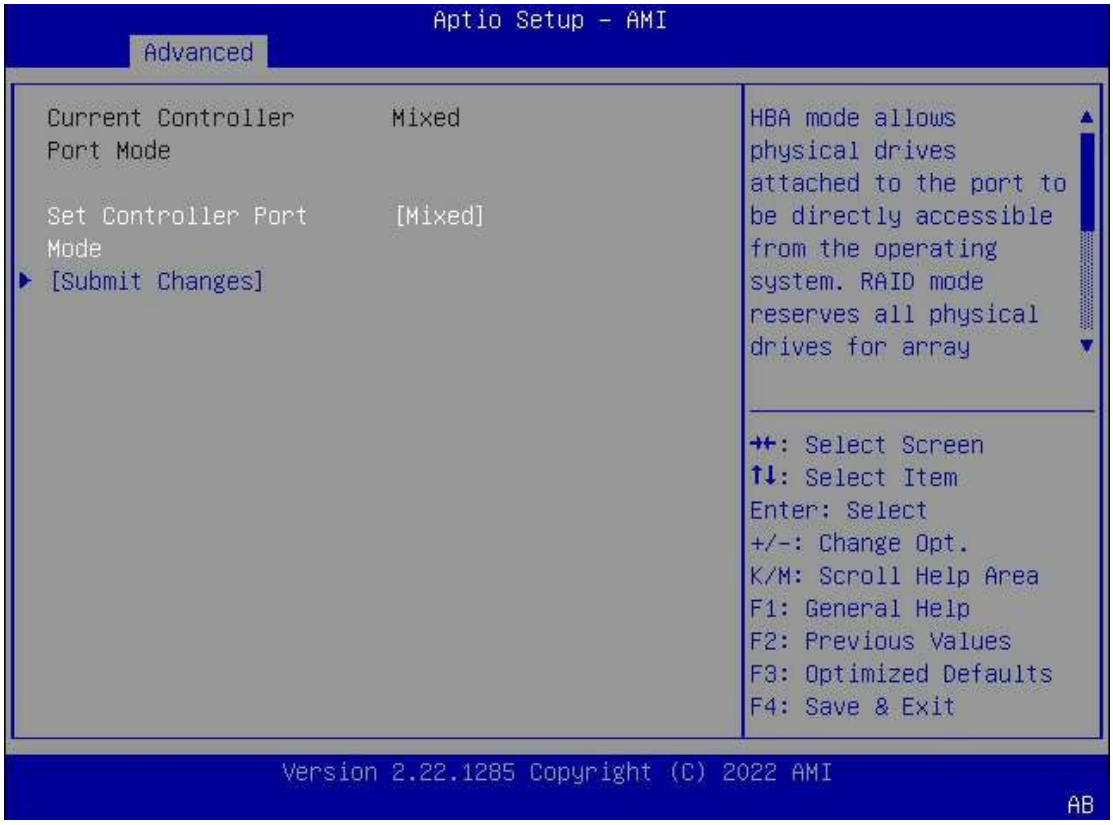

3. Use the arrow keys to select **Set Controller Port Mode**, and then press **Enter**. The **Set Controller Port Mode** dialog box is displayed, see Figure 3-58.

Figure 3-58 Set Controller Port Mode Dialog Box

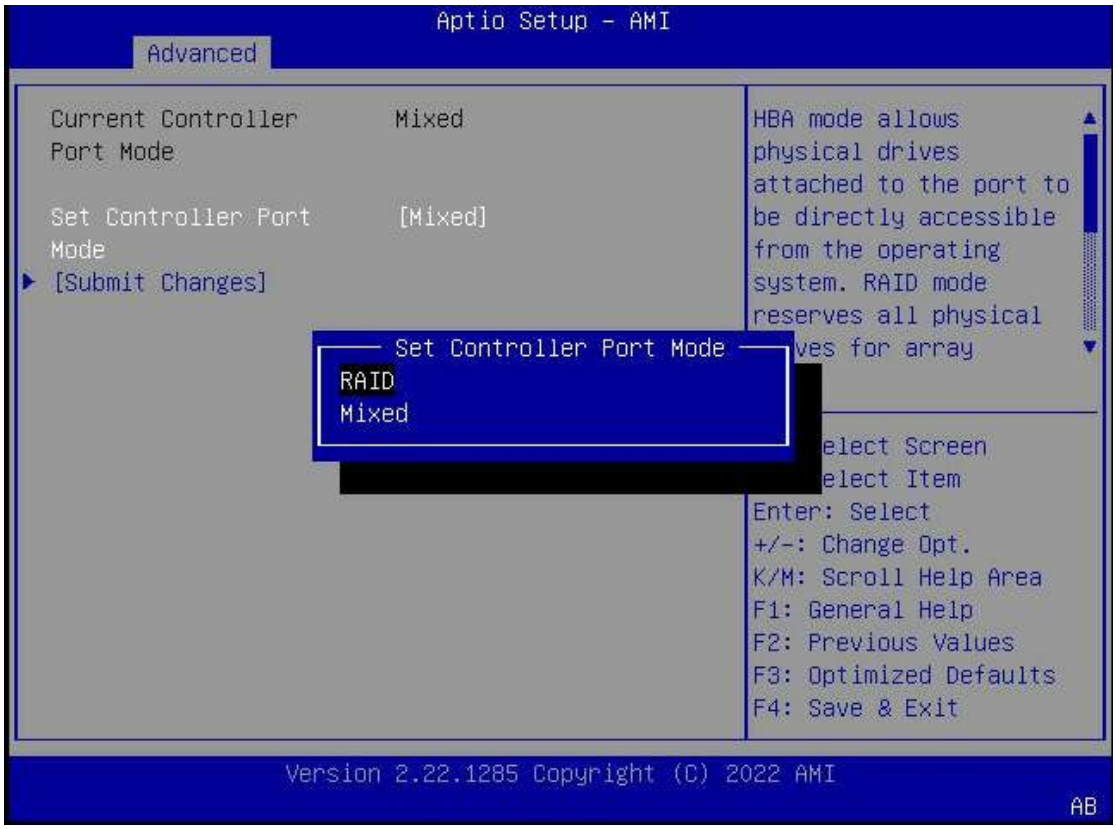

4. Use the arrow keys to select the port mode to be set, and then press **Enter**, see Figure 3-59.

Figure 3-59 Selecting a Port Mode

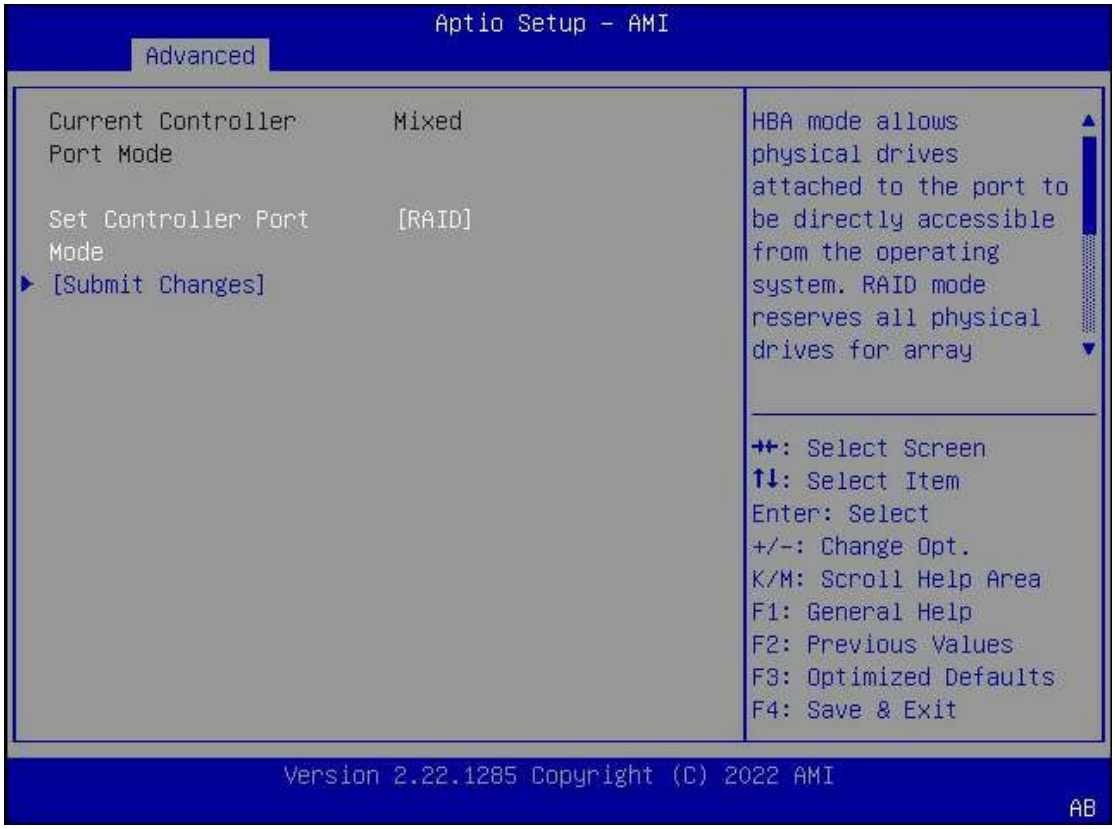

5. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 3-60.

Figure 3-60 Successful Setting

6. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

- Setting the Mode of a Port Separately

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-61.

Figure 3-61 Configure Controller Settings Screen

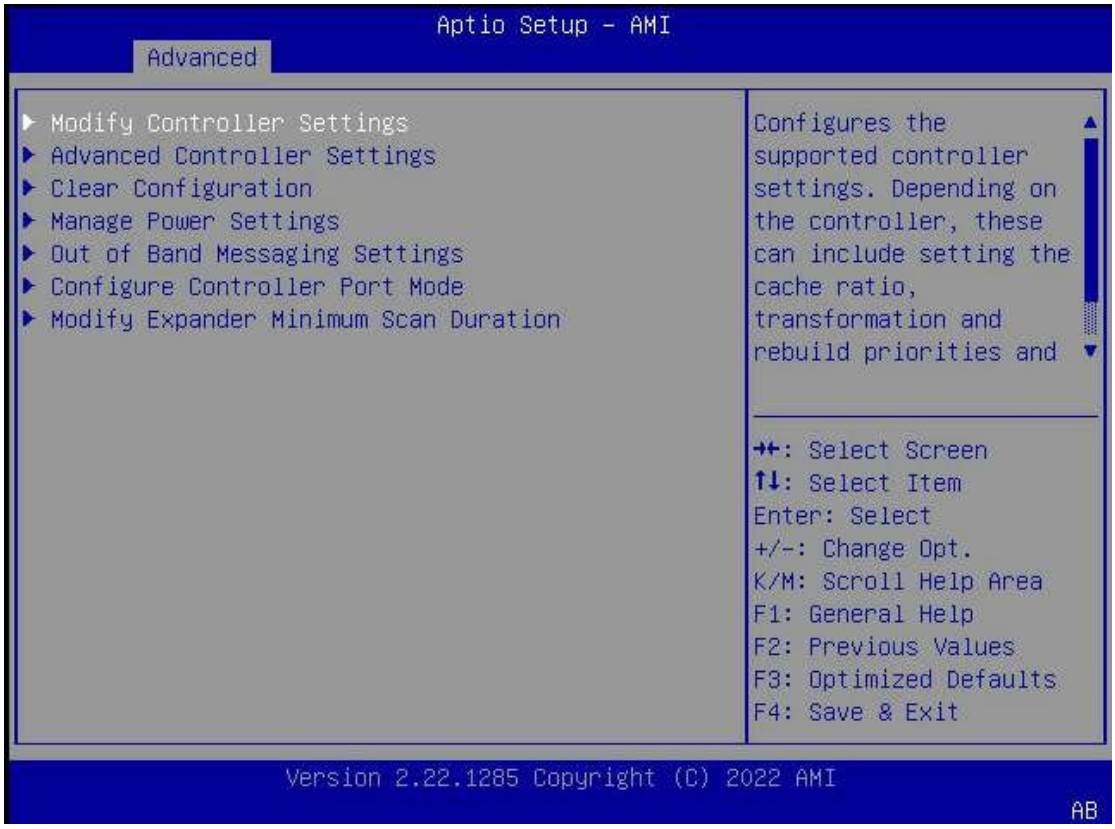

2. Use the arrow keys to select **Modify Controller Settings**, and then press **Enter**. The **Modify Controller Settings** screen is displayed, see Figure 3-62.

Figure 3-62 Modify Controller Settings Screen

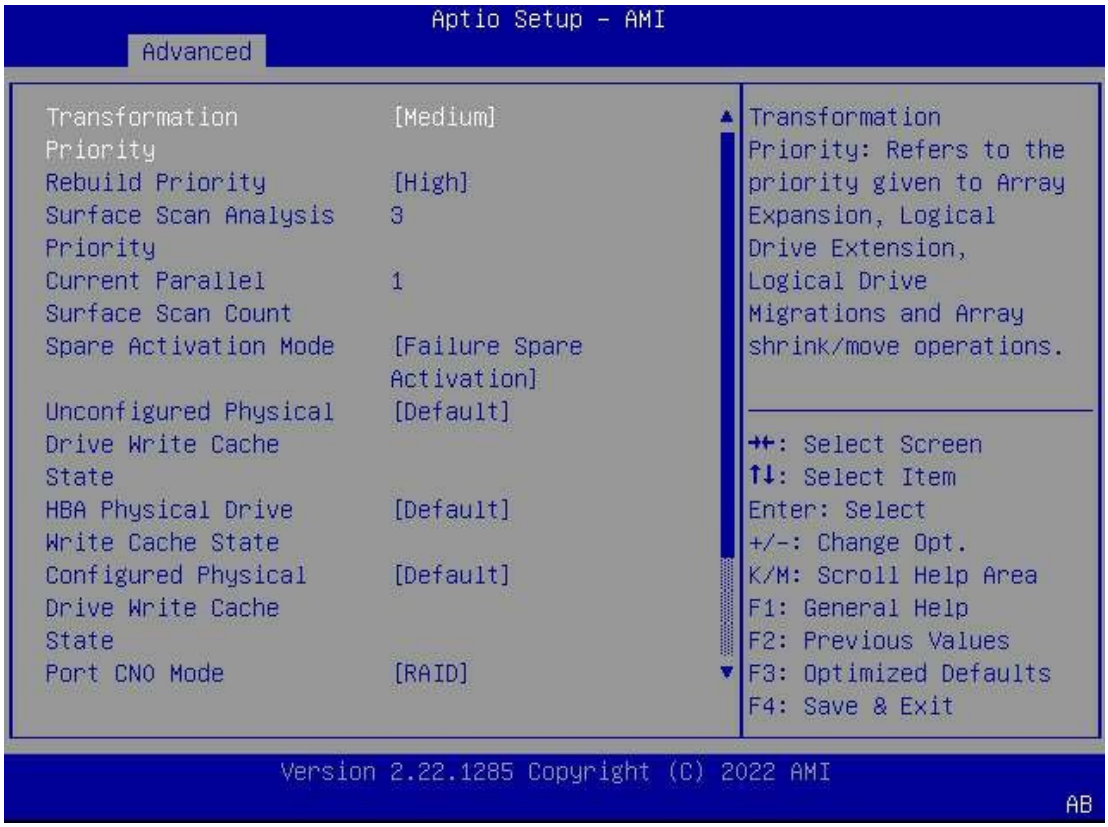

3. Use the arrow keys to select **Port CN0 Mode**, and then press **Enter**. The **Port CN0 Mode** dialog box is displayed, see Figure 3-63.

Figure 3-63 Port CNO Mode Dialog Box

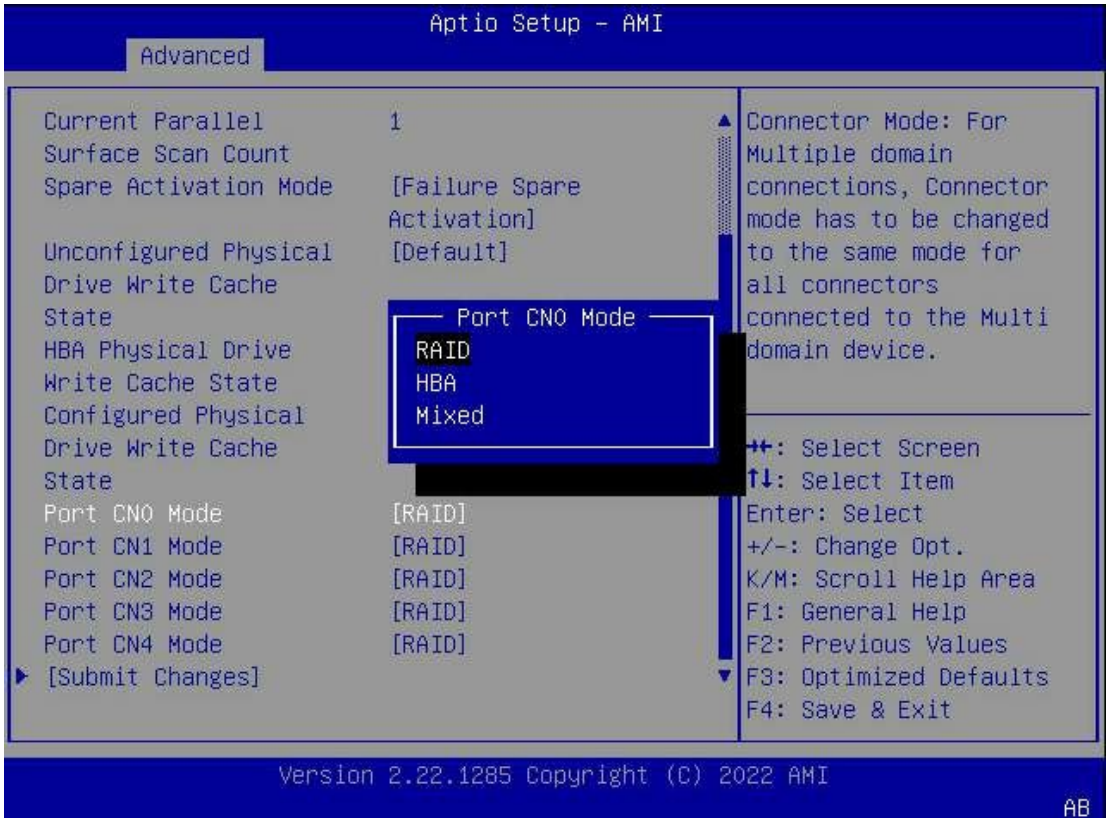

4. Use the arrow keys to select the port mode to be set, and then press **Enter**, see Figure 3-64.

Figure 3-64 Setting the Port Mode

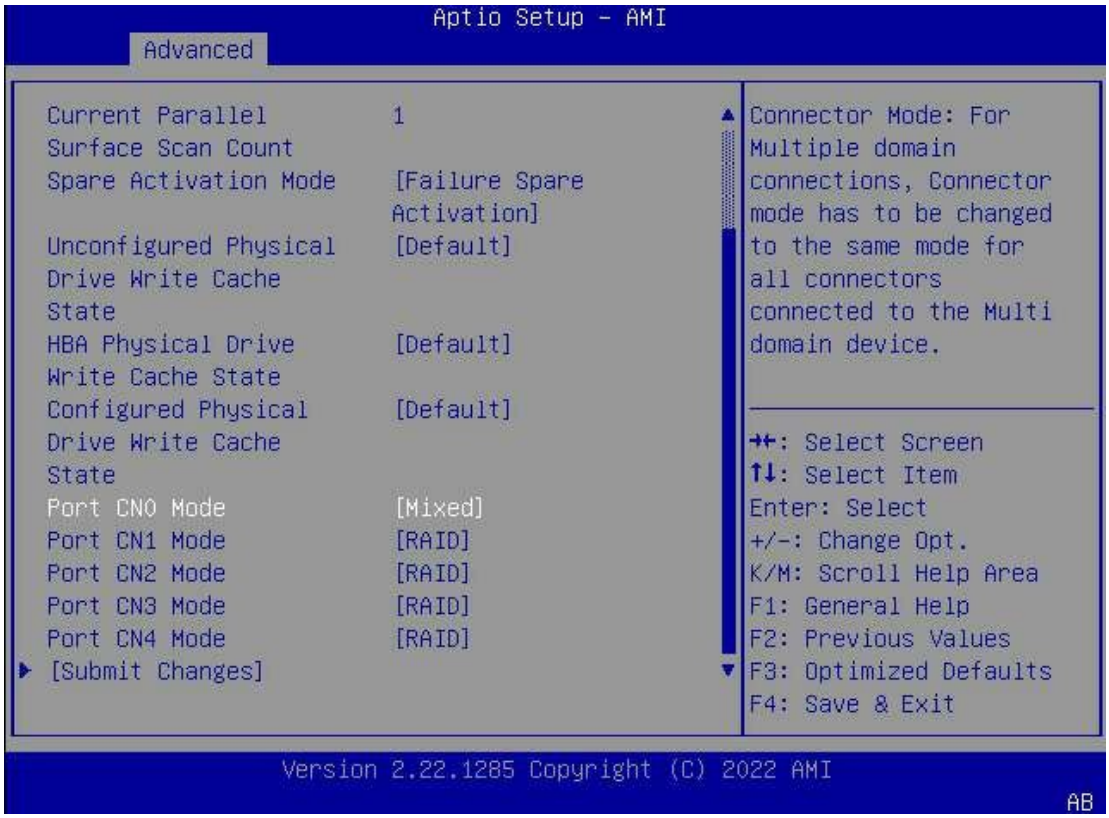

5. Repeat Step 3 and Step 4 to set the mode of another port, see Figure 3-65.

Figure 3-65 Setting the Port Mode

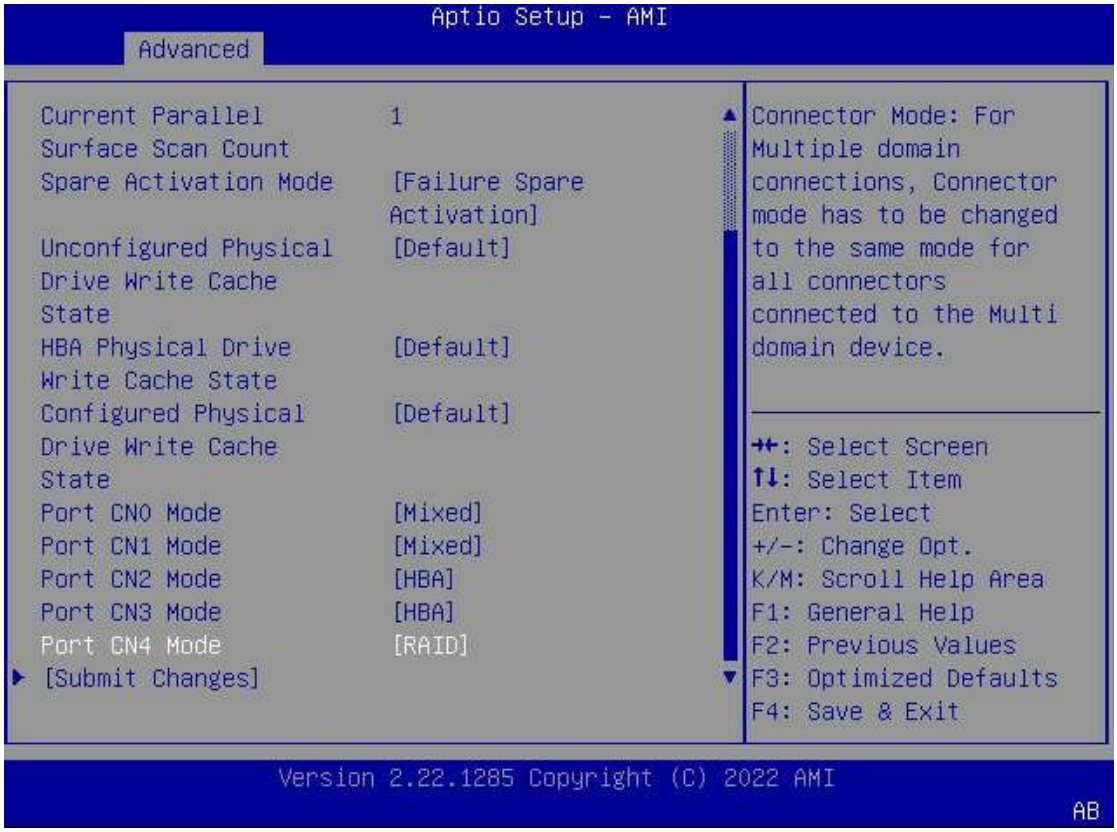

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 3-66.

Figure 3-66 Port Mode Set Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.2 Locating a Disk

Abstract

After the indicator of a disk is lit, you can locate the disk so that you can easily replace or maintain it. You can locate a physical disk or multiple disks in a disk group.

Steps

- Locating a Single Physical Disk
  1. On the controller management screen, use the arrow keys to select **Disk Utilities**, and then press **Enter**. The screen for the list of physical disks mounted on the RAID controller card is displayed, see Figure 3-67.

Figure 3-67 Screen for the List of Disks Mounted on the RAID Controller Card

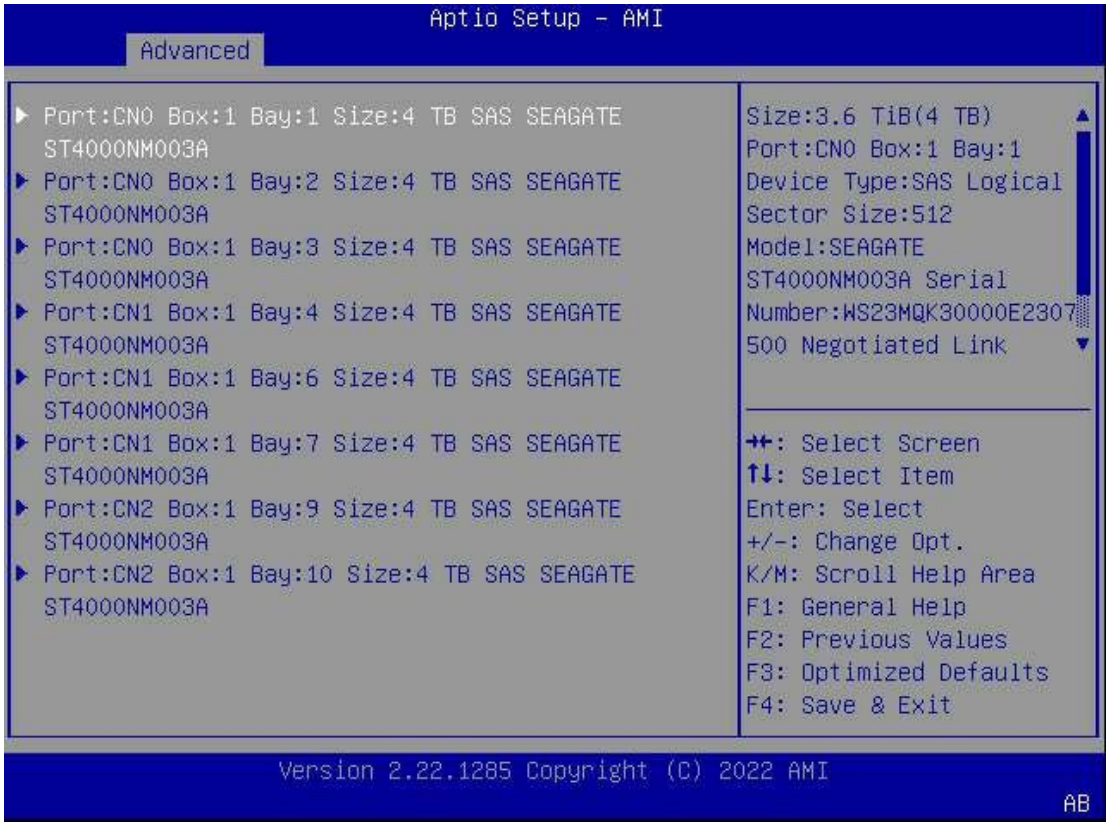

2. Use the arrow keys to select the disk to be located, and then press **Enter**. The screen for setting a disk is displayed, see Figure 3-68.

Figure 3-68 Setting a Disk

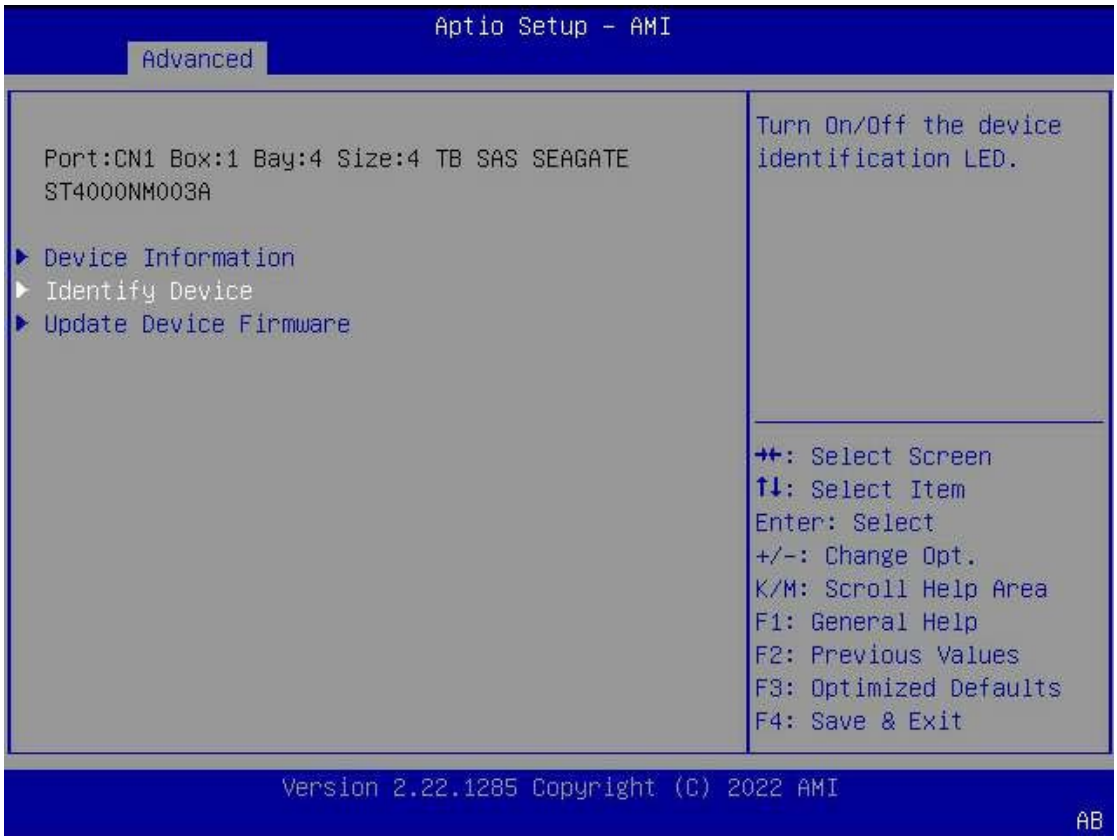

3. Use the arrow keys to select **Identify Device**, and then press **Enter**. The **Identify Device** screen is displayed, see Figure 3-69.

Figure 3-69 Identify Device Screen

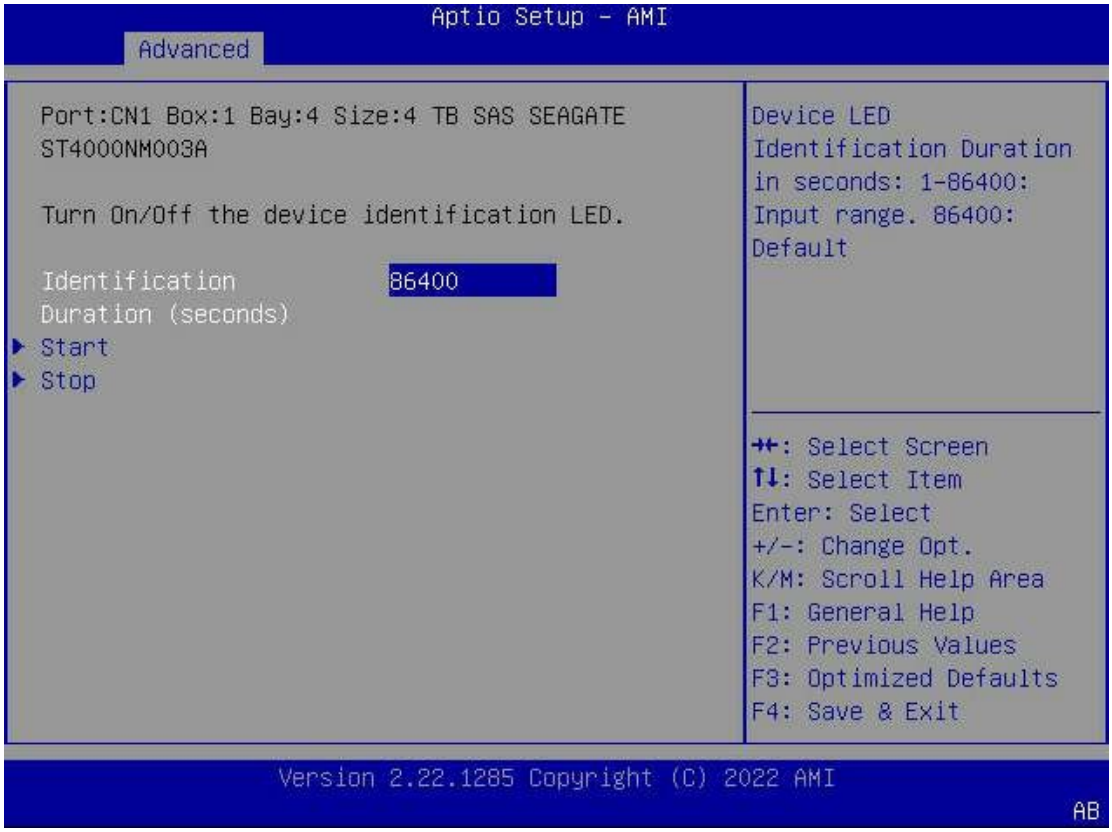

4. Next to **Identification Duration (seconds)**, enter the time of the lighting delay (unit: seconds), select **Yes**, and press **Enter** for confirmation, see Figure 3-70.

Figure 3-70 Lighting Delay

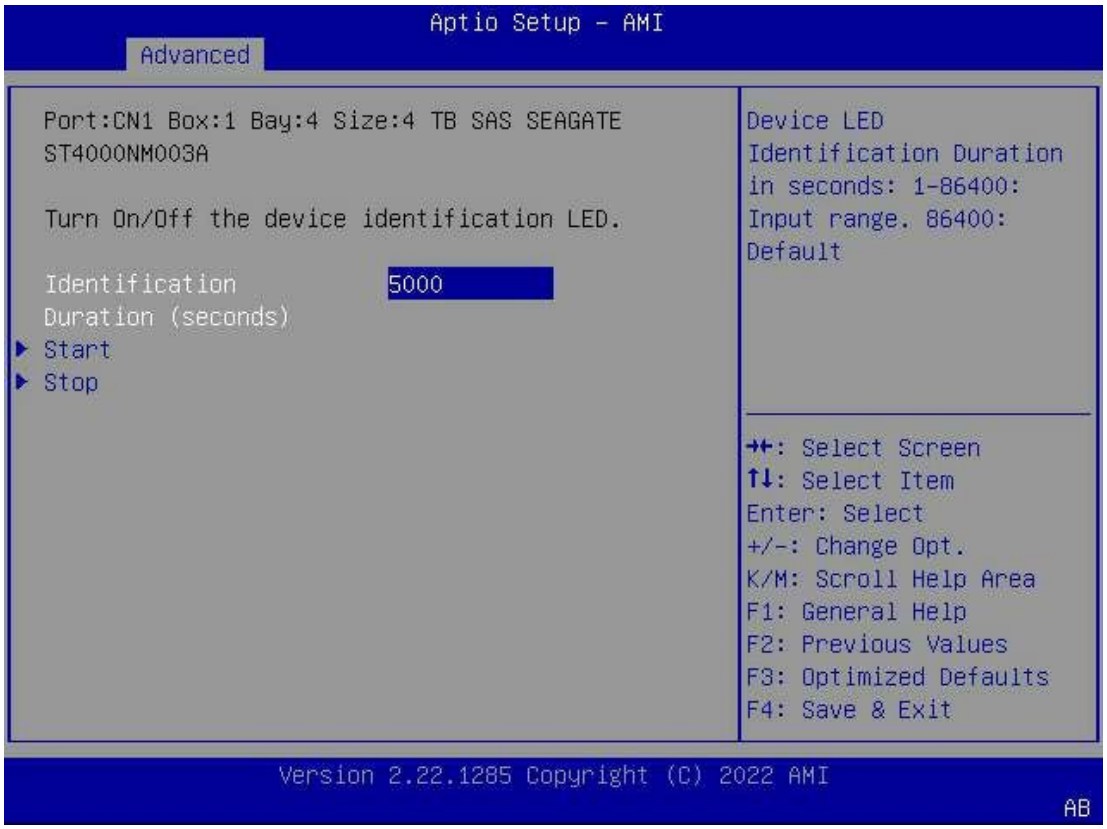

5. Use the arrow keys to select **Start**, and then press **Enter**. The red indicator of the corresponding disk is lit and continuously flashes, see Figure 3-71.

Figure 3-71 Starting Locating a Disk

Note

The indicator flashing duration is the configured lighting delay. After the lighting delay is reached, the default value 86400 s (24 hours) is restored, and the indicator goes out.

6. (Optional) To go out the disk indicator and end the locating, press **Esc** to return to the locating screen, use the arrow keys to select **Stop**, and press **Enter**, see Figure 3-72.

Figure 3-72 Ending Disk Locating

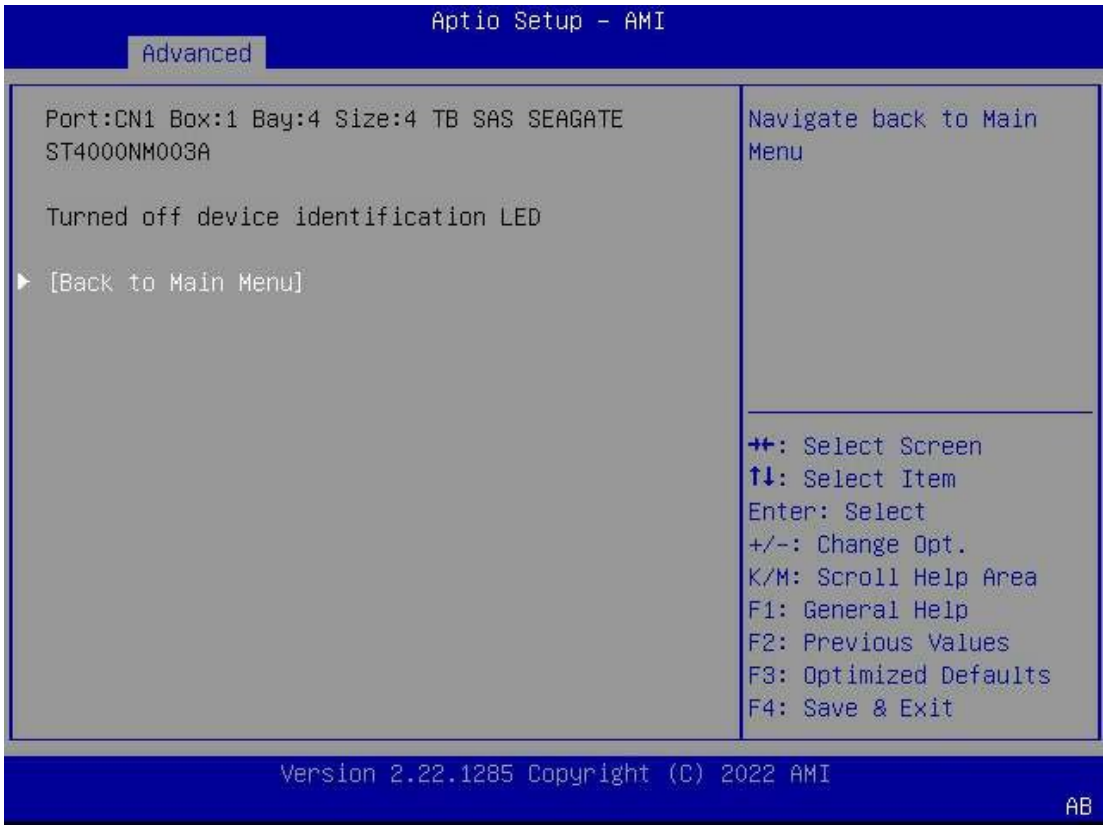

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

- Locating Multiple Disks in a Disk Group

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-73.

Figure 3-73 Array Configuration Screen

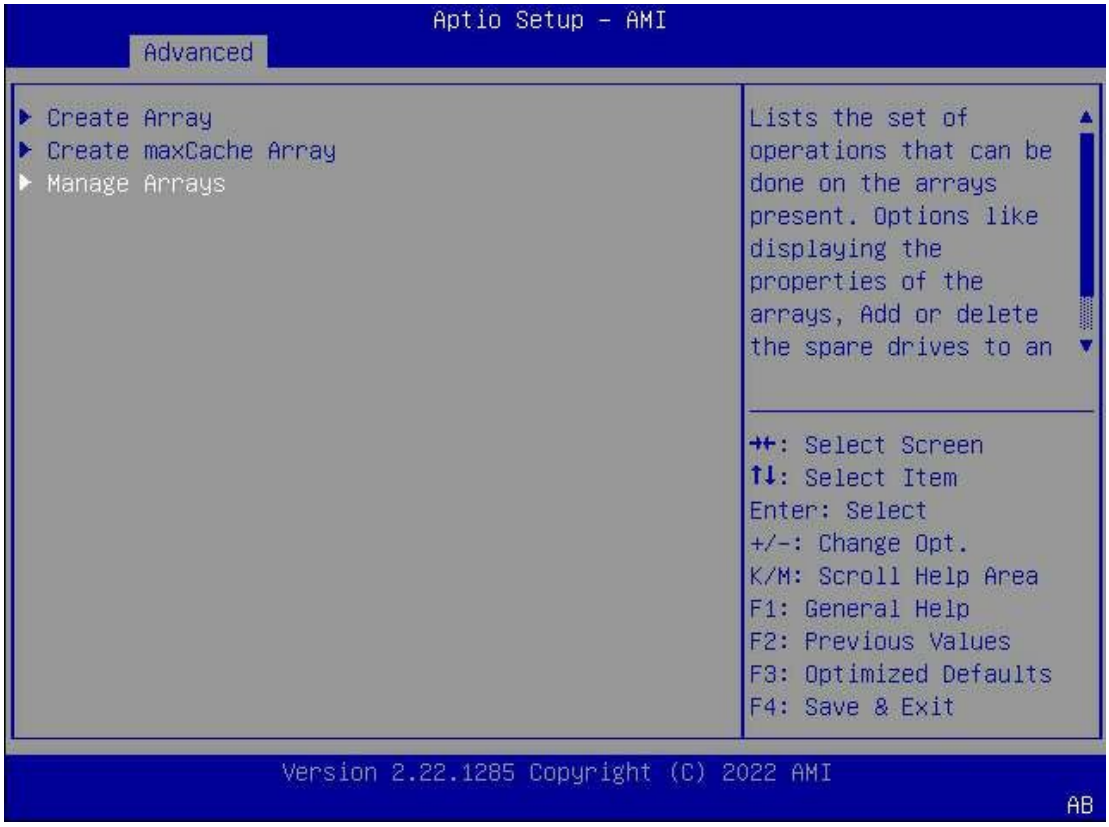

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-74.

Figure 3-74 Managing Arrays Screen

3. Use the arrow keys to select the array in which the disks to be located are placed, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-75.

Figure 3-75 Managing the Selected Array

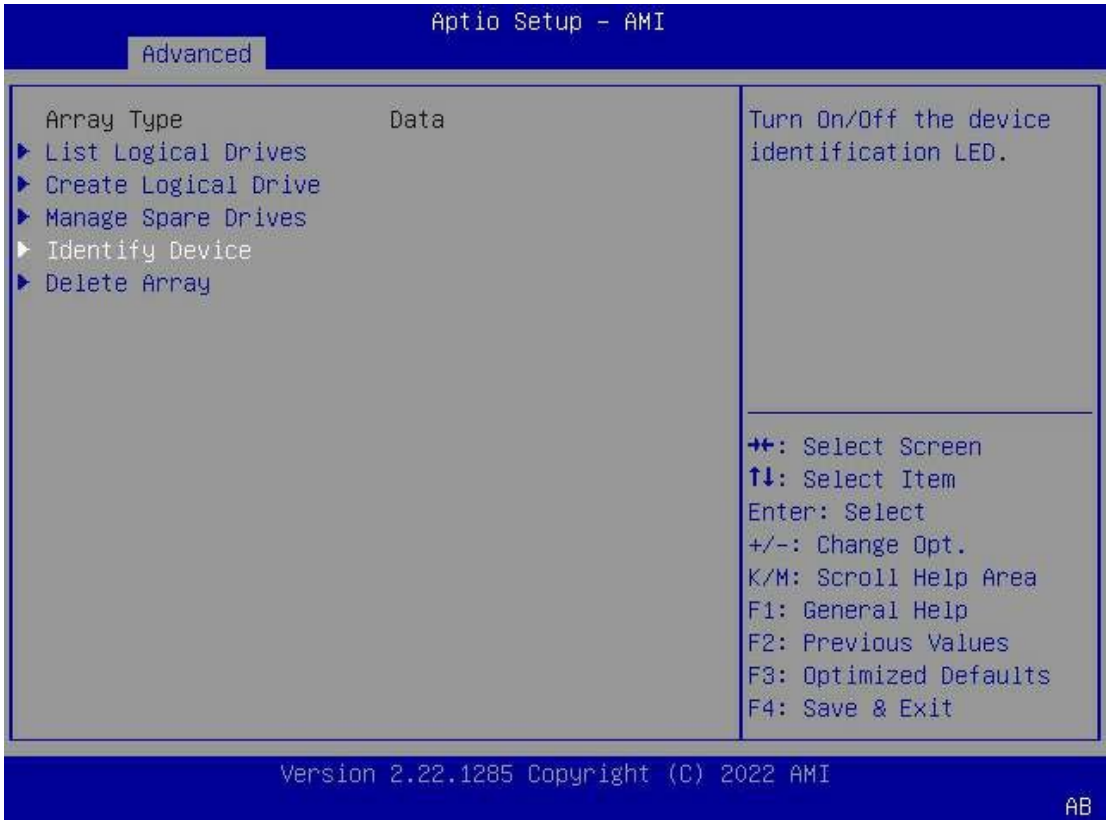

4. Use the arrow keys to select **Identify Device**, and then press **Enter**. The **Identify Device** screen is displayed, see Figure 3-76.

Figure 3-76 Identify Device Screen

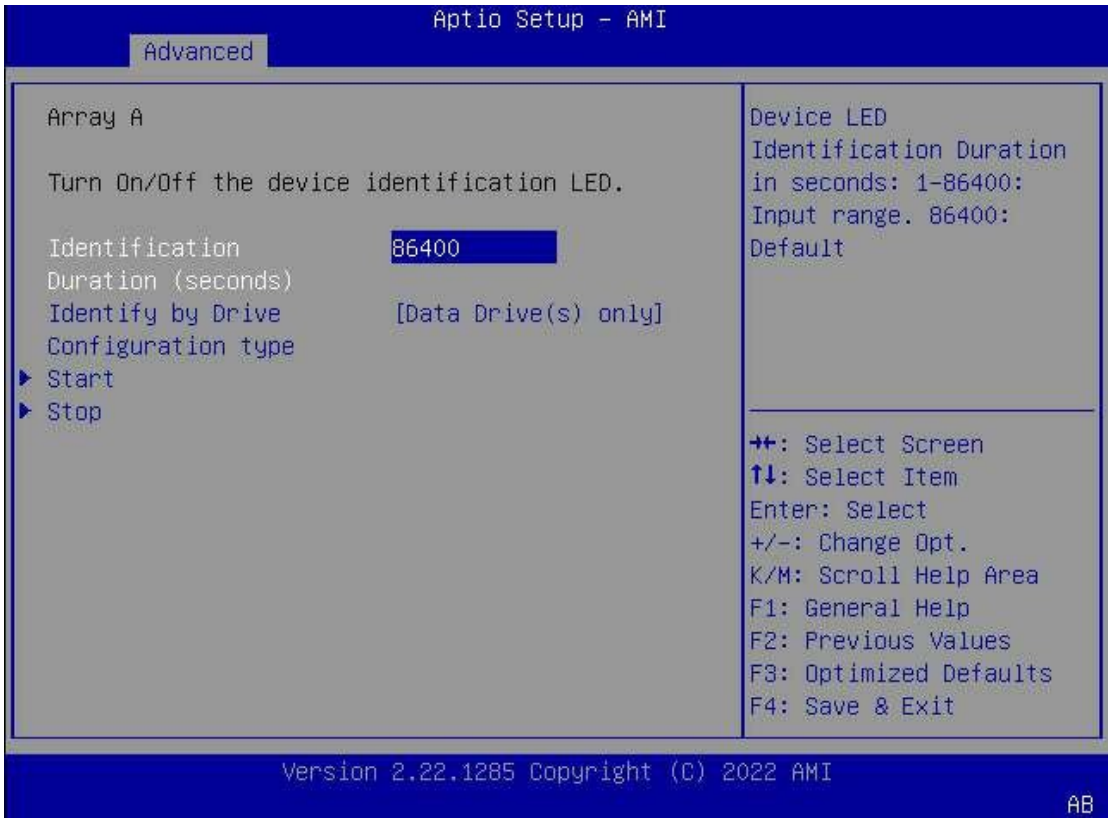

5. Next to **Identification Duration (seconds)**, enter the time of the lighting delay (unit: seconds), and then press **Enter** for confirmation, see Figure 3-77.

Figure 3-77 Lighting Delay

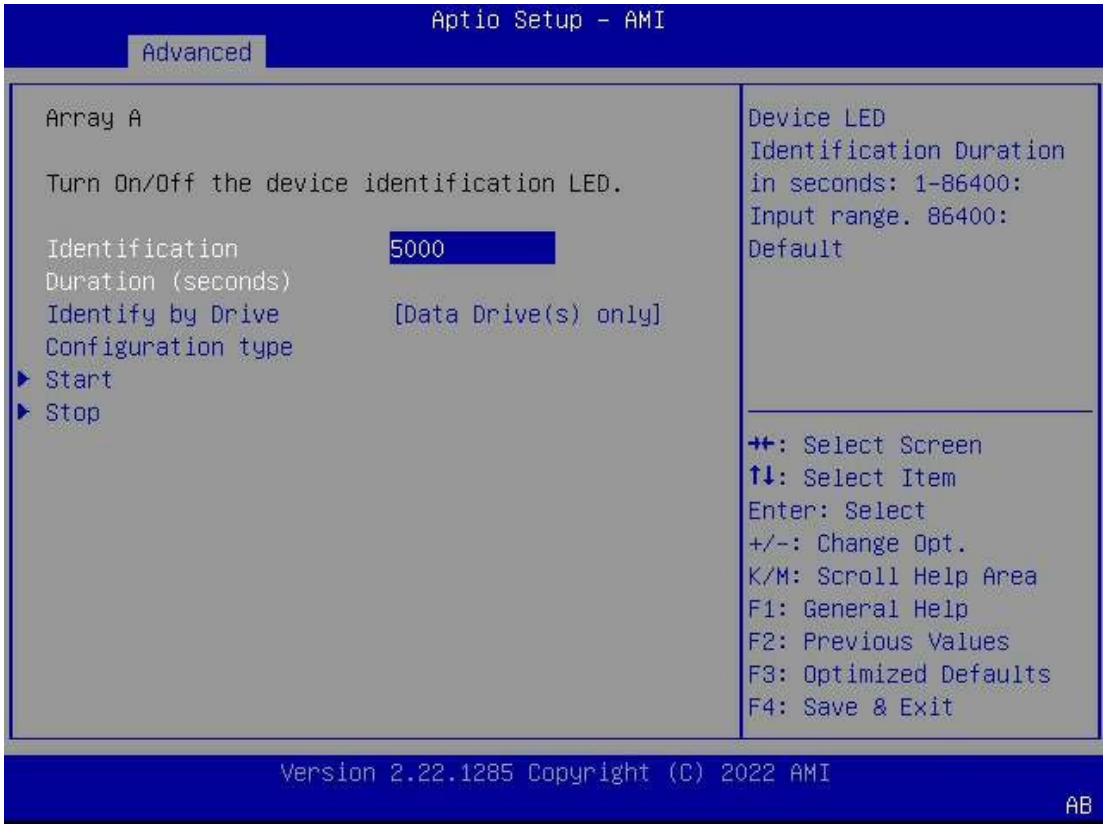

6. Use the arrow keys to select **Start**, and then press **Enter**. The red indicators of all the disks that belong to the array are lit and continuously flash, see Figure 3-78.

Figure 3-78 Starting Locating a Disk

Note

The indicators of the hot spare disks belonging to the array are also lit and flash continuously at the same time. The indicator flashing duration is the configured lighting delay. After the lighting delay is reached, the default value 86400 s (24 hours) is restored, and the indicator goes out.

7. (Optional) To go out the disk indicator and end the locating, press **ESC** to return to the locating screen, use the arrow keys to select **Stop**, and press **Enter**, see Figure 3-79.

Figure 3-79 Ending Disk Locating

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.3 Creating a Hot Spare Disk

Abstract

A hot spare disk improves the data security of a RAID array. For a description of the hot spare disk types supported by a SmartIOC 2100 RAID controller card, refer to Table 3-10.

Table 3-10 Hot Spare Disk Types

Type	Description
Dedicated	● This type of hot spare disks is exclusive to the specified one or more disk groups of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a dedicated hot spare disk temporarily takes over the faulty disk.
Auto Replace	● This type of hot spare disks provides the hot standby function for a disk group of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a hot spare disk of this type automatically replaces the faulty disk.

Prerequisite

There are sufficient idle disks on the server.

Context

When creating a hot spare disk, pay attention to the following points:

- Multiple hot spare disks can be created for a disk group, but only one type of hot spare disk can be set at a time. That is, either **Dedicated** or **Auto Replace** is specified.
- An idle disk can be set as a hot spare disk. The disk that has been used to create a RAID volume cannot be set as a hot spare disk.
- The hot spare disk must be of the same type as that of any member disk in the corresponding disk group. That is, all of them are SATA disks or SAS disks, and the hot spare disk's capacity must not be less than the maximum capacity of the member disks.
- Disk groups at all levels except RAID 0 support hot spare disks.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-80.

Figure 3-80 Array Configuration Screen

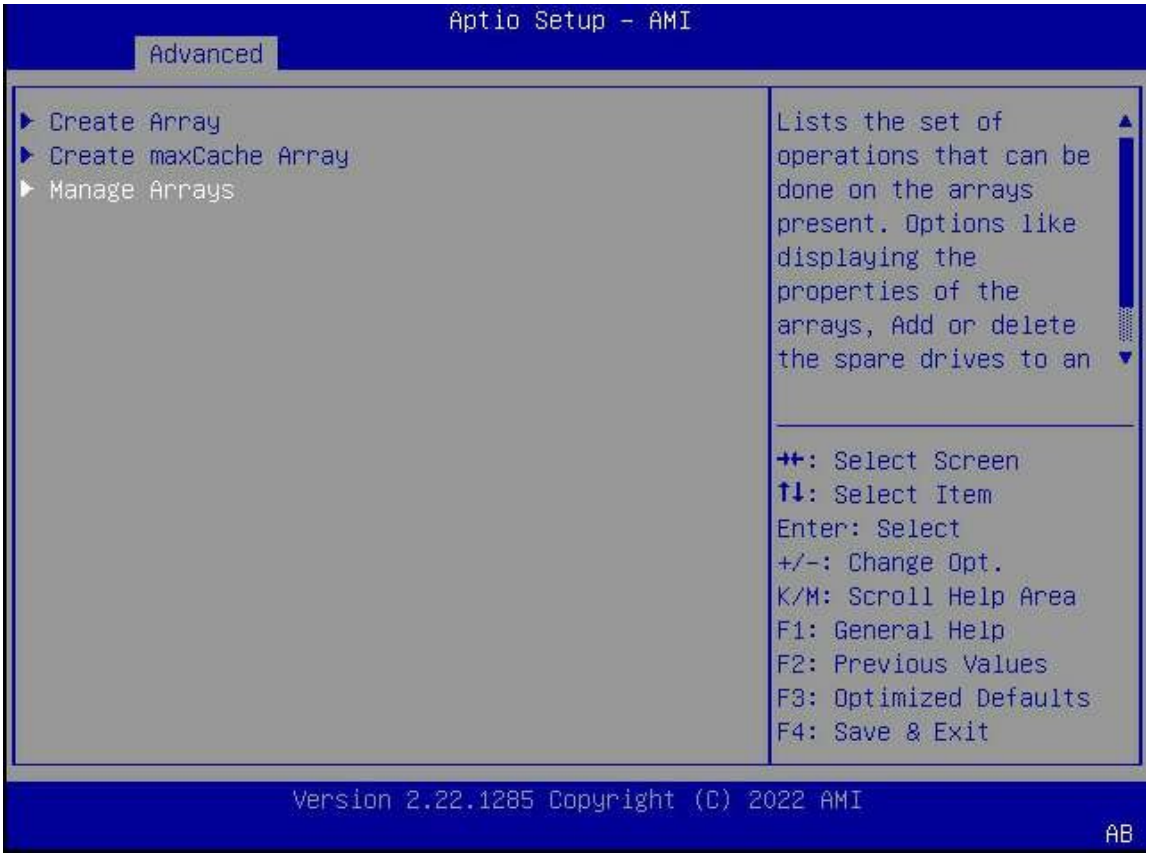

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-81.

Figure 3-81 Managing Arrays Screen

3. Use the arrow keys to select the array for which you need to create the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-82.

Figure 3-82 Managing the Selected Array

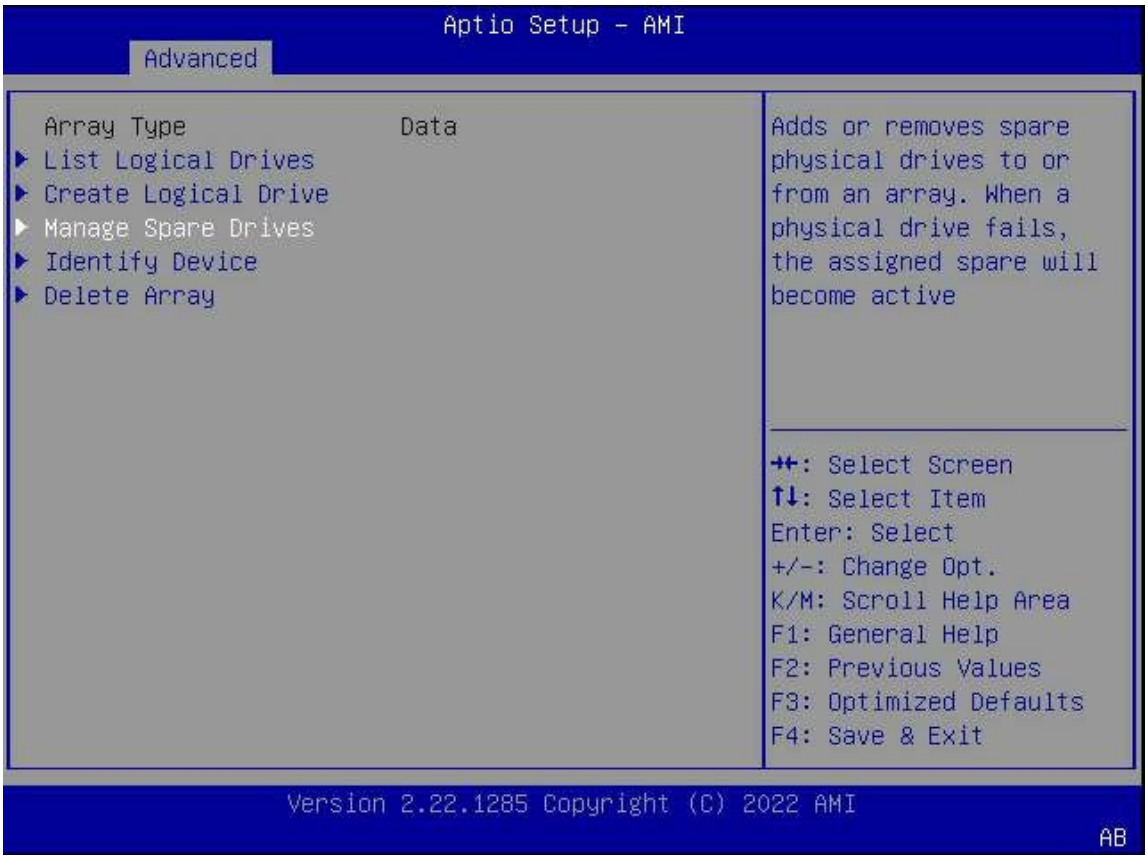

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 3-83.

Figure 3-83 Manage Spare Drives Screen

5. In accordance with your actual conditions, use the arrow keys to select the type of the hot spare disk to be created, and then press **Enter**. The screen for selecting a hot spare disk is displayed, see Figure 3-84.

Figure 3-84 Selecting a Hot Spare Disk

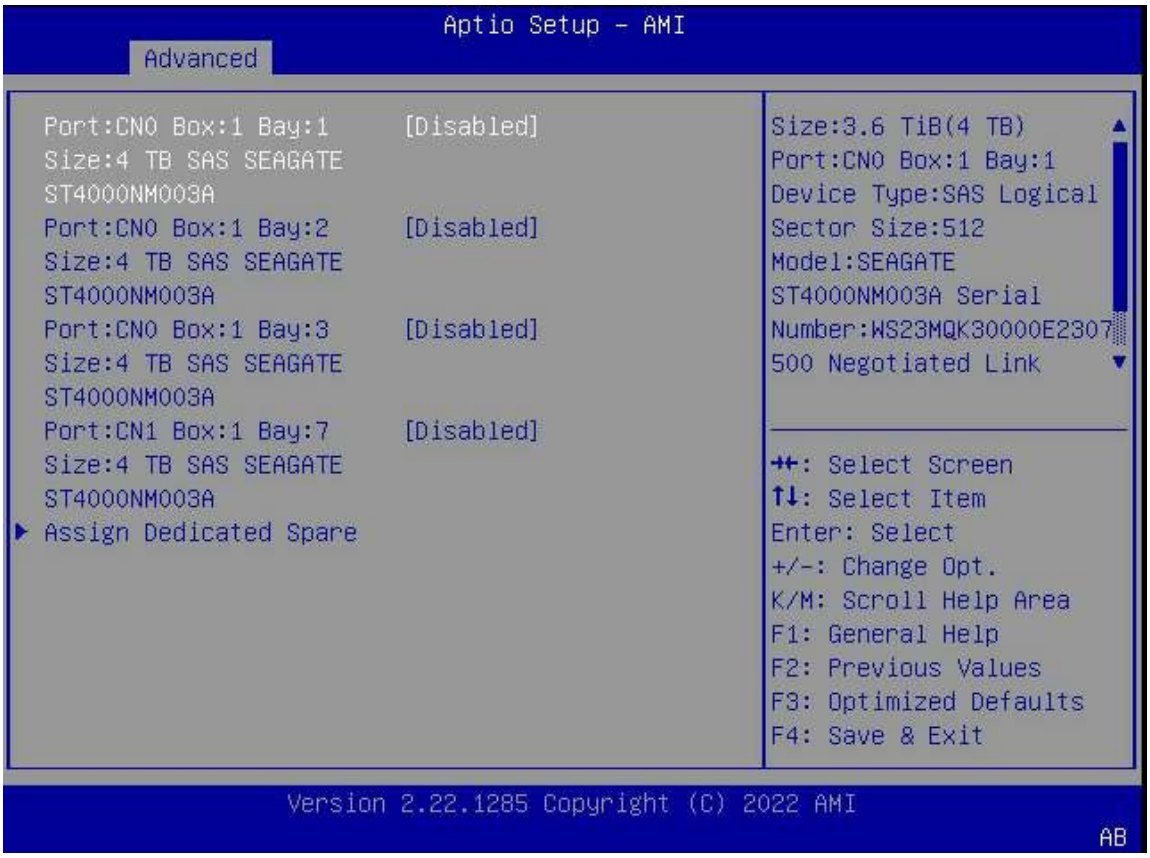

6. Use the arrow keys to select the disk to be set as a hot spare disk, press **Enter**, and set the status of the disk to **Enabled**, see Figure 3-85.

Figure 3-85 Setting Disk Status

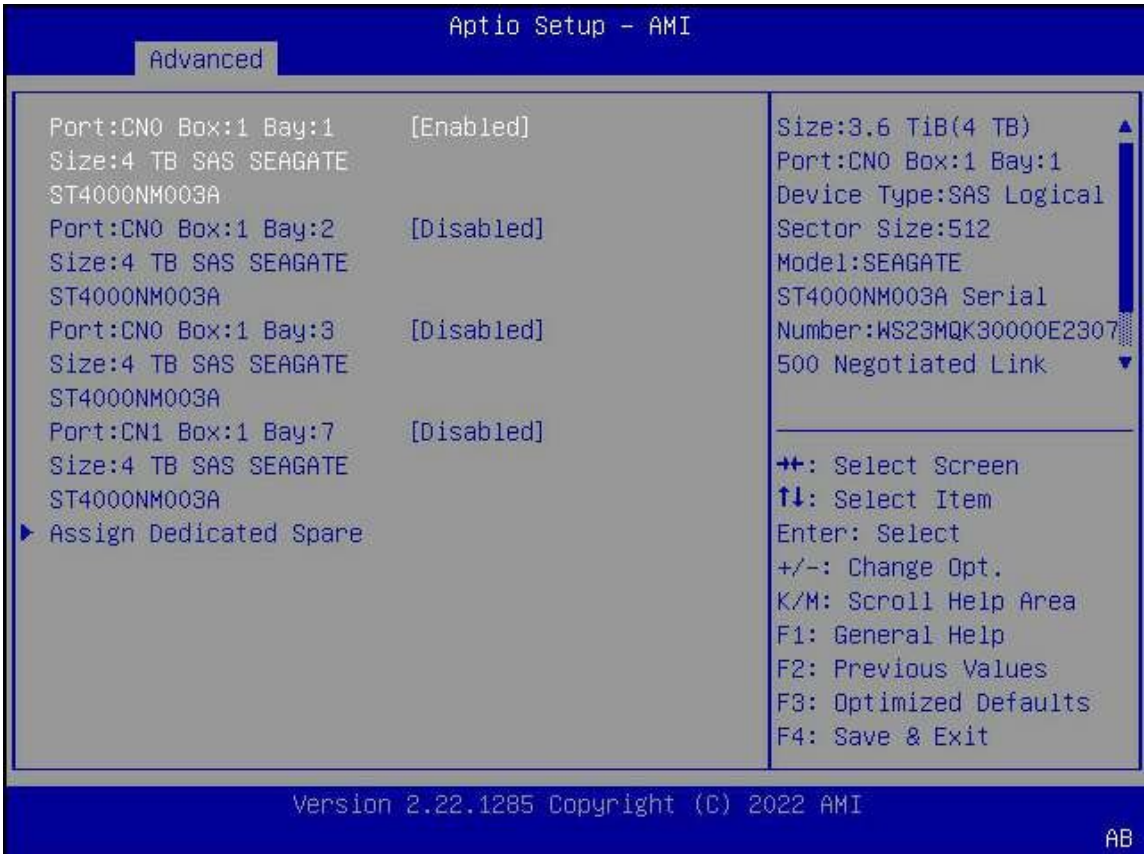

7. Use the arrow keys to select **Assign Dedicated Spare**, and then press **Enter**. The hot spare disk is successfully created, see Figure 3-86.

Figure 3-86 Hot Spare Disk Created Successfully

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.4 Changing a Hot Spare Disk

Abstract

A SmartIOC 2100 RAID controller card supports modifying the type of a hot spare disk, namely, allowing type change between **Dedicated** and **Auto Replace**.

Only one type of hot spare disk can be set at a time. That is, **Dedicated** and **Auto Replace** cannot be specified at the same time. This procedure uses changing a hot spare disk of the Dedicated type to that of the Auto Replace type as an example to describe how to perform a type change.

Note

A hot spare disk of the **Auto Replace** type can be changed to that of the **Dedicated** type by referring to this procedure.

Prerequisite

A hot spare disk is already set to the **Dedicated** type. For details, refer to "3.5.3 Creating a Hot Spare Disk".

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-87.

Figure 3-87 Array Configuration Screen

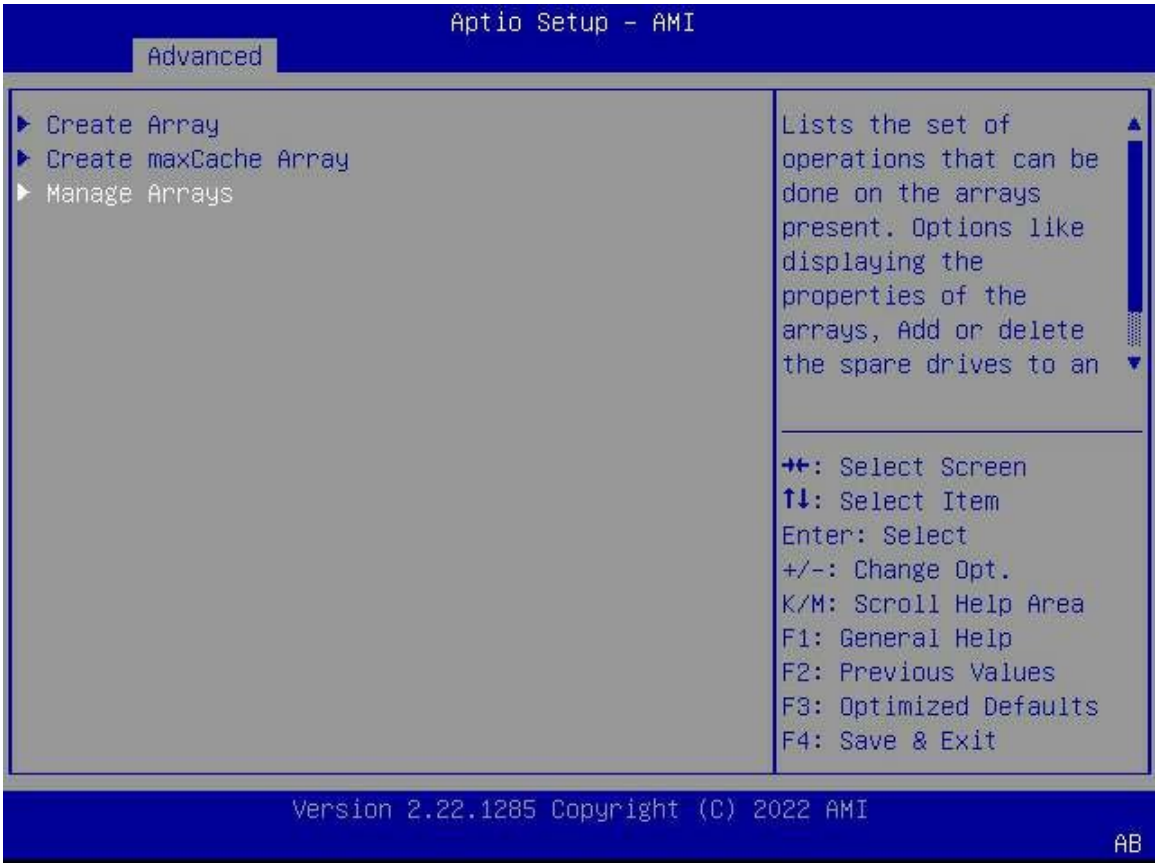

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-88.

Figure 3-88 Managing Arrays Screen

3. Use the arrow keys to select the array for which you need to modify the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-89.

Figure 3-89 Managing the Selected Array

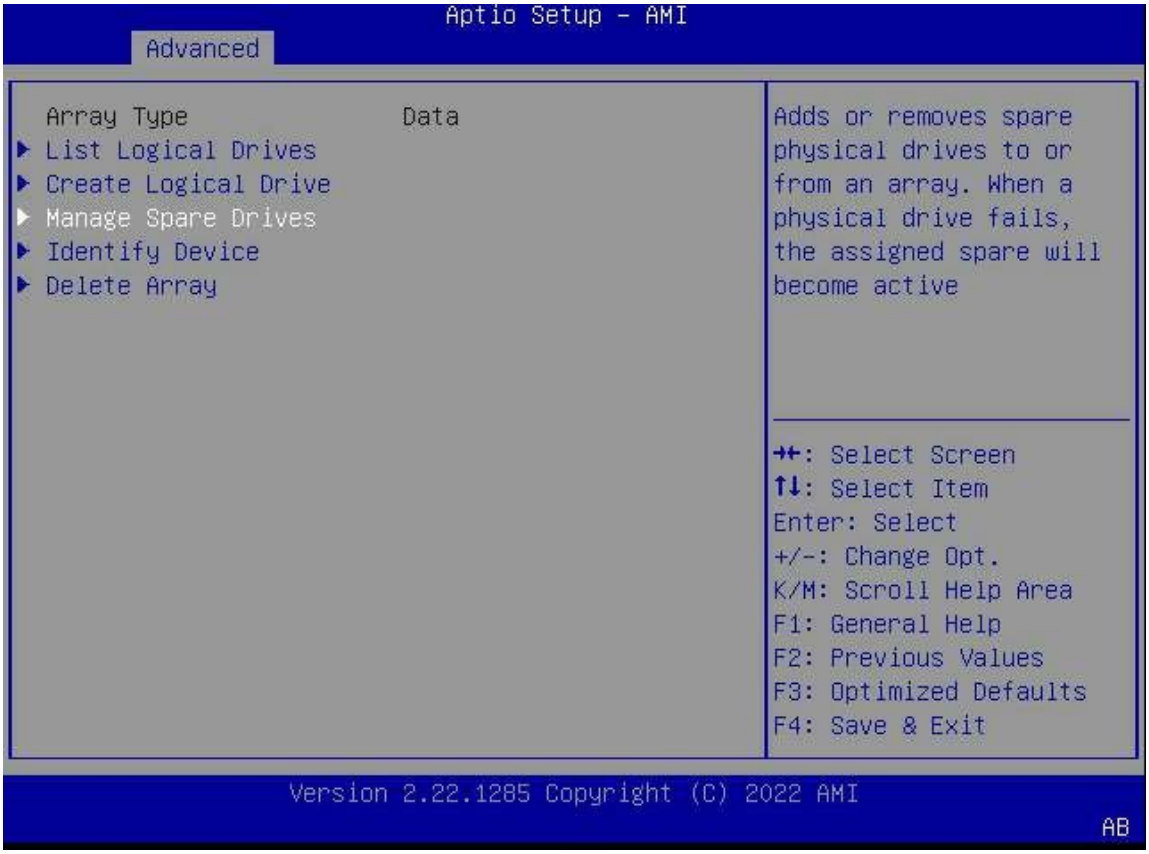

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 3-90.

Figure 3-90 Manage Spare Drives Screen

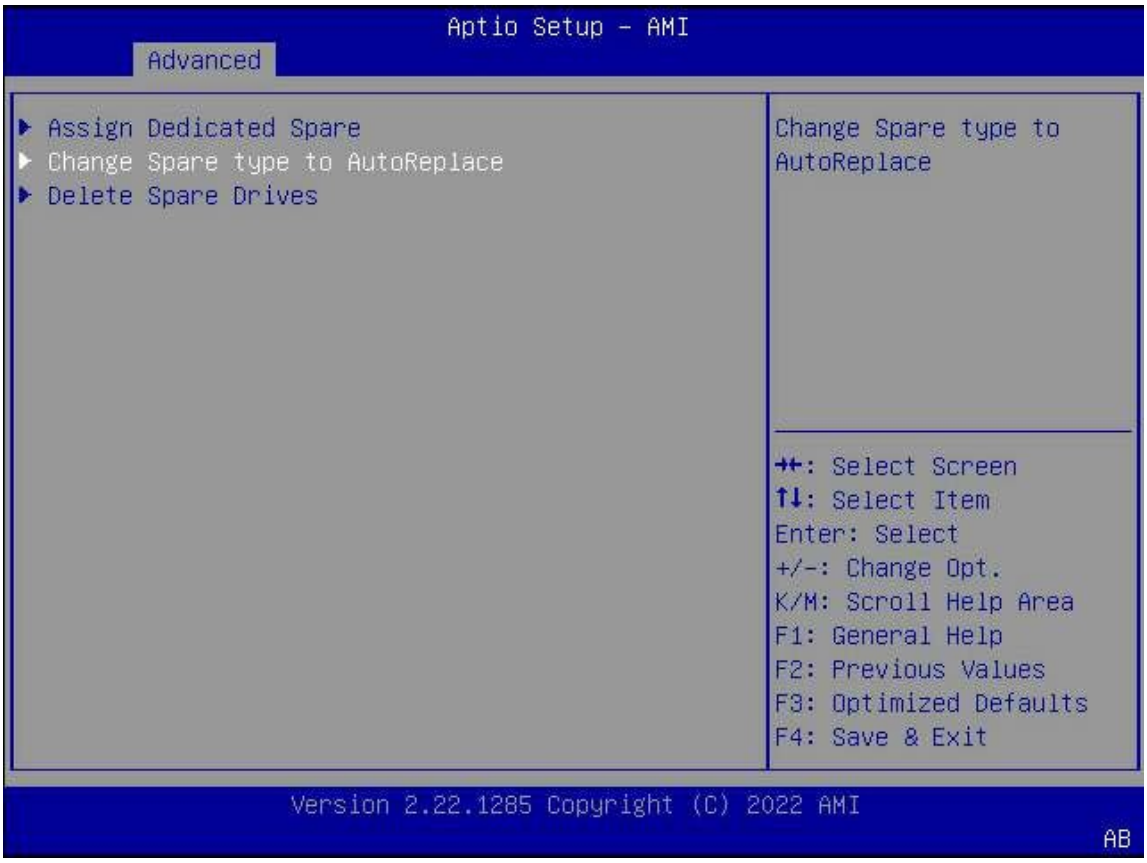

5. Use the arrow keys to select **Change Spare type to Auto Replace**, and then press **Enter**.
   The **Change Spare type to Auto Replace** screen is displayed, see Figure 3-91.

Figure 3-91 Change Spare Type to Auto Replace Screen

6. Select **Submit Changes**, and then press **Enter**. The type of hot spare disk is changed successfully, see Figure 3-92.

Figure 3-92 Hot Spare Disk Type Changed Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.5 Deleting a Hot Spare Disk

Abstract

When the number of disks of a server cannot meet the requirements, you can delete an existing hot spare disk and restore it to a common disk.

Prerequisite

A hot spare disk is already created. For details, refer to “3.5.3 Creating a Hot Spare Disk”.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-93.

Figure 3-93 Array Configuration Screen

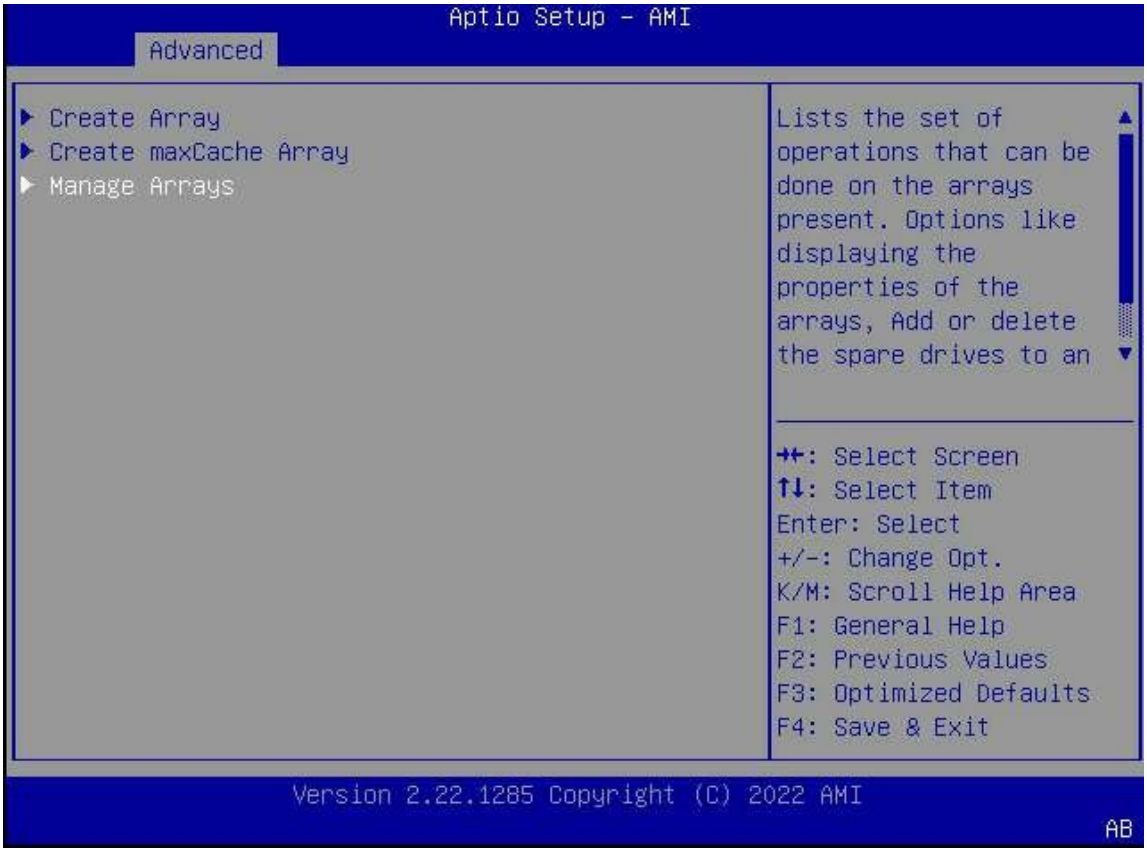

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-94.

Figure 3-94 Managing Arrays Screen

3. Use the arrow keys to select the array for which you need to delete the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-95.

Figure 3-95 Managing the Selected Array

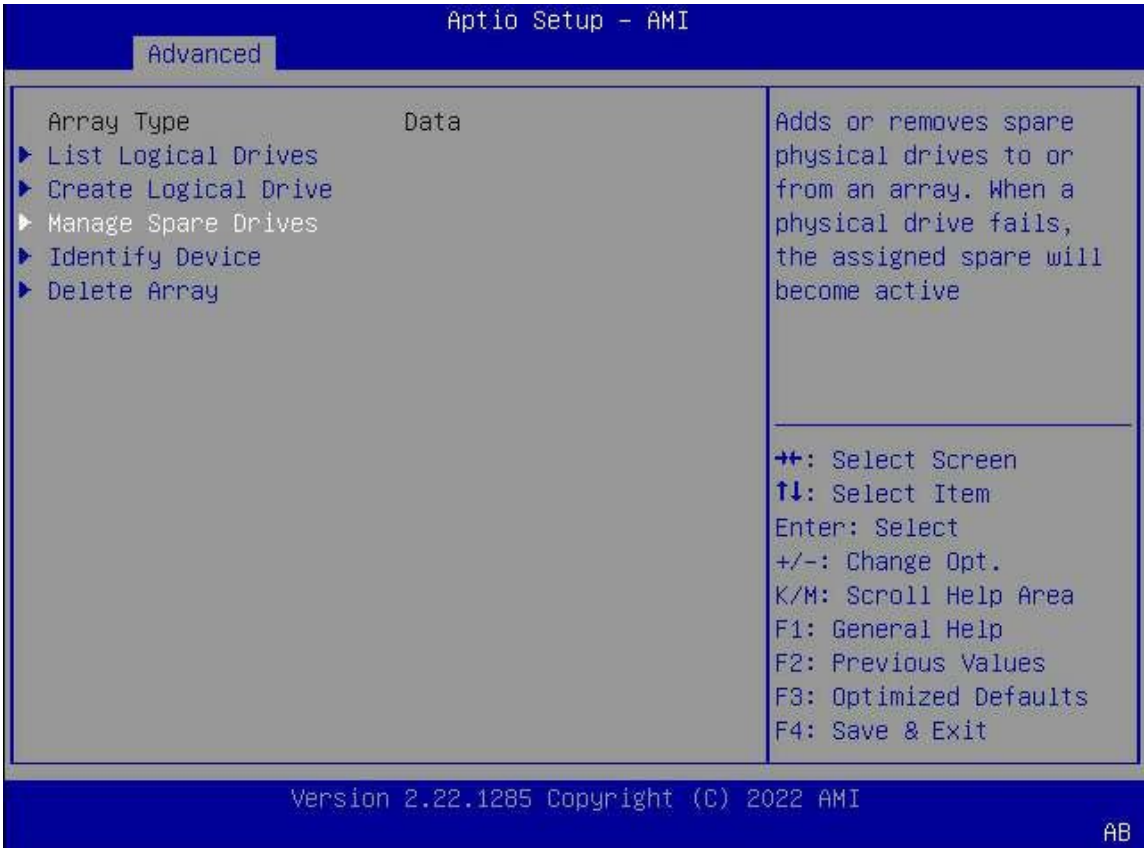

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 3-96.

Figure 3-96 Manage Spare Drives Screen

5. Use the arrow keys to select **Delete Spare Drives**, and then press **Enter**. The **Delete Spare Drives** screen is displayed, see Figure 3-97.

Figure 3-97 Delete Spare Drives Screen

6. Use the arrow keys to select the hot spare disk to be deleted, press **Enter**, and then set the status of the disk to **Enabled**, see Figure 3-98.

Figure 3-98 Setting Disk Status

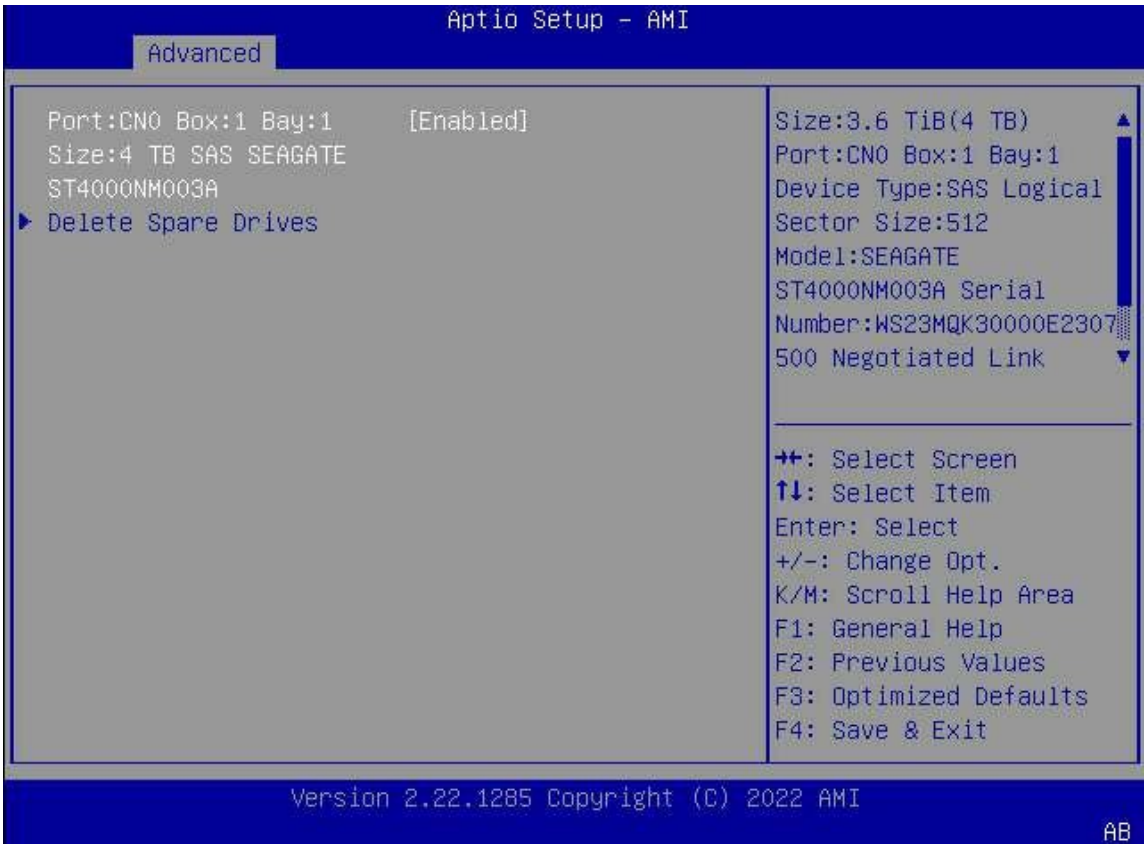

7. Use the arrow keys to select **Delete Spare Drive**, and then press **Enter**. The hot spare disk is successfully deleted, see Figure 3-99.

Figure 3-99 Hot Spare Disk Deleted Successfully

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.6 Configuring the Performance or Power Mode

Abstract

This procedure describes how to configure the power mode for a SmartIOC 2100 RAID controller card. For a description of the power modes supported by the SmartIOC 2100 RAID controller card, refer to Table 3-11.

Table 3-11 Power Mode Descriptions

Power Mode	Description
Minimum Power	In this mode, the static settings of the power are adjusted to the possible lowest value, and the power is dynamically reduced based on the working load.
Maximum Performance	In this mode, the static settings of the power are adjusted to the possible highest value, and the power is not dynamically reduced based on the working load.

Steps

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-100.

Figure 3-100 Configure Controller Settings Screen

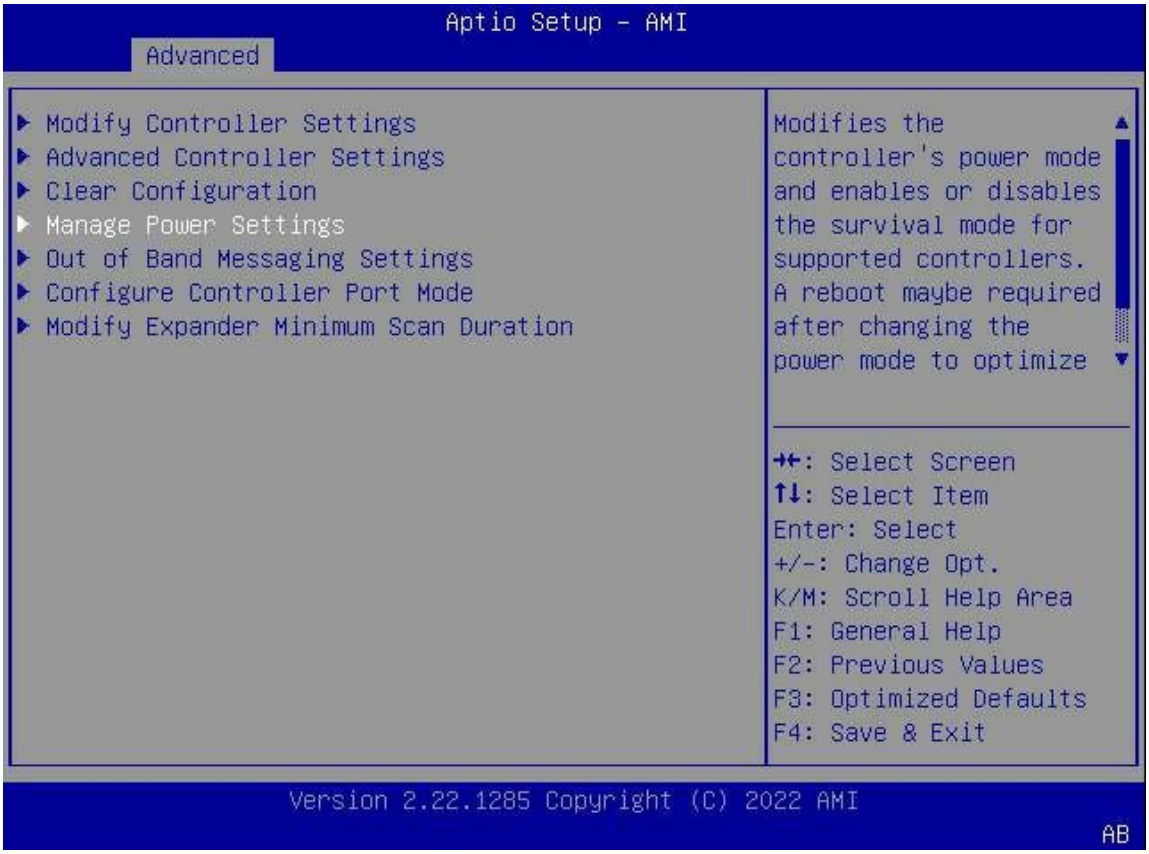

2. Use the arrow keys to select **Manage Power Settings**, and then press **Enter**. The **Manage Power Settings** screen is displayed, see Figure 3-101.

Figure 3-101 Manage Power Settings Screen

3. Use the arrow keys to select **Power Mode**, and then press **Enter**. The **Power Mode** menu is displayed, see Figure 3-102.

Figure 3-102 Power Mode Menu

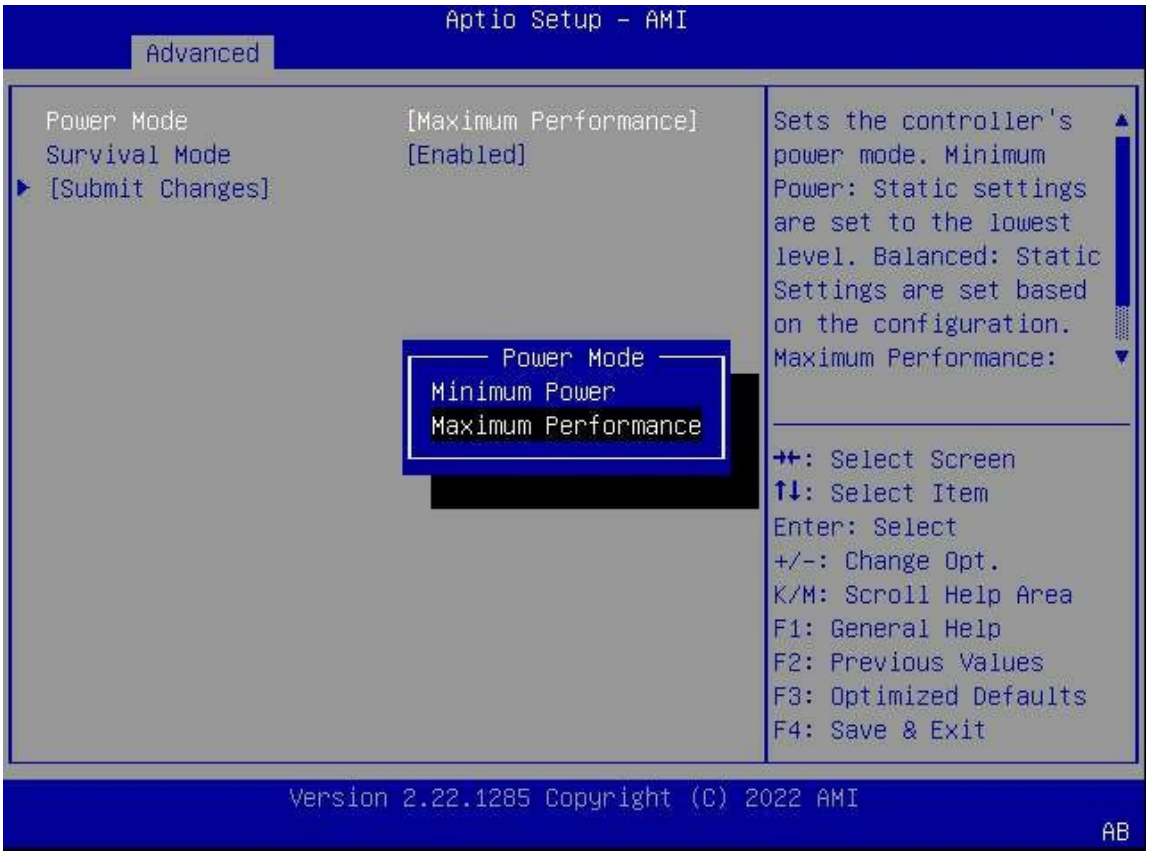

4. In accordance with your actual conditions, use the arrow keys to select the power mode to be applied, and then press **Enter** for confirmation, see Figure 3-103.

Figure 3-103 Selecting a Power Mode

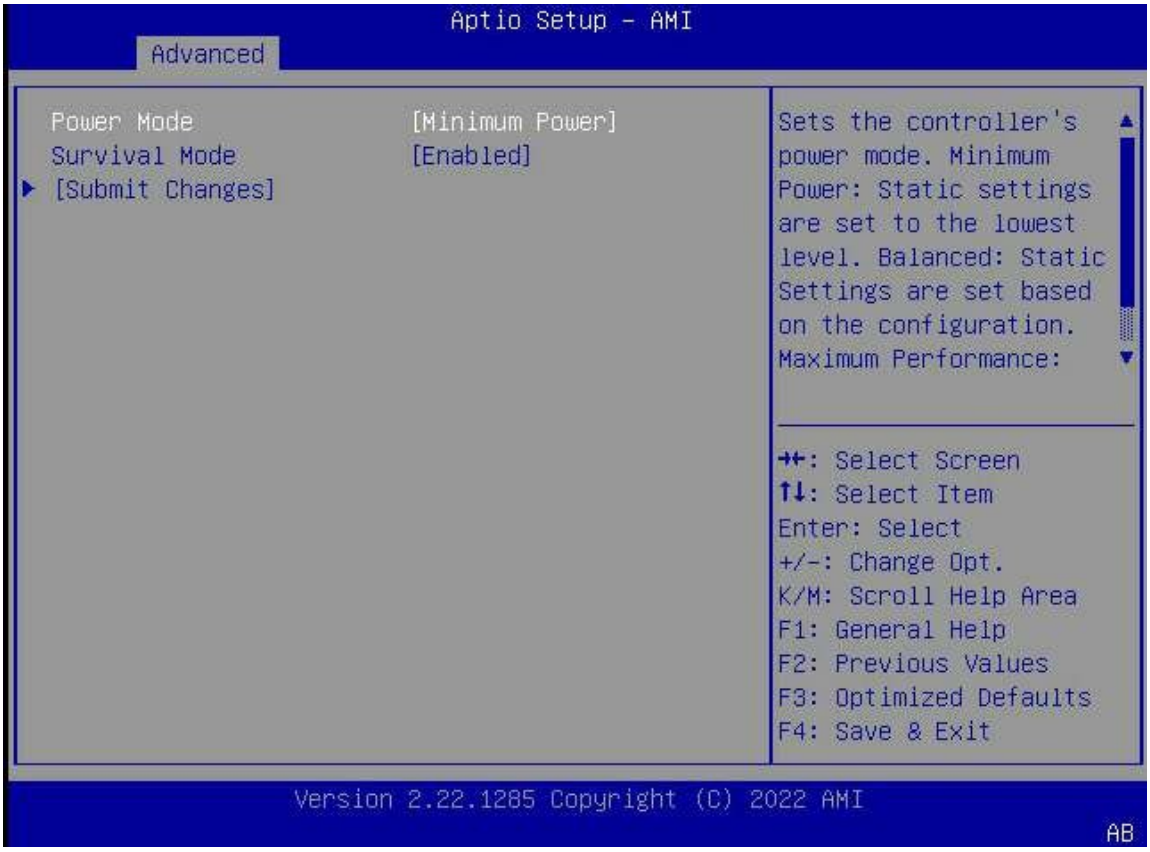

5. (Optional) Determine whether to enable **Survival Mode** as required.

Note

By default, **Survival Mode** is enabled, indicating that when the working temperature of the power supply exceeds the threshold, the RAID controller card is allowed to switch to the energy saving mode, but it may cause performance deterioration.

a. Use the arrow keys to select **Survival Mode**, and then press **Enter**. The **Survival Mode** menu is displayed, see Figure 3-104.

Figure 3-104 Survival Mode Menu

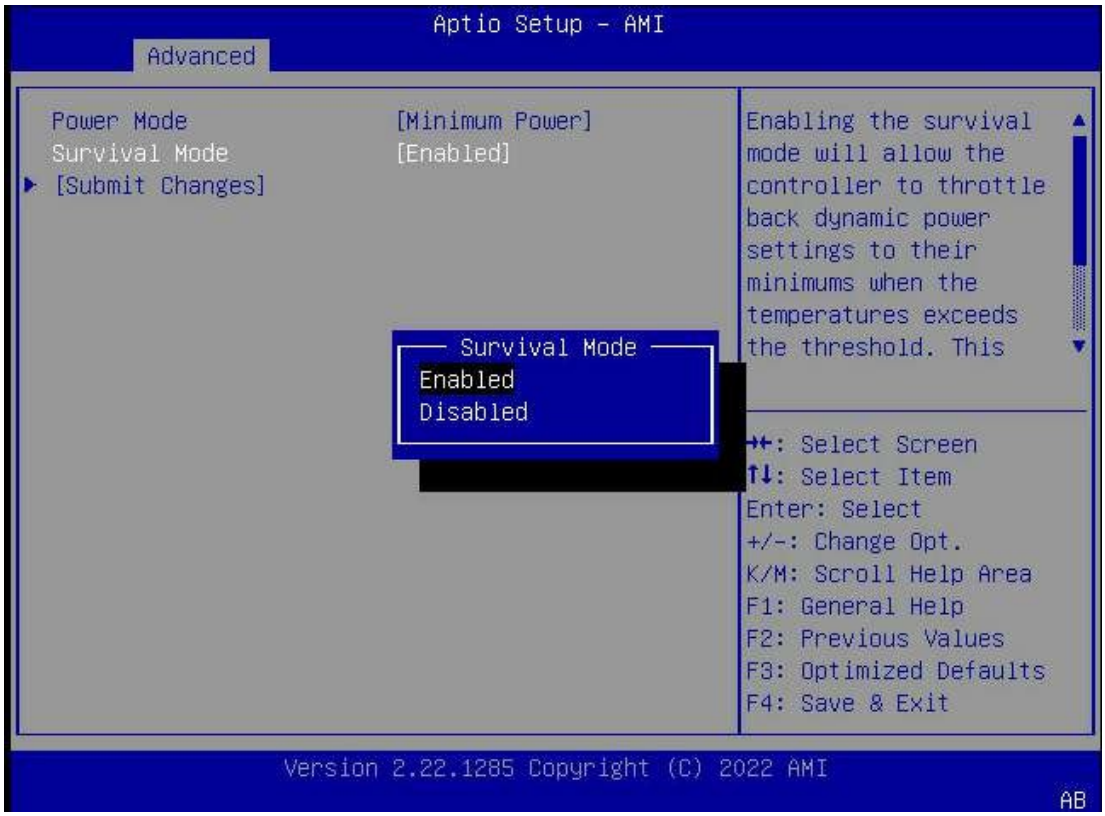

b. Use the arrow keys to select **Disabled**, and then press **Enter**. The **Survival Mode** is disabled, see Figure 3-105.

Figure 3-105 Disabling Survival Mode

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The power mode is set successfully, see Figure 3-106.

Figure 3-106 Power Mode Configured Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.7 Deleting a RAID volume

Abstract

When a server no longer needs a RAID volume, you can delete the RAID volume to release the disk space.

Notice

- The data that is lost during deletion of the RAID volume cannot be restored. Therefore, you must make sure that you have backed up important data before deleting the volume.
- If the RAID logical volume to be deleted is the only logical volume on the current array, the array is also deleted after the RAID logical volume is deleted.

Prerequisite

A RAID volume is created successfully. For details, refer to “3.3.2 Creating a RAID Volume"

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-107.

Figure 3-107 Array Configuration Screen

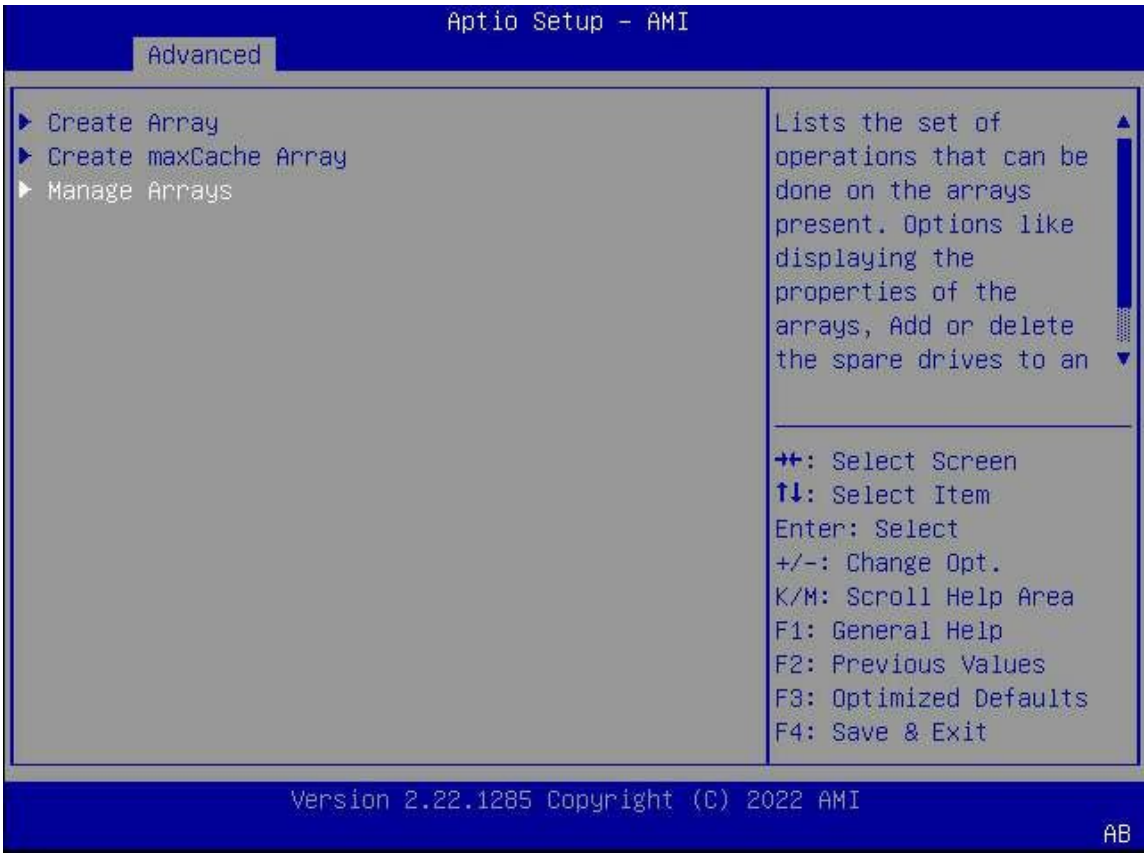

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-108.

Figure 3-108 Managing Arrays Screen

3. Use the arrow keys to select the array for which you need to delete a logical volume, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-109.

Figure 3-109 Managing the Selected Array

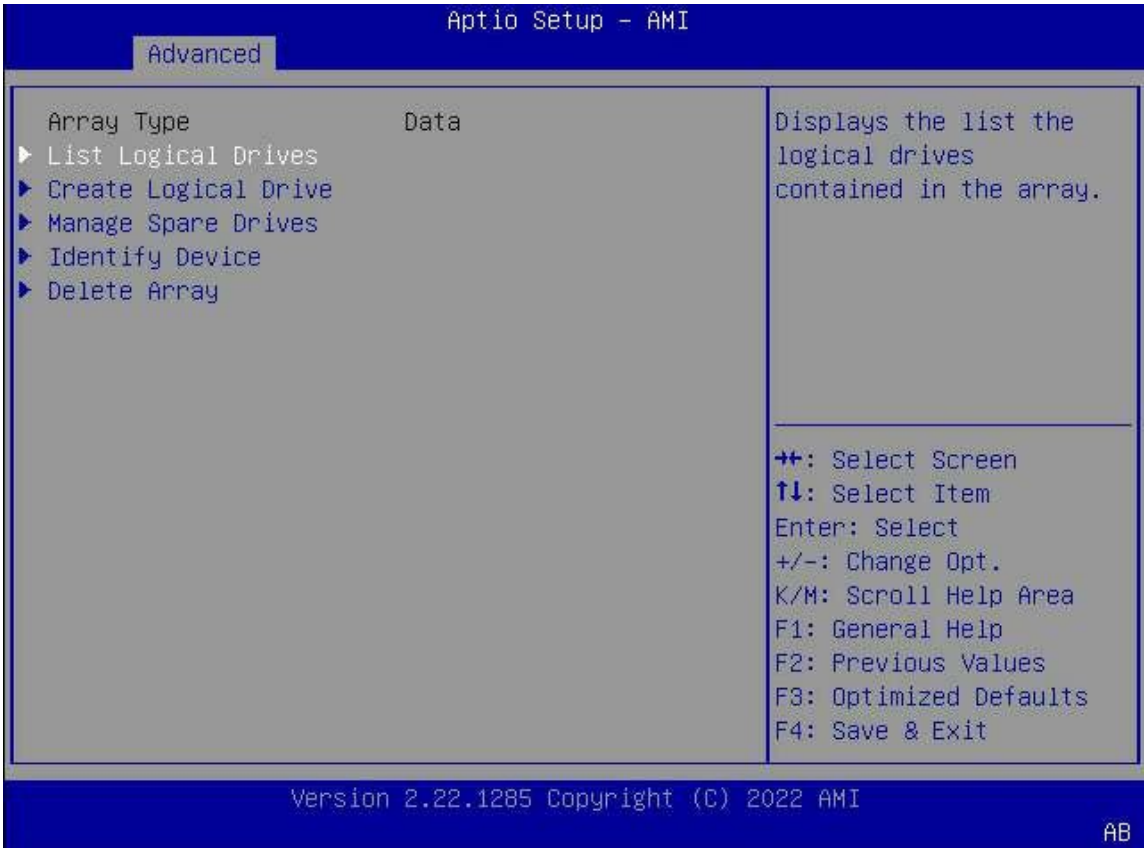

4. Use the arrow keys to select **List Logical Drives**, and then press **Enter**. The **List Logical Drives** screen is displayed, see Figure 3-110.

Figure 3-110 List Logical Drives Screen

5. Use the arrow keys to select the logical volume to be deleted, and then press **Enter**. The screen for managing the selected logical volume is displayed, see Figure 3-111.

Figure 3-111 Managing the Selected Logical Volume

6. (Optional) To view the details of the logical volume, use the arrow keys to select **Logical Drive Details**, and then press **Enter**, see Figure 3-112.

Figure 3-112 Viewing Logical Volume Details Screen

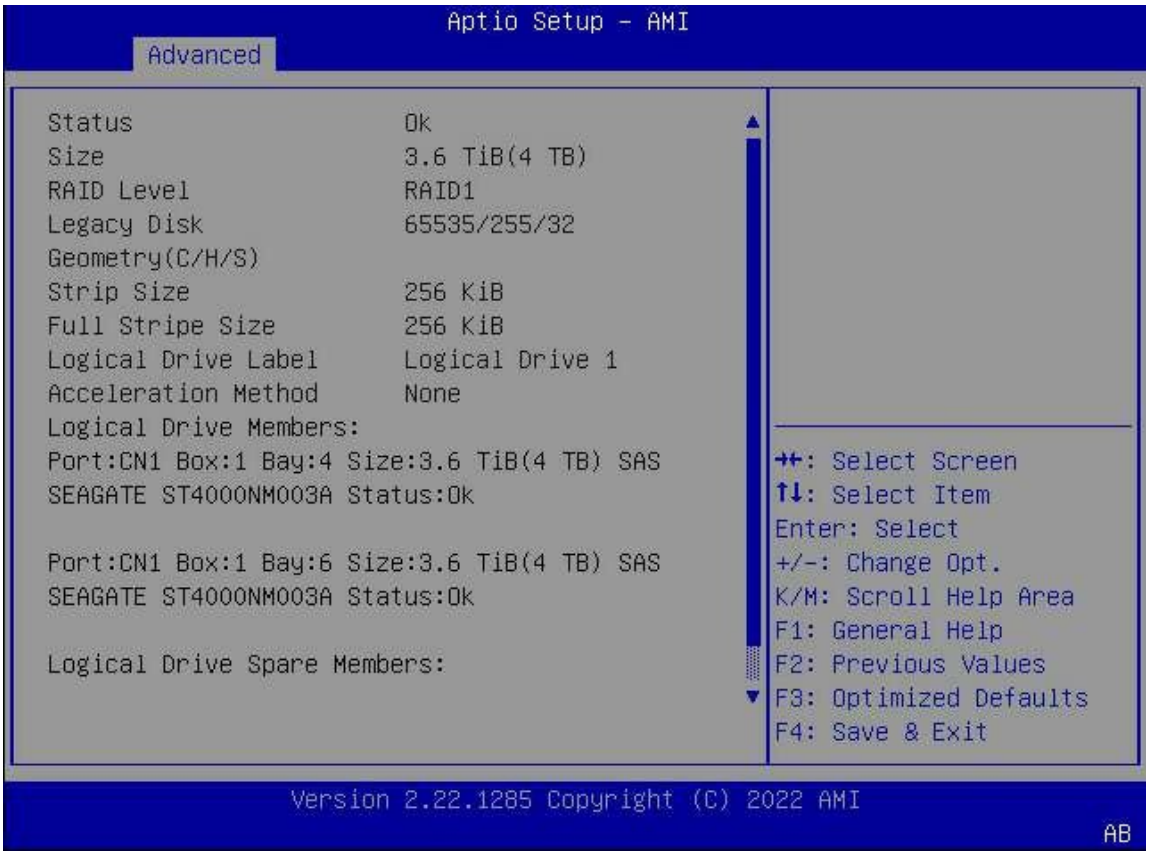

7. On the logical volume management screen, use the arrow keys to select **Delete Logical Drive**, and then press **Enter**. The confirmation screen for RAID volume deletion is displayed, see Figure 3-113.

Figure 3-113 Confirmation Screen for RAID Volume Deletion

8. Select **Submit Changes**, and then press **Enter**. The logical volume is deleted successfully, see Figure 3-114.

Figure 3-114 Logical Volume Deleted Successfully

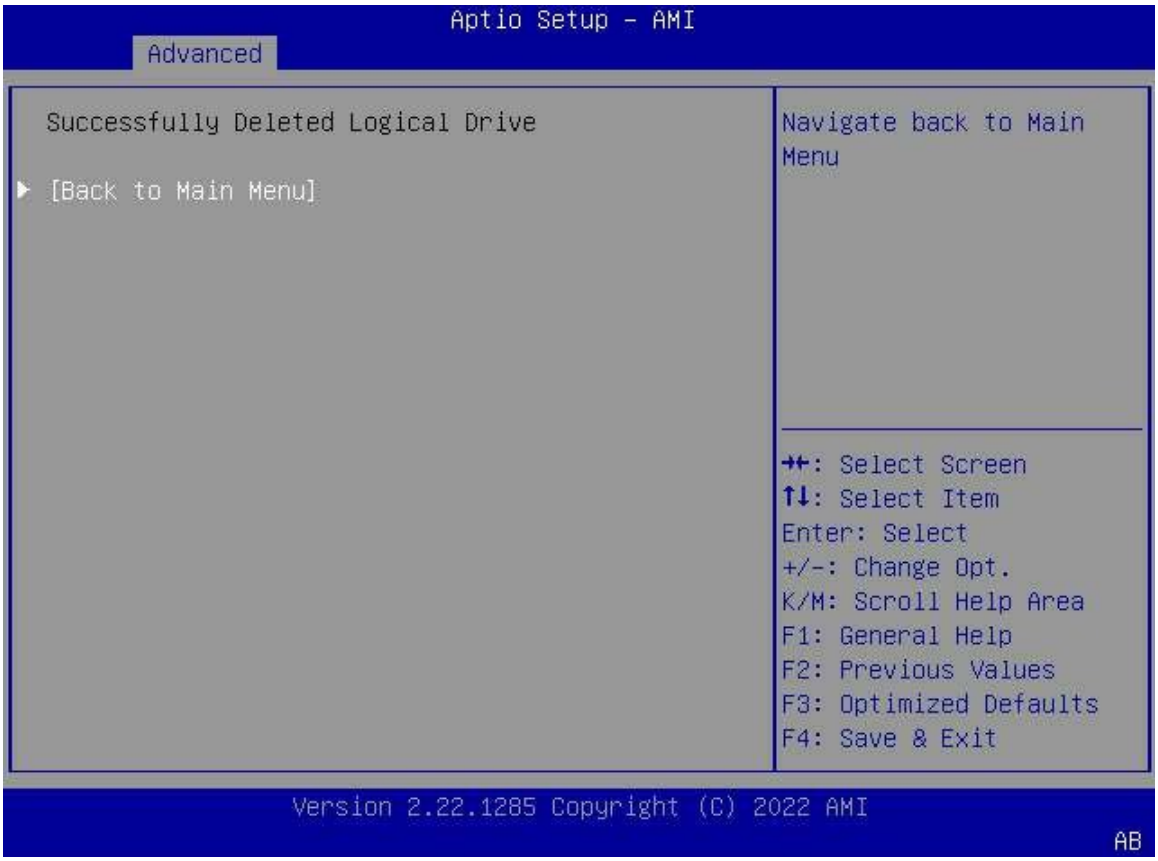

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.8 Deleting a Disk Group

Abstract

When a server no longer needs a disk group (array), you can delete it to release the disk space.

Notice

- When an array is deleted, the RAID logical volume built on it is also deleted, and the data lost during the deletion cannot be restored. Therefore, it is required to make sure that you have backed up important data before the deletion.
- If the array to be deleted is the only array of the RAID controller card, the configurations related to the RAID controller card are also cleared, and the default configurations are restored.

Prerequisite

A RAID volume is created successfully to form a disk group. For details, refer to “3.3.2 Creating a RAID Volume”.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 3-115.

Figure 3-115 Array Configuration Screen

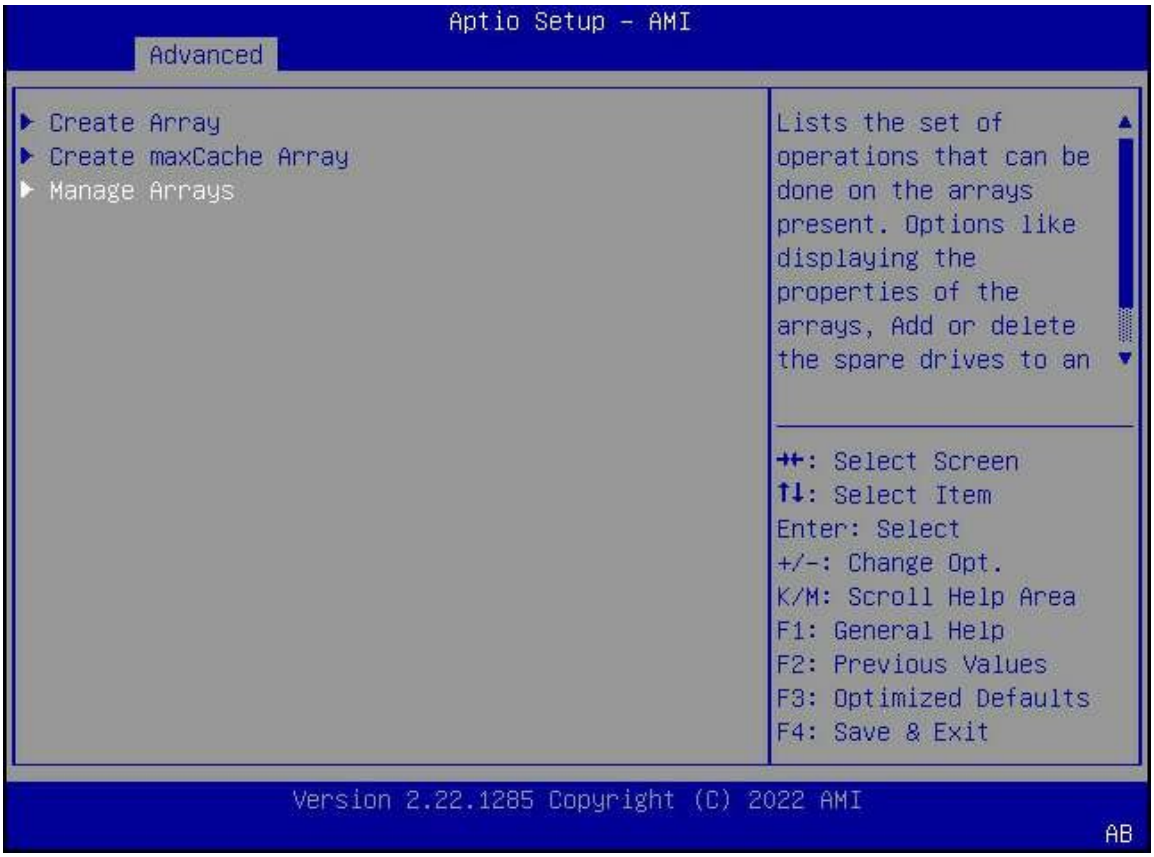

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 3-116.

Figure 3-116 Managing Arrays Screen

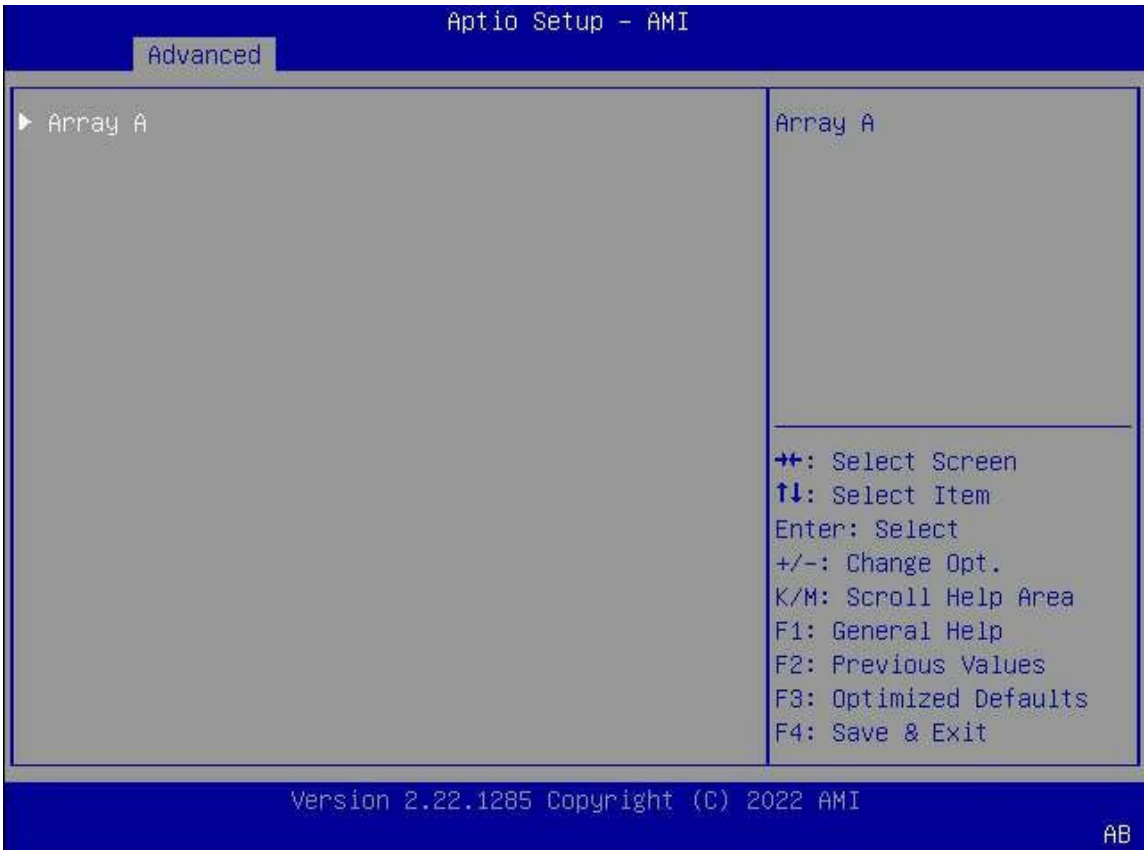

3. Use the arrow keys to select the array to be deleted, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 3-117.

Figure 3-117 Managing the Selected Array

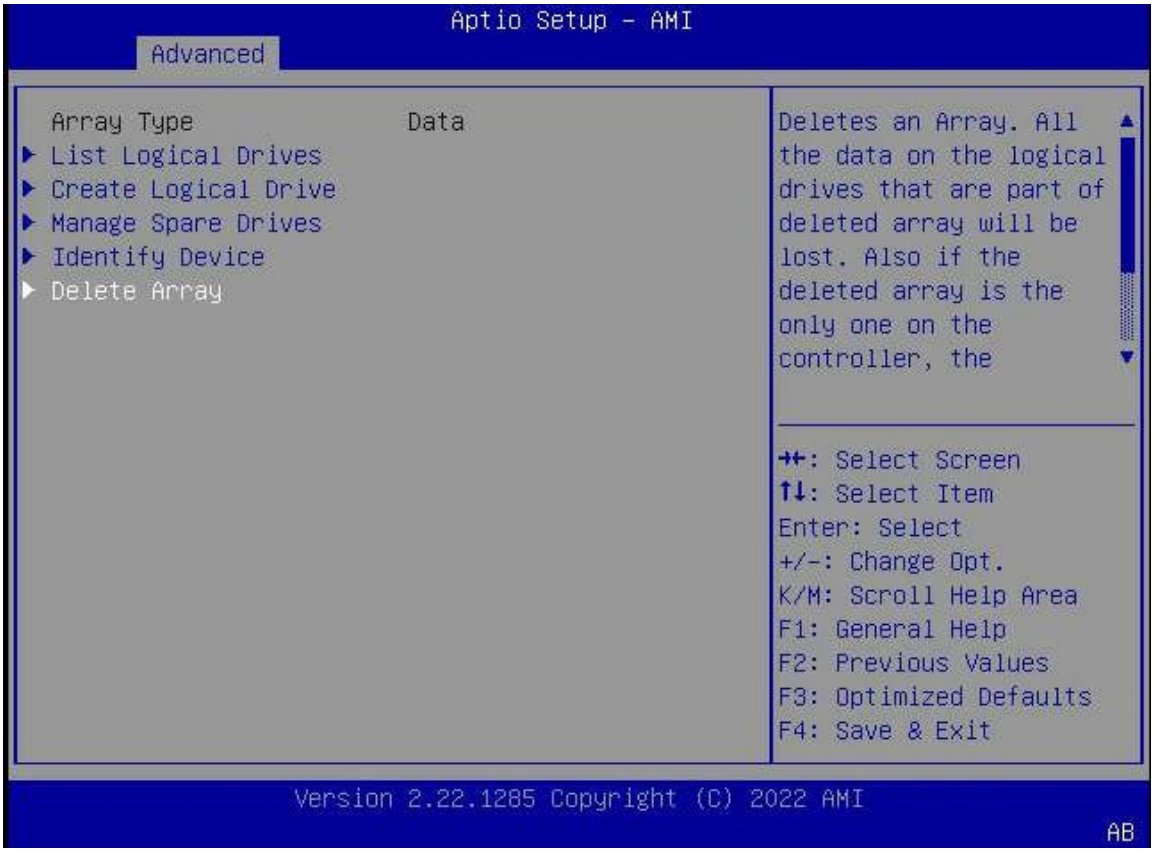

4. Use the arrow keys to select **Delete Array**, and then press **Enter**. The **Delete Array** screen is displayed, see Figure 3-118.

Figure 3-118 Delete Array Screen

5. Use the arrow keys to select **Submit Changes**, and press **Enter**. The array is deleted successfully, see Figure 3-119.

Figure 3-119 Array Deleted Successfully

6. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.9 Clearing RAID Configuration Information

Abstract

This procedure describes how to clear all configuration information that is already created on a SmartIOC 2100 RAID controller card.

Notice

The data that is lost during clearing of the configuration information on the RAID controller card cannot be restored. Therefore, it is required to make sure that you have backed up important data before the clearing operation.

Prerequisite

A RAID volume is created successfully and it has the corresponding RAID configuration information. For details, refer to “3.3.2 Creating a RAID Volume”.

Steps

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-120.

Figure 3-120 Configure Controller Settings Screen

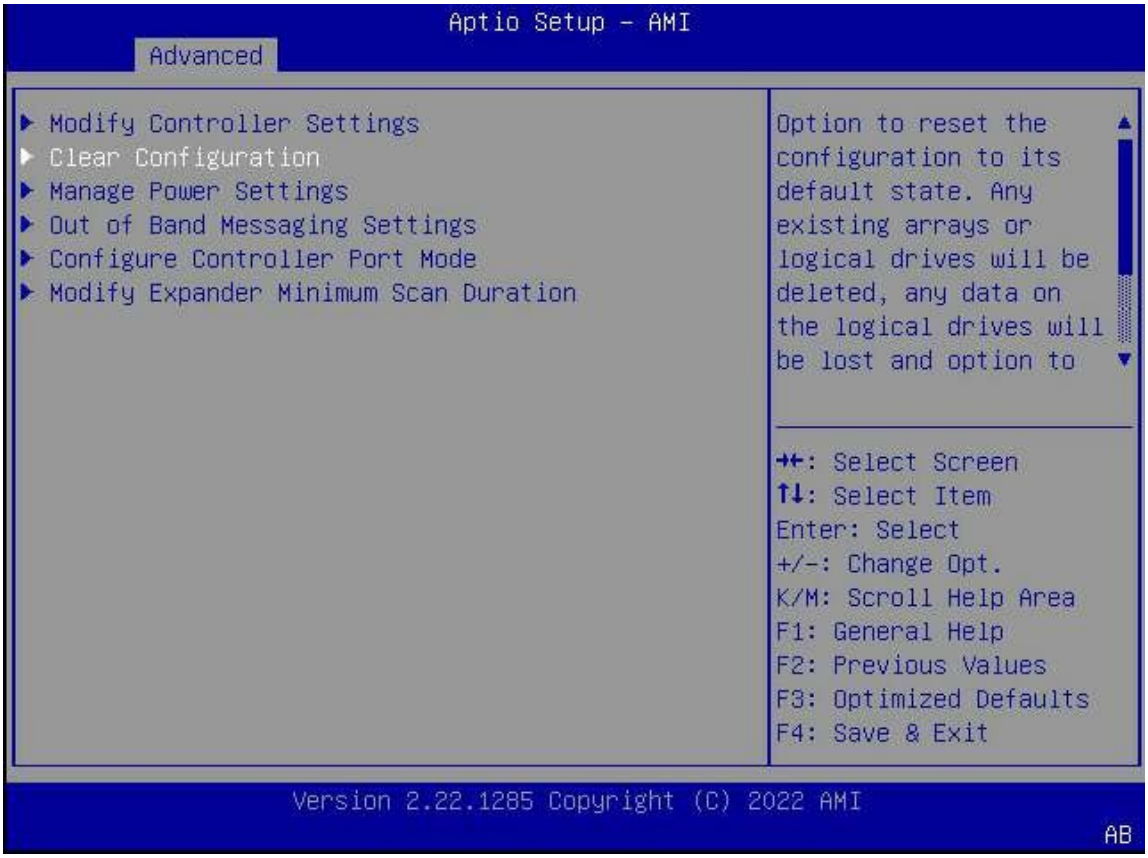

2. Use the arrow keys to select **Clear Configuration**, and then press **Enter**. The **Clear Configuration** screen is displayed, see Figure 3-121.

Figure 3-121 Clear Configuration Screen

3. Use the arrow keys to select **Delete All Array Configuration**, and then press **Enter**. The **Delete All Array Configuration** screen is displayed, see Figure 3-122.

Figure 3-122 Delete All Array Configuration Screen

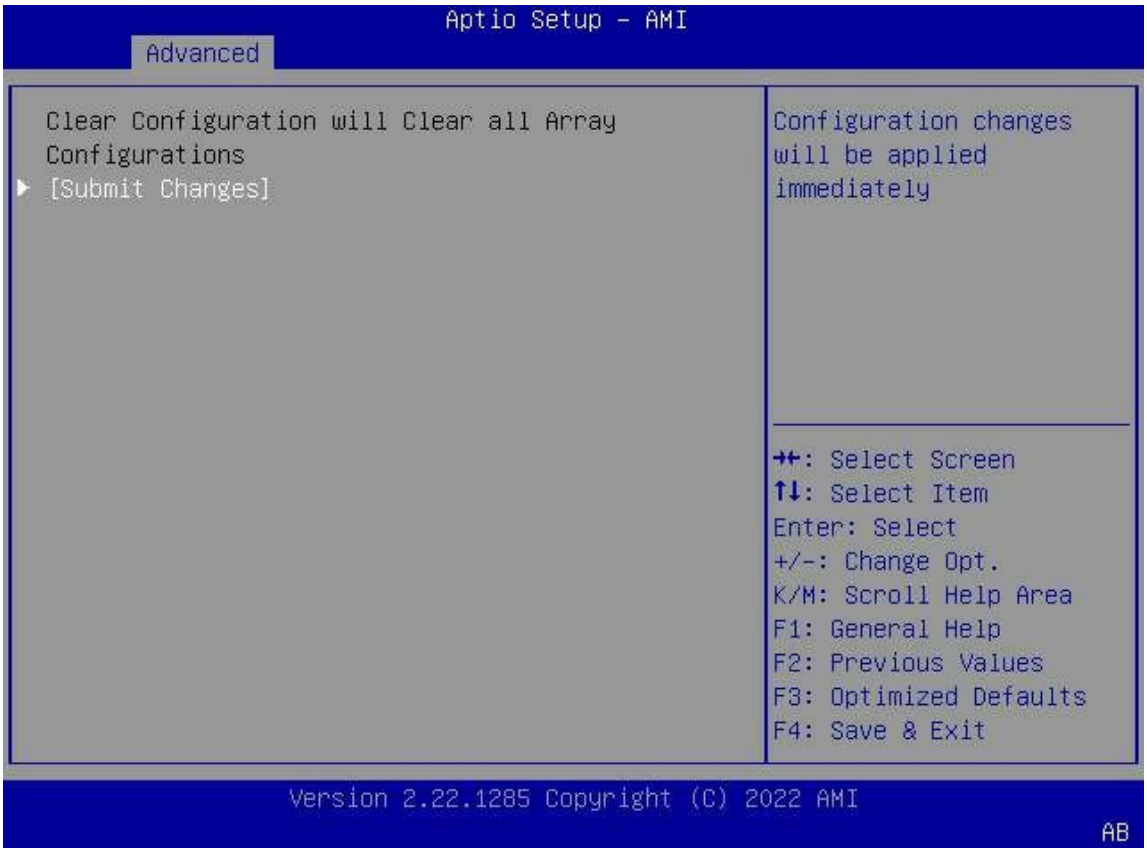

4. Use the arrow keys to select **Submit Changes**, and press **Enter**. The configuration information is cleared successfully, see Figure 3-123.

Figure 3-123 Configuration Cleared Successfully

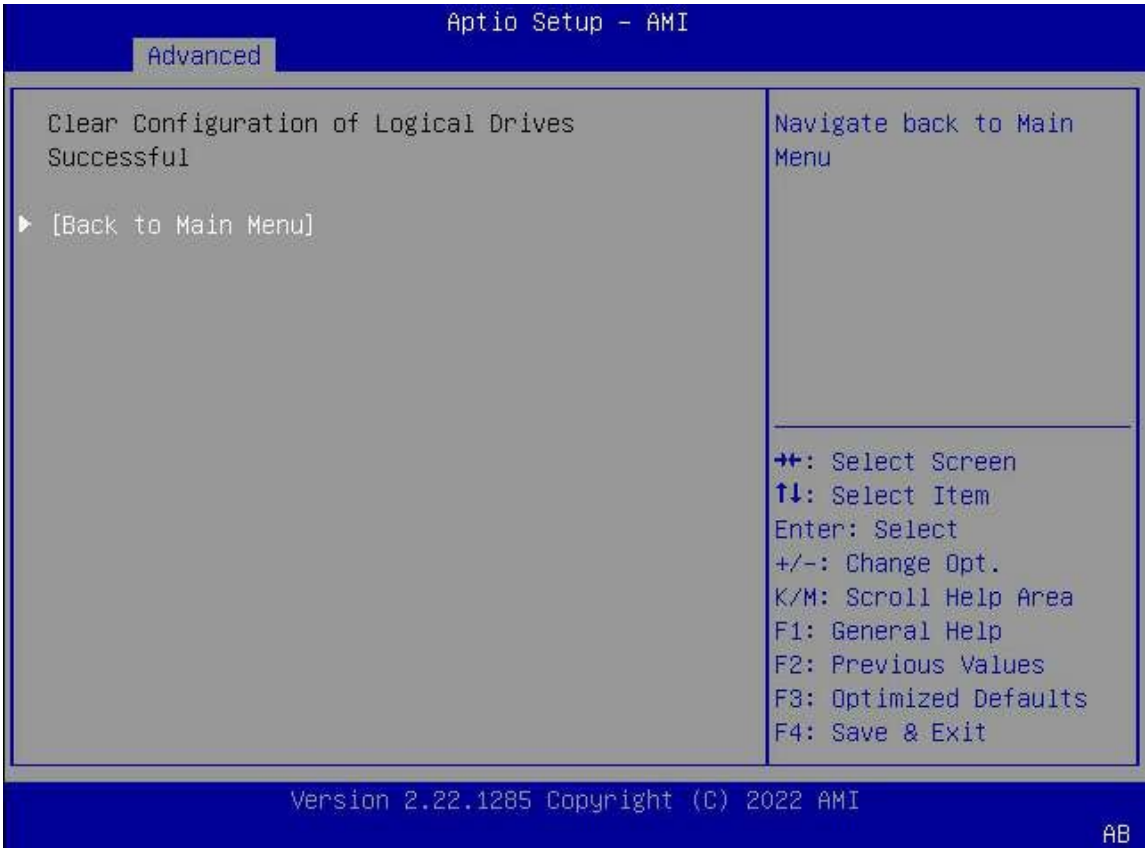

5. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.5.10 Configuring a Pass-Through Disk

Abstract

In accordance with the actual RAID configuration, you can configure a pass-through disk in either of the following ways:

- Setting the mode of all ports in batches:
  If the disks connected to all ports are not used to create a RAID volume, the mode of these ports can be set to HBA in a unified manner.
- Setting the mode of ports separately:
  If the disks connected to some ports have been used to create a RAID volume, the mode of these ports can be set to Mixed and that of other ports can be set to HBA.

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.

- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

- Setting the Mode of All Ports in Batches
  1. On the controller management screen, use the arrow keys to select **Controller Information**, and then press **Enter**. The **Controller Information** screen is displayed. The value of **Controller Mode** is **Mixed**, see Figure 3-124.

Figure 3-124 Controller Information Screen

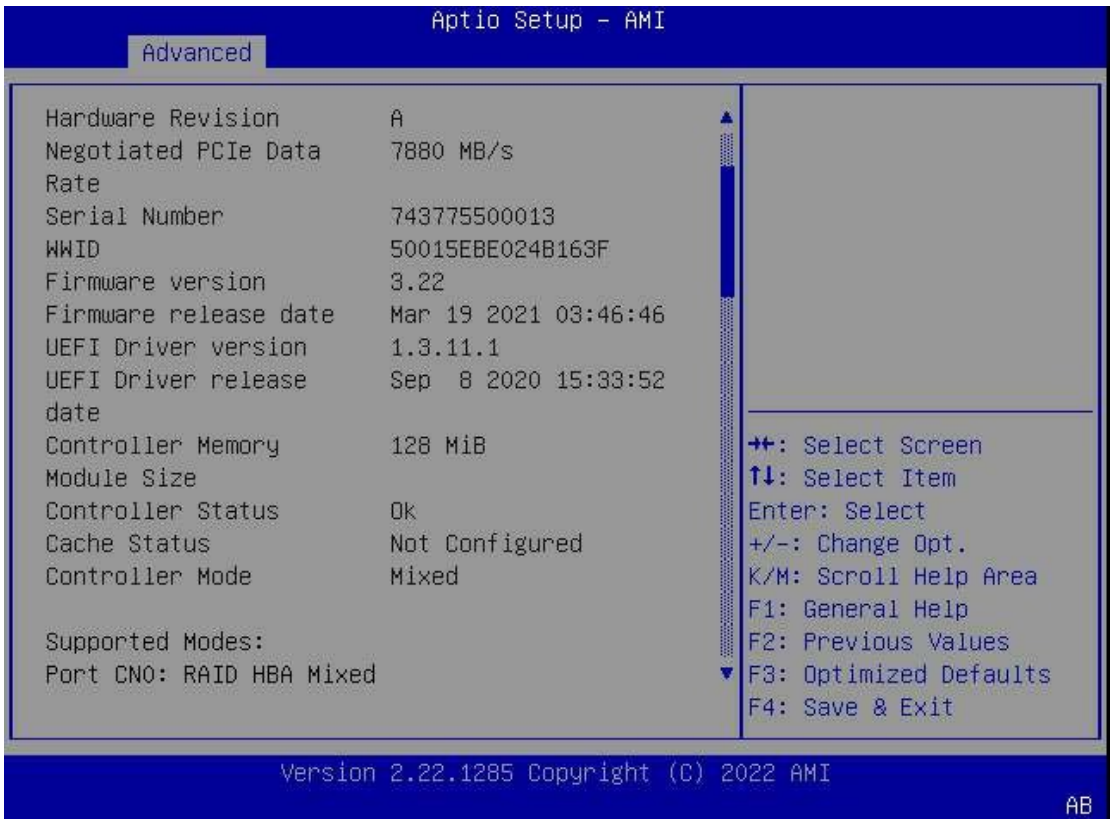

  2. Press **Esc** to return to the controller management screen. Use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-125.

Figure 3-125 Configure Controller Settings Screen

3. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**.
   The **Configure Controller Port Mode** screen is displayed, see Figure 3-126.

Figure 3-126 Configure Controller Port Mode Screen

4. Use the arrow keys to select **Set Controller Port Mode**, and then press **Enter**. The **Set Controller Port Mode** dialog box is displayed, see Figure 3-127.

Figure 3-127 Set Controller Port Mode Dialog Box

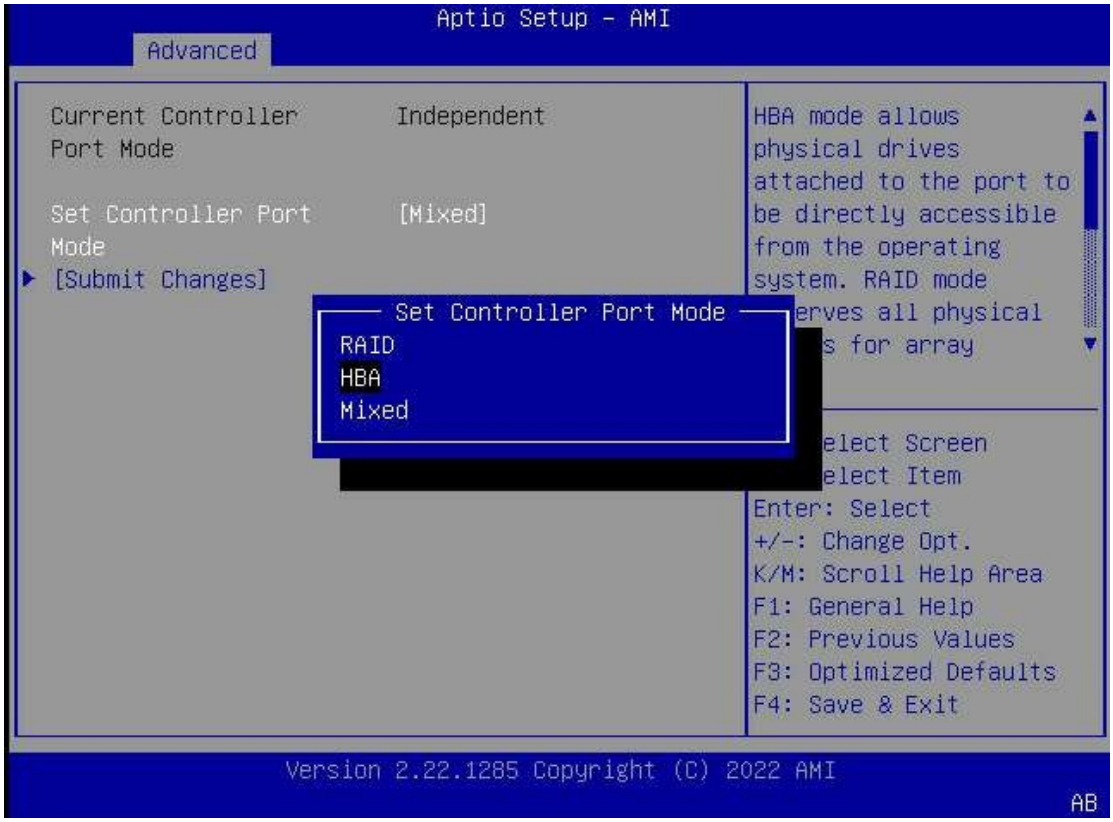

5. Use the arrow keys to select **HBA**, and then press **Enter**. The mode of all ports is set to **HBA**, see Figure 3-128.

Figure 3-128 Setting the Port Mode to HBA

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 3-129.

Figure 3-129 Port Mode Set Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen, see Figure 3-130.

Figure 3-130 Controller Management Screen

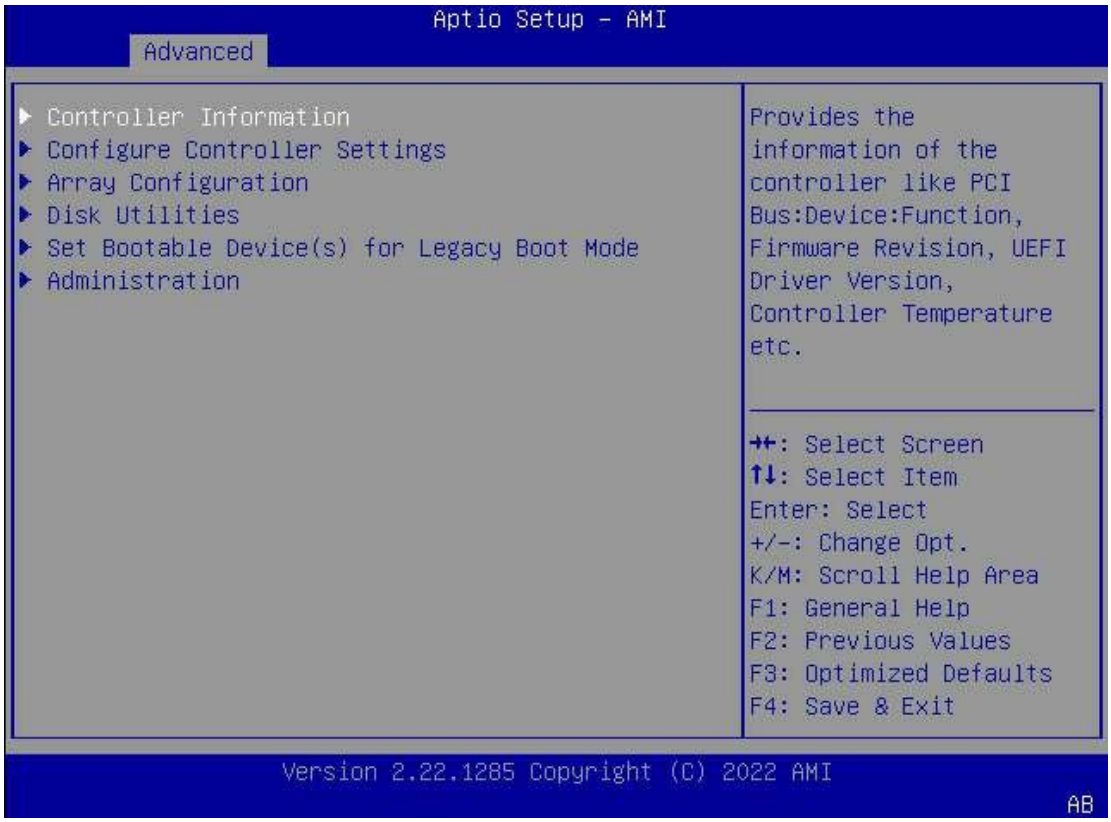

8. Use the arrow keys to select **Controller Details**, and then press **Enter**. On the displayed **Controller Information** screen, view the port mode after the modification, see Figure 3-131. The value of **Controller Mode** is **HBA**.

Figure 3-131 Controller Information Screen

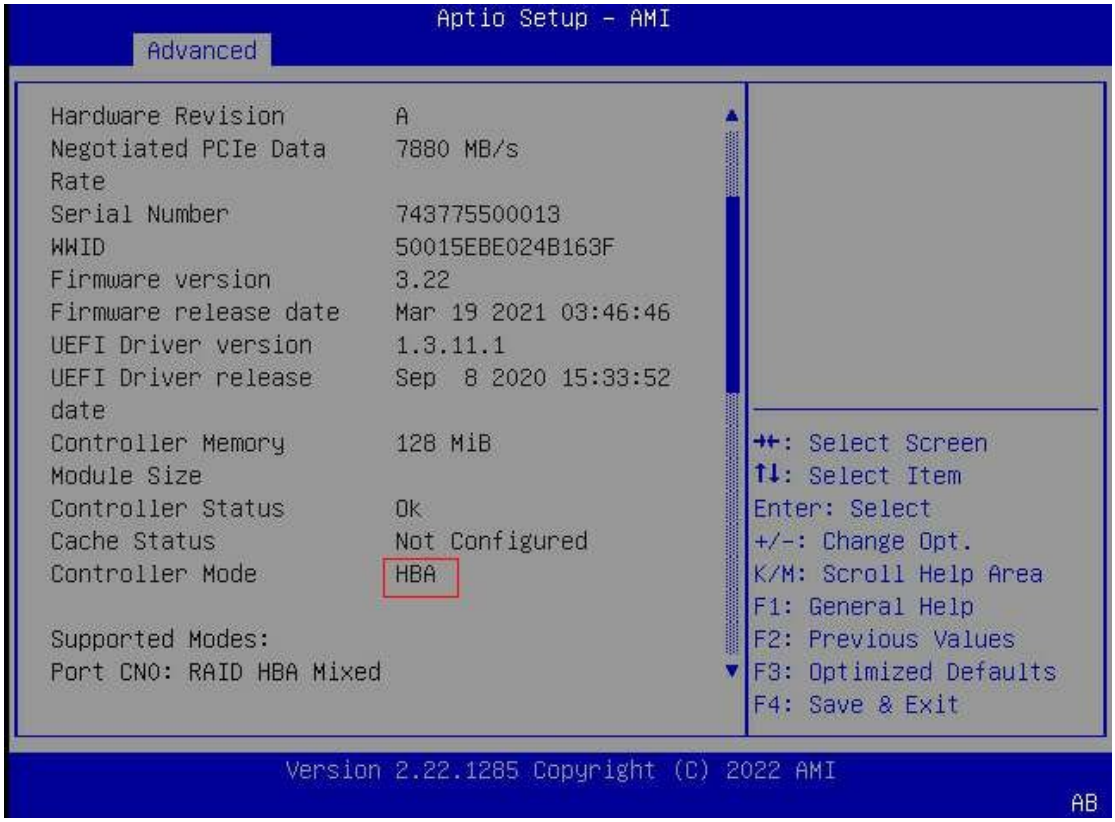

- Setting the Mode of a Port Separately

  1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 3-132.

Figure 3-132 Configure Controller Settings Screen

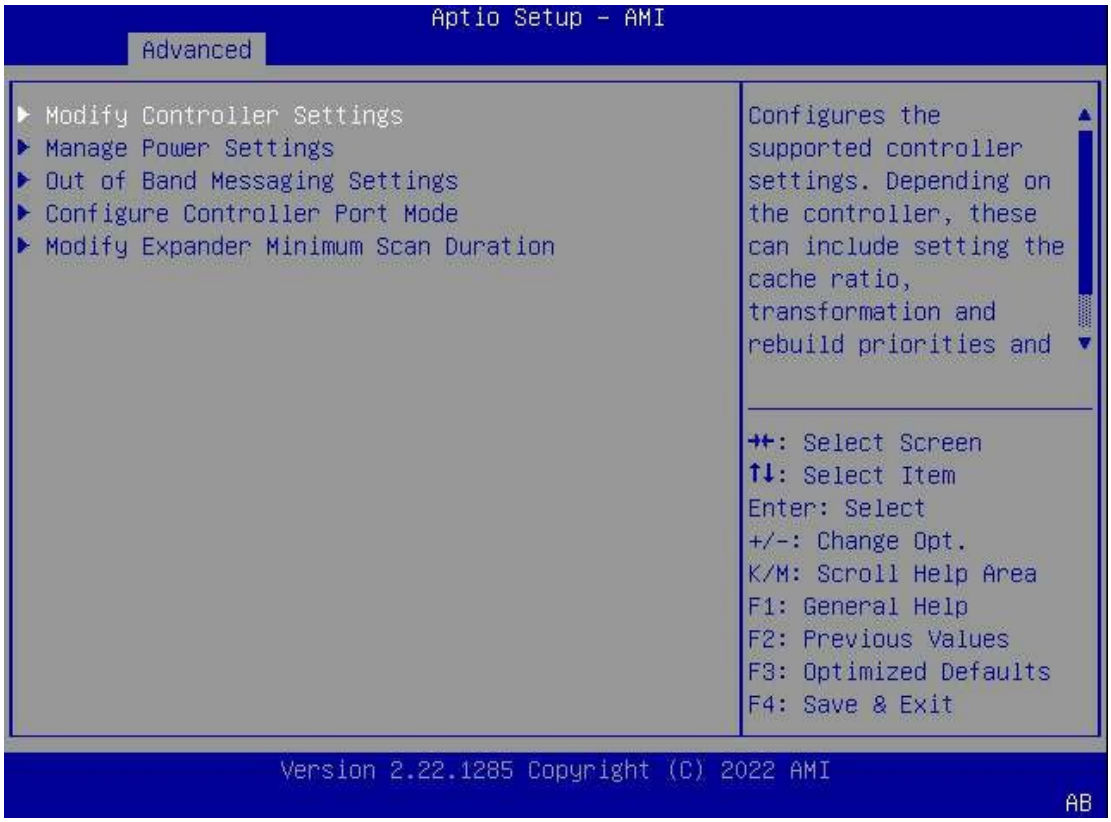

2. Use the arrow keys to select **Modify Controller Settings**, and then press **Enter**. The **Modify Controller Settings** screen is displayed, see Figure 3-133.

Figure 3-133 Modify Controller Settings Screen

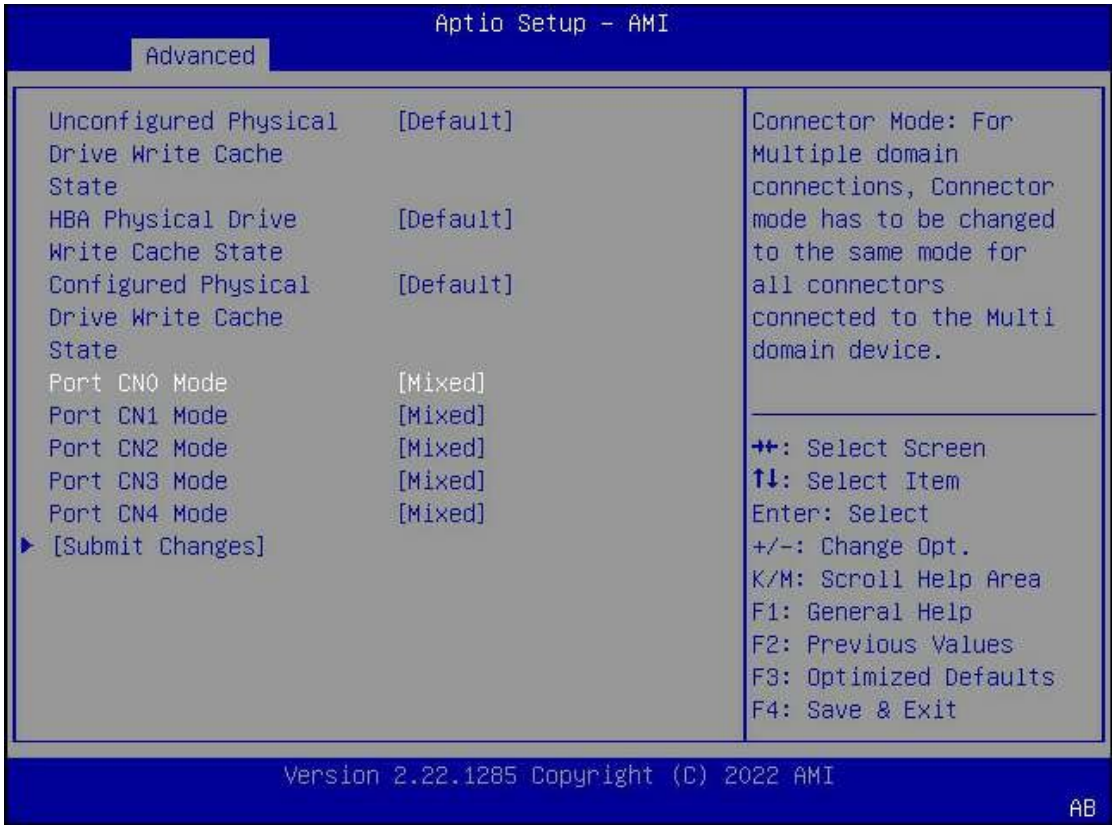

3. Use the arrow keys to select **Port CN0 Mode**, and then press **Enter**. The **Port CN0 Mode** dialog box is displayed, see Figure 3-134.

Figure 3-134 Port CNO Mode Dialog Box

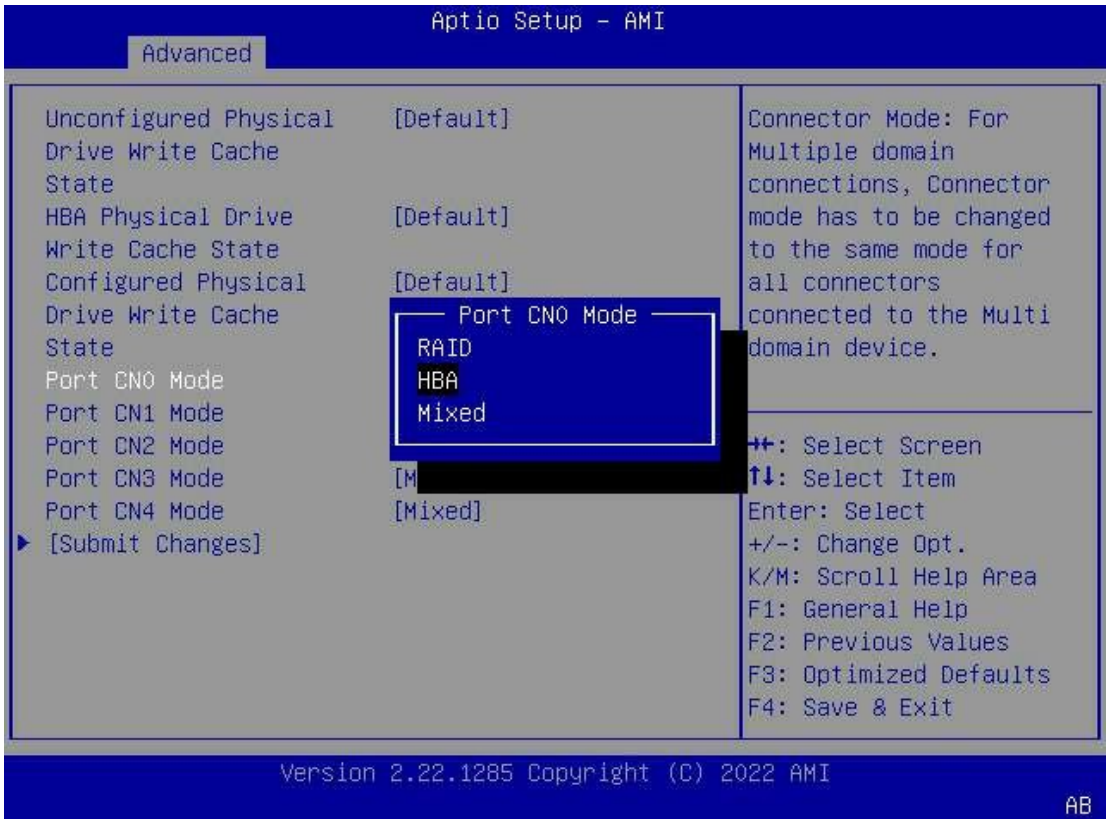

4. Use the arrow keys to select **HBA**, and then press **Enter**. The mode of the port whose connected disk is not used to create a RAID volume is set to **HBA**, see Figure 3-135.

Figure 3-135 Setting the Port Mode to HBA

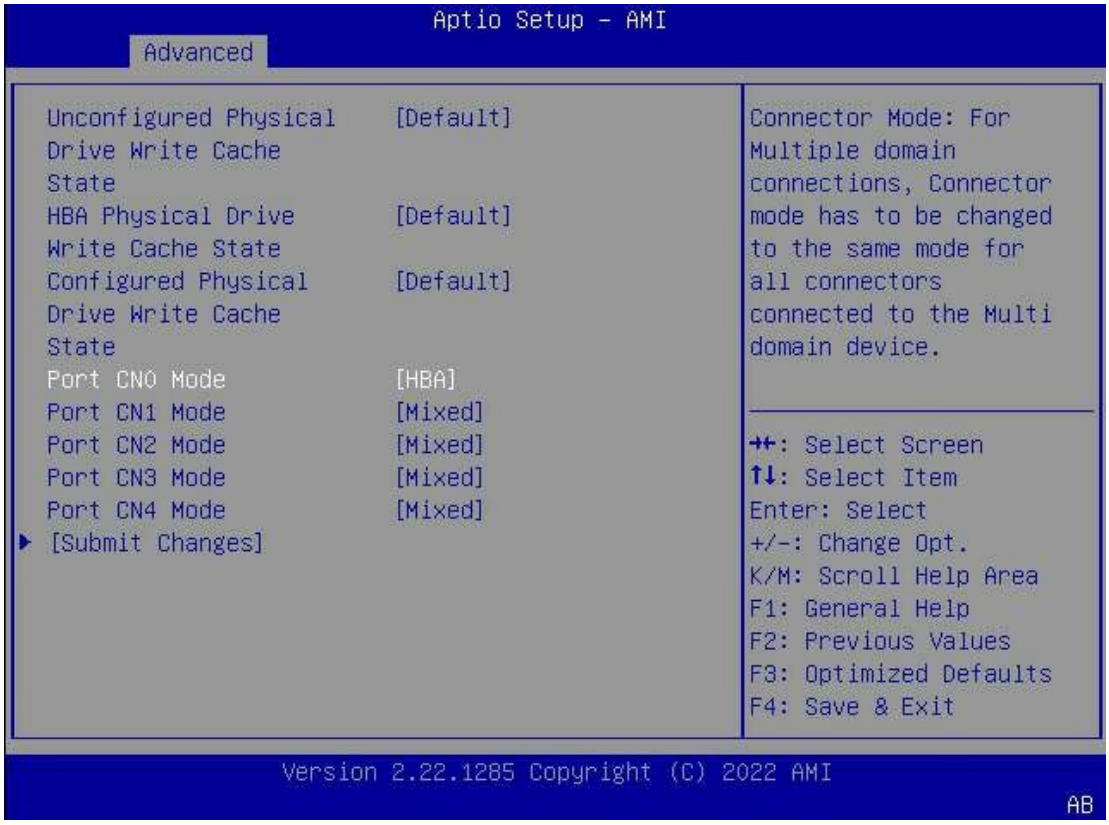

5. Repeat Step 3 and Step 4 to set the mode of another port whose connected disk is not used to create a RAID volume to **HBA**.
6. Use the arrow keys to select a port whose connected disk is already used to create a RAID volume (using **Port CN2 Mode** as an example), and then press **Enter**. The port mode setting menu is displayed, see Figure 3-136.

Figure 3-136 Port Mode Setting Menu

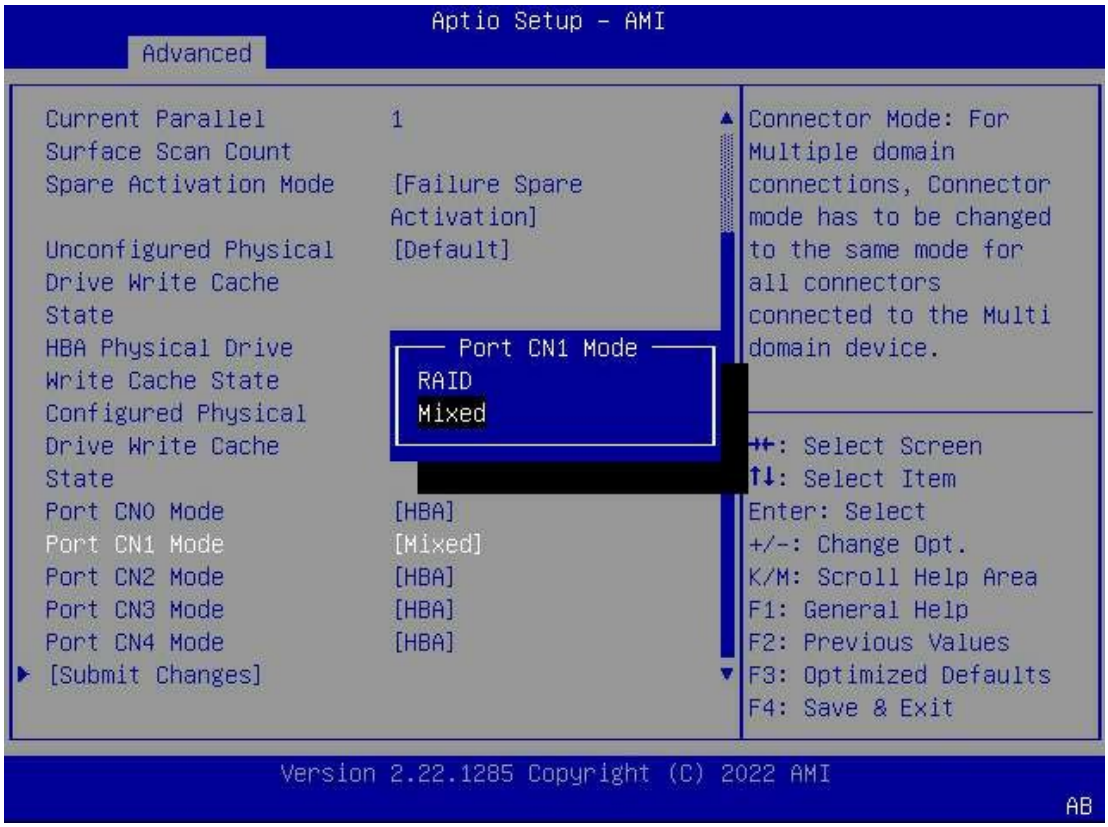

7. Use the arrow keys to select **Mixed**, and then press **Enter**. The mode of the port whose connected disk is used to create a RAID volume is set to **Mixed**, see Figure 3-137.

Figure 3-137 Setting the Port Mode to Mixed

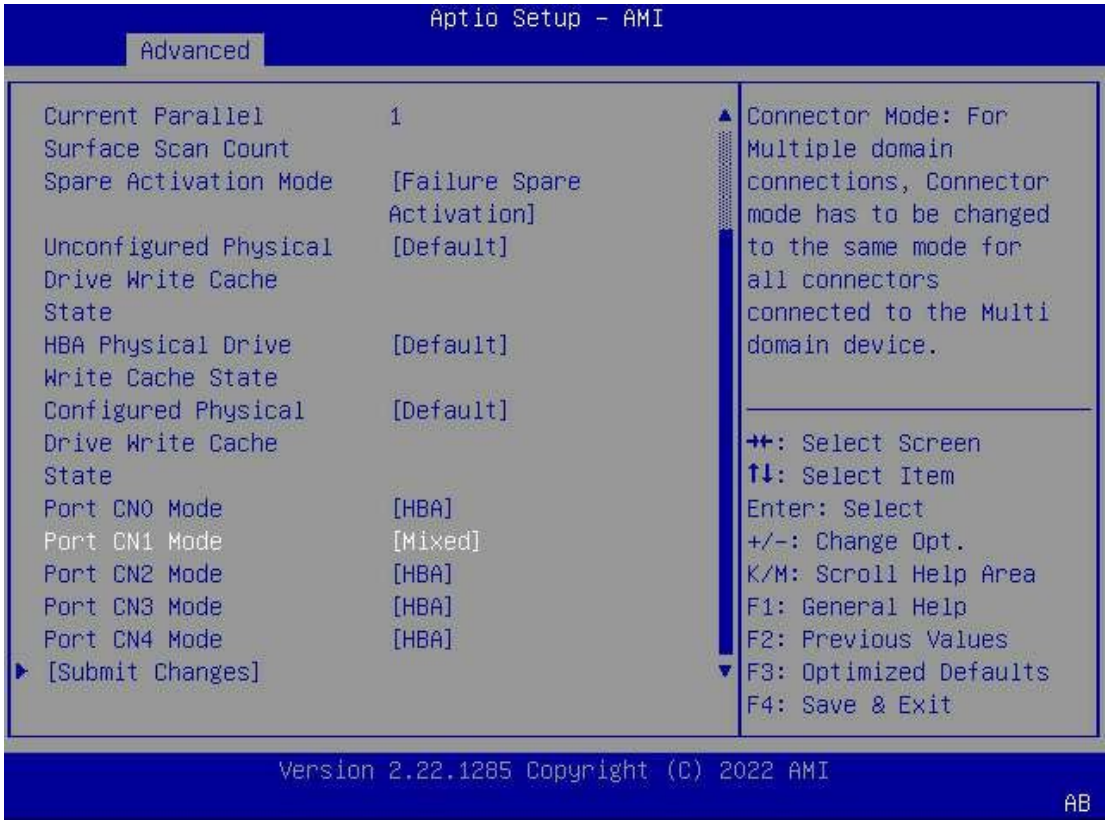

Note

The mode of the ports whose connected disks are already used to create a RAID volume cannot be set to HBA.

8. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 3-138.

Figure 3-138 Port Mode Set Successfully

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

3.6 Typical Scenarios for Replacing a Disk (Legacy Mode)

For a description of the common scenarios for replacing disks in a RAID volume on a SmartIOC 2100 RAID card in legacy mode, refer to Table 3-12.

Table 3-12 Common Scenarios for Replacing a Disk in a RAID Volume on a SmartIOC 2100 RAID Controller Card

Scenario	Description
Scenario 1	A newly inserted disk is converted into a RAID member disk. For details, refer to "3.6.1 Converting a Newly Inserted Disk Into a RAID Member Disk".
Scenario 2	After a faulty SmartROC 2100 RAID card is replaced, all the member disks in the RAID 1 array managed by the original faulty RAID card are moved to a new RAID card. For details, refer to “3.6.2 Moving All Member Disks of a RAID 1 Volume”.

3.6.1 Converting a Newly Inserted Disk Into a RAID Member Disk

Abstract

A newly inserted disk needs to be converted into a RAID member disk in the following two cases:

- The newly inserted disk is a foreign disk.
- The disk in a slot is removed and inserted back.

Inserting a Foreign Disk as a New Disk

When a disk in a RAID volume created on a server is faulty and needs replacement, remove the faulty disk from the disk slot on the server, and insert the prepared disk into the disk slot of the faulty disk.

After the disk is replaced, the configuration utility of the RAID controller card automatically synchronizes data on the newly inserted disk in the RAID volume.

Note

Data is automatically synchronized to the newly inserted disk no matter whether it carries RAID information or not.

Installing a Disk in the Original Slot After Removing It from the Slot

After a disk on a server is used to create a RAID volume, if the disk is removed from its slot and then inserted back, the RAID controller card configuration utility automatically rebuild the disk.

Note

The RAID 0 volume does not support the above functions.

3.6.2 Moving All Member Disks of a RAID 1 Volume

Abstract

If a SmartIOC 2100 RAID controller card on a server fails and needs to be replaced, all the member disks in the RAID 1 volume on the faulty RAID controller card need to be moved to a new SmartIOC 2100 RAID controller card.

Notice

It is risky to move the member disks of the RAID volume, and therefore it is recommended that you contact VANTAGEO technical support for help.

Steps

1. Shut down the server, and replace the faulty SmartIOC 2100 RAID controller card with a new one.
2. Connect all member disks of the RAID 1 volume be moved to the new SmartIOC 2100 RAID controller card.
3. Power on the server again and start the server system.
4. Start the BIOS configuration utility. For details, refer to “3.2.1 Starting the Configuration Utility”.
5. Query RAID volume information. For details, refer to “3.4.1 Querying RAID Volume Information”.
6. Contact VANTAGEO technical support to move member disks.

3.7 Typical Scenarios for Replacing a Disk (UEFI Mode)

For a description of the common scenarios for replacing a disk in a RAID volume on a SmartIOC 2100 RAID card in UEFI mode, refer to Table 3-13.

Table 3-13 Common Scenarios for Replacing a Disk in a RAID Volume on a SmartIOC 2100 RAID Controller Card

Scenario	Description
Scenario 1	When a RAID 0 member disk is faulty, the RAID controller card is reconfigured. For details, refer to “3.7.1 A RAID 0 Member Disk Fails”.
Scenario 2	When a member disk of a logical volume with no hot spare disk configured is faulty, the faulty disk is replaced. For details, refer to "3.7.2 A Member Disk of a RAID Redundant Logical Volume (Without a Configured Hot Spare Disk) Fails".
Scenario 3	When a member disk of a logical volume with a hot spare disk configured is faulty, the faulty disk is replaced. For details, refer to "3.7.3 A Member Disk of a RAID Redundant Logical Volume (with a Configured Hot Spare Disk) Fails".

3.7.1 A RAID 0 Member Disk Fails

RAID 0 does not support data redundancy or backup. As a result, data cannot be restored after a fault occurs in the RAID 0 logical volume. It is necessary to install a new disk and reconfigure the RAID array.

3.7.2 A Member Disk of a RAID Redundant Logical Volume (Without a Configured Hot Spare Disk) Fails

Abstract

If a fault occurs in a member disk of a redundant logical volume (with no hot spare disk configured) on a SmartIOC 2100 RAID controller card, the SmartIOC 2100 RAID controller card can automatically restore the data after the faulty disk is replaced with a new disk. During the process, the member disk status may be **OK**, but the logical disk status may be **Failed**. In this case, you need to restore the logical disk status.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The screen for configuring an array is displayed.
2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The screen for managing arrays is displayed.
3. Use the arrow keys to select the array for which you need to manage the logical volume, and then press **Enter**. The screen for managing the selected array is displayed.
4. Use the arrow keys to select **List Logical Drives**, and then press **Enter**. The screen for managing logical volumes is displayed.
5. Use the arrow keys to select the logical volume to be corrected, and then press **Enter**. The screen for managing the selected logical volume is displayed.
6. Use the arrow keys to select **Re-Enable Logical Drive**, and then press **Enter**. The screen for restoring logical volume status is displayed.
7. Press **Enter**. The status of the logical disk is restored.

Verification

On the logical volume management screen, use arrow keys to select **Logical Drive Details**, and then press **Enter**. The logical volume details are displayed. Verify that the logical volume status is **Ok**.

3.7.3 A Member Disk of a RAID Redundant Logical Volume (with a Configured Hot Spare Disk) Fails

When a fault occurs in a member disk of a redundant logical volume (with a hot spare disk configured) on a SmartIOC 2100 RAID controller card, the RAID controller card automatically replaces the faulty disk with the hot spare disk and restores the data.

- When the hot spare disk is of the **Dedicated** type, the RAID controller card temporarily replaces the faulty disk with the hot spare disk and automatically restores the data. After the faulty disk is replaced with a new disk, the hot spare disk is restored to Hot Spare status.

- When the hot spare disk is of the **Auto Replace** type, the RAID controller card immediately replaces the faulty disk with the hot spare disk and automatically restores the data. After the new disk is inserted, the new disk becomes a hot spare disk.

Chapter 4
VT SmartROC 3100 RAID Controller Card

Table of Contents

A VT SmartROC 3100 RAID controller card is integrated with the BIOS configuration utility. By using the program, you can configure the disks that are supported by the VT SmartROC 3100 RAID controller card to a RAID volume of a specific level.
A VT SmartROC 3100 RAID controller card supports the caching function, which can improve read/write performance.

Note

The operation screens of the BIOS configuration utility displayed in this chapter are for reference only, and may not be the same as the actual ones.

4.1 Capability Features

For a description of the capabilities of a VT SmartROC 3100 RAID controller card, refer to Table 4-1.

Table 4-1 Descriptions of the Capabilities of a VT SmartROC 3100 RAID Controller Card

Capability Item	Capability Parameter
Product form	Mezz card
Controller chip	PMC PM8236

Capability Item	Capability Parameter
Host interface	PCIe 3.0x8
SAS interface	12 Gb SAS
Number of ports	16+2
Drive interface	SAS and SATA
Drive type	HDD and SSD
Whether drives are hot swappable	Supported
Maximum number of RAID groups	64
Number of drives	238
RAID level	RAID 0, RAID 1, RAID 10, RAID 5, RAID 50, RAID 6, RAID 60
JBOD mode	Supported
Cache	2 GB/4 GB
Cache protection	Super capacitor
Out-of-band management	Supported
Consistency check/verification and fix	Supported
Online capacity expansion	Supported
Online RAID level migration	Supported
Automatic rebuild	Supported
Manufacturer tool support	arcconf

4.2 Initial Configuration (Legacy Mode)

Figure 4-1 shows the initial configuration flow of a VT SmartROC 3100 RAID controller card.

Figure 4-1 Initial Configuration Flow of a VT SmartROC 3100 RAID Controller Card

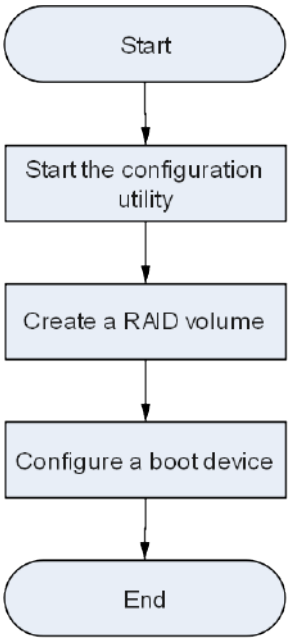

4.2.1 Starting the Configuration Utility

Abstract

This procedure describes how to start the BIOS configuration utility of a VT SmartROC 3100 RAID controller card to log in to the management screen and complete the subsequent initial and common configurations.

Prerequisite

The boot mode is already set to **Legacy** in BIOS. For details, refer to "5.3.1 Setting the Boot Mode to Legacy".

Steps

1. Start the server system.
2. During the POST process, press **Ctrl+A** to start the BIOS configuration utility of the VT SmartROC 3100 RAID controller card. The screen as shown in Figure 4-2 is displayed.

Figure 4-2 BIOS Configuration Utility Screen

3. In the **Options** area, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Configuration Menu** screen is displayed, see Figure 4-3.

Figure 4-3 Configuration Menu Screen

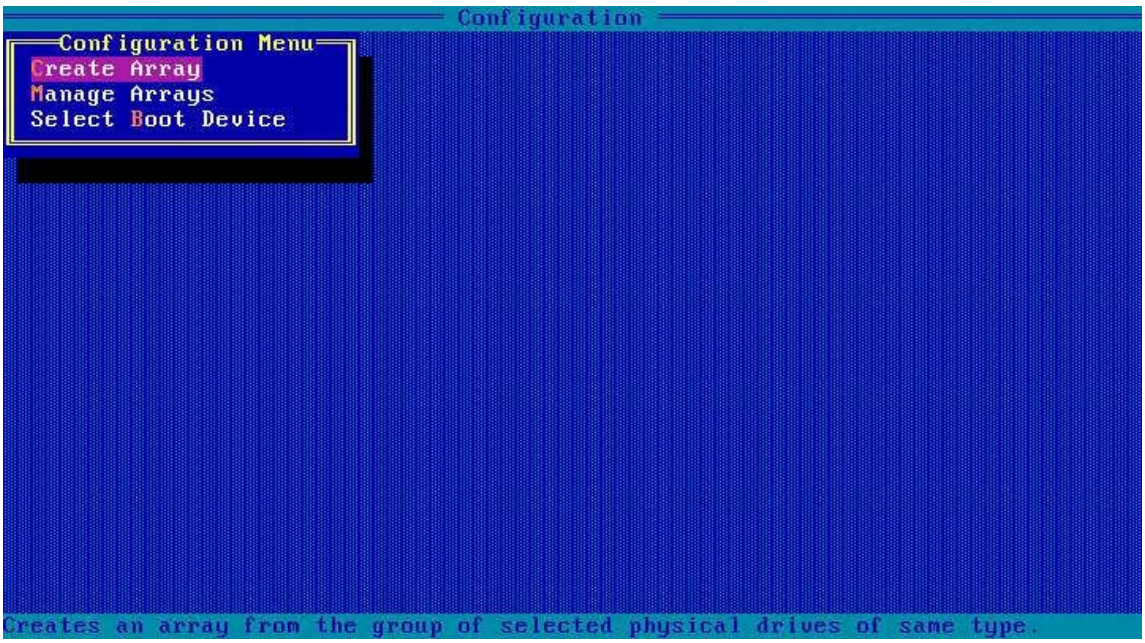

4.2.2 Creating a RAID Volume

Abstract

You can create RAID volumes at different levels as required.

The procedures for creating RAID volumes at different levels are similar. This procedure uses a RAID 0 volume in legacy mode as an example.

Prerequisite

Sufficient SATA and SAS disks are installed on the server.

Context

For a description of the number of disks required to create a RAID volume, refer to Table 4-2.

Table 4-2 Number of Disks Required for Creating a RAID Volume

RAID Level	Description
RAID 0	RAID 0 requires at least one disks.
RAID1	RAID 1 requires at least two disks. Disks with different capacities can be used in a RAID 1 volume, but the logical capacity of each member disk depends on the space of the disk with the smallest capacity.
RAID 5	RAID 5 requires at least three disks.
RAID 6	RAID 6 requires at least four disks.
RAID 10	RAID 10 requires at least four disks. A RAID 10 volume consists of at least two RAID 1 volumes. For example, if there are four disks to be used in RAID 10 mode, you need to add them to two drive groups, each of which is mounted with two disks in RAID 1 mode.
RAID 50	RAID 50 requires at least six disks. A RAID 50 volume consists of at least two RAID 5 volumes. For example, if there are six disks to be used in RAID 50 mode, you need to add them to two drive groups, each of which is mounted with three disks in RAID 5 mode.
RAID 60	RAID 60 requires at least eight disks. A RAID 60 volume consists of at least two RAID 6 volumes. For example, if there are eight disks to be used in RAID 60 mode, you need to add them to two drive groups, each of which is mounted with four disks in RAID 6 mode.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Create Array**, and then press **Enter**. On the displayed **Select drives to create Array** screen, all the disks that can be used to create a RAID volume are displayed, see Figure 4-4.

Figure 4-4 Select Drives to Create Array Screen

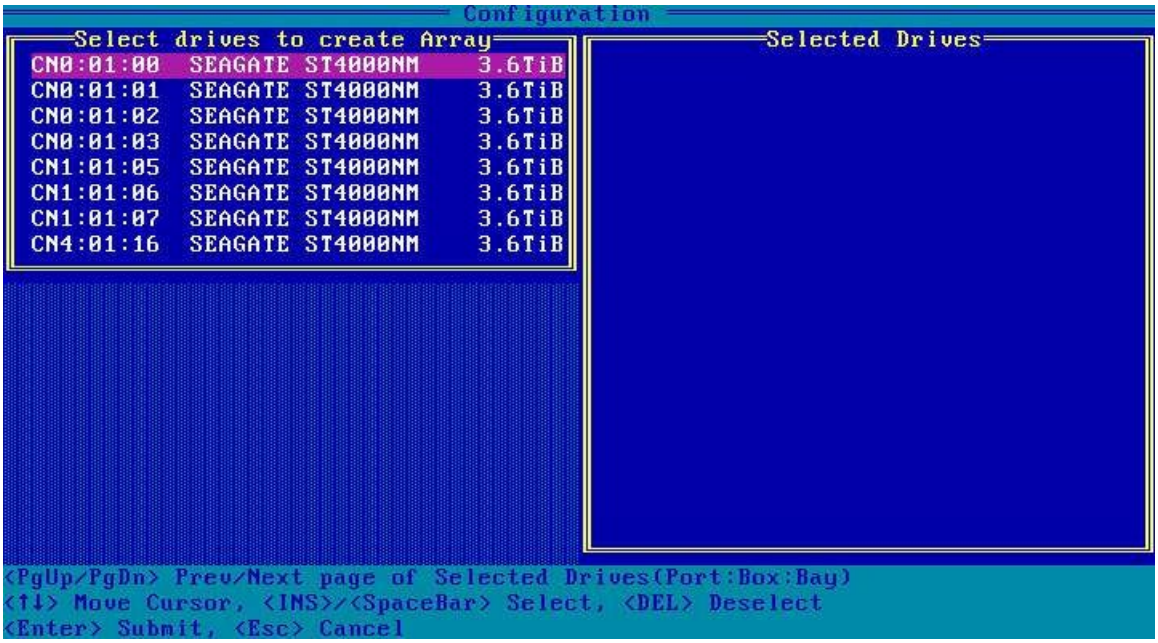

2. Use the arrow keys to select the disks to be used to create the RAID volume, and then press **Insert** to add these disks to the **Selected Drives** list, see Figure 4-5.

Figure 4-5 Selected Drives List

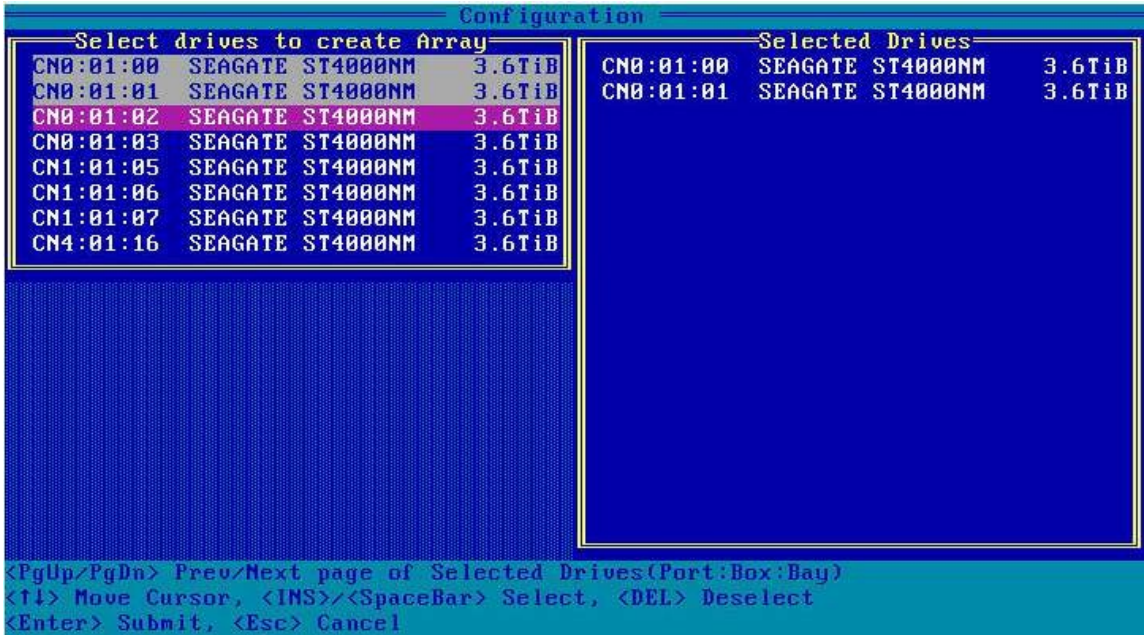

Note

- The disks for creating a RAID volume must be of the same type. It is forbidden to select disks with interface types such as SATA and SAS at the same time.
- Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

3. Press **Enter**. The **Create Logical Drive** screen is displayed, see Figure 4-6.

Figure 4-6 Create Logical Drive Screen

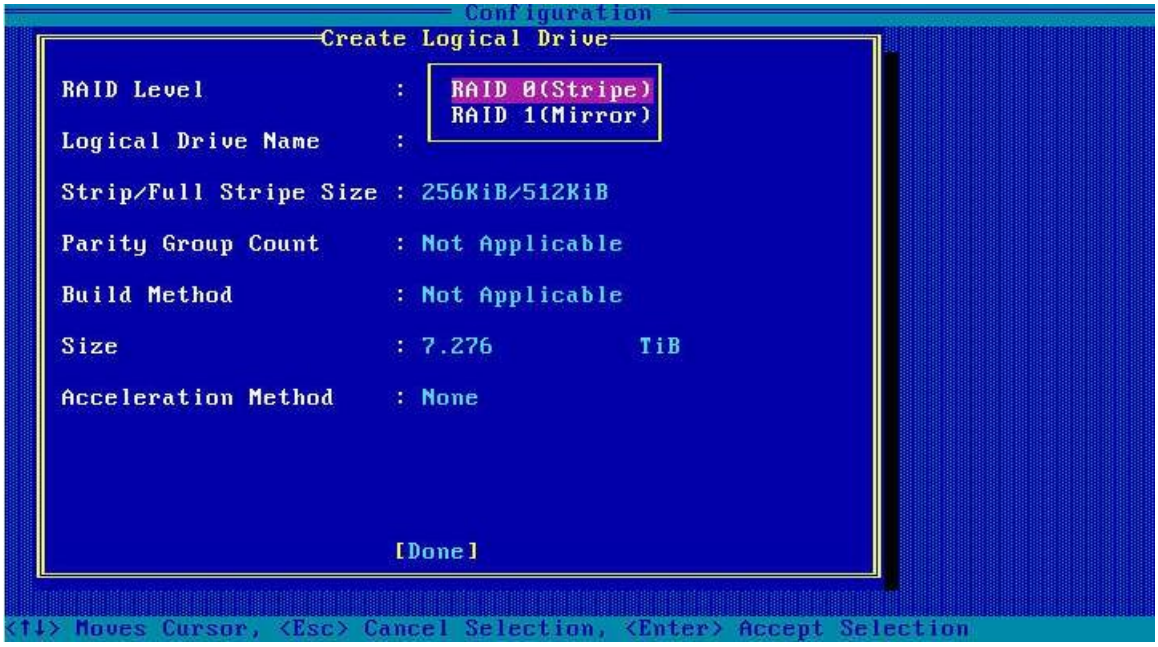

4. Use **Tab**/**Tab+Shift** to select the parameters that you want to modify. In the displayed operation box, use the arrow keys to select the related parameters, and then press **Enter** for confirmation. For a description of the parameters on the **Create Logical Drive** screen, refer to Table 4-3.

Table 4-3 Descriptions of the Parameters on the Create Logical Drive Screen

Parameter	Description
RAID Level	Sets a RAID level, for example, **RAID 0(Stripe)**.
Logical Drive Name	Sets the RAID name, for example, "0".
Strip/Full Stripe Size	The stripe size should be equal to the size of average disk IO requests generated by server applications. In the optimum status, only one IO operation is executed for each IO request. The recommended stripe size configurations are as follows: ● For a Web server, 8 KB is recommended. ● For a groupware server (such as an email server), 16 KB is recommended. ● For a database server, 16 KB or 32 KB is recommended. ● For a file server, 32 KB or 64 KB is recommended ● For a video file server, 64 KB, 128 KB, or 256 KB is recommended.
Parity Group Count	Configures logical-device parity groups in accordance with the number of physical devices in the array. It is not applicable to all RAID levels.

Parameter	Description
Build Method	Sets the RAID initialization method, which is used to determine how the logical devices prepare for read and write, and how long the initialization takes. ● default: When the logical devices can be accessed by the operating system, parity blocks are initialized at the back end. A lower RAID level can achieve faster parity initialization. ● RPI: The data and parity blocks at the front end are overwritten. Before the parity initialization procedure is completed, logical devices remain invisible and unavailable to the operating system. All parity groups are initialized in parallel, but the initialization of a single parity group (RAID 5) is faster. The RAID level does not affect the performance during the RAID initialization. ● Not Applicable: unavailable. Keep the default configuration unless otherwise specified.
Size	Displays the storage size of the disk array in accordance with the total storage space of the disks added to the disk array. By default, the RAID created uses all the available disk space.
Acceleration Method	Sets the caching mode for the RAID volume. ● IO Bypass: This option is valid only when the RAID logical volume is formed by SSDs. ● Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time. ● None: disables the controller cache. Neither **IO Bypass** nor **Controller Cache** is used. Keep the default configuration unless otherwise specified.

5. Use **Tab** to select **Done**, and then press **Enter** to create the RAID volume, see Figure 4-7.

Figure 4-7 Creating a RAID Volume

6. The RAID volume is created successfully, see Figure 4-8. Press any key to return to the **Configuration Menu** screen.

Figure 4-8 RAID Volume Created Successfully

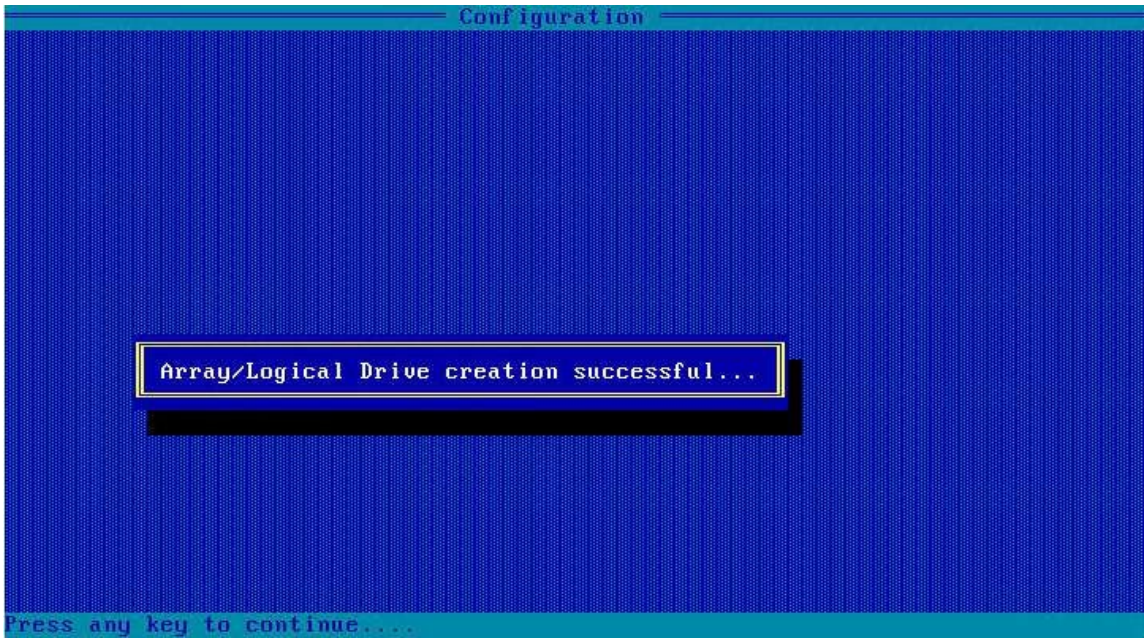

4.2.3 Configuring a Boot Device

Abstract

After a RAID volume is created, if you need to install an operating system on the RAID volume, and there are multiple RAID volumes on the RAID controller card, you must set the RAID volume as a boot device.

Prerequisite

A RAID volume is created successfully. For details, refer to “4.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** page is displayed, see Figure 4-9.

Figure 4-9 List of Arrays Screen

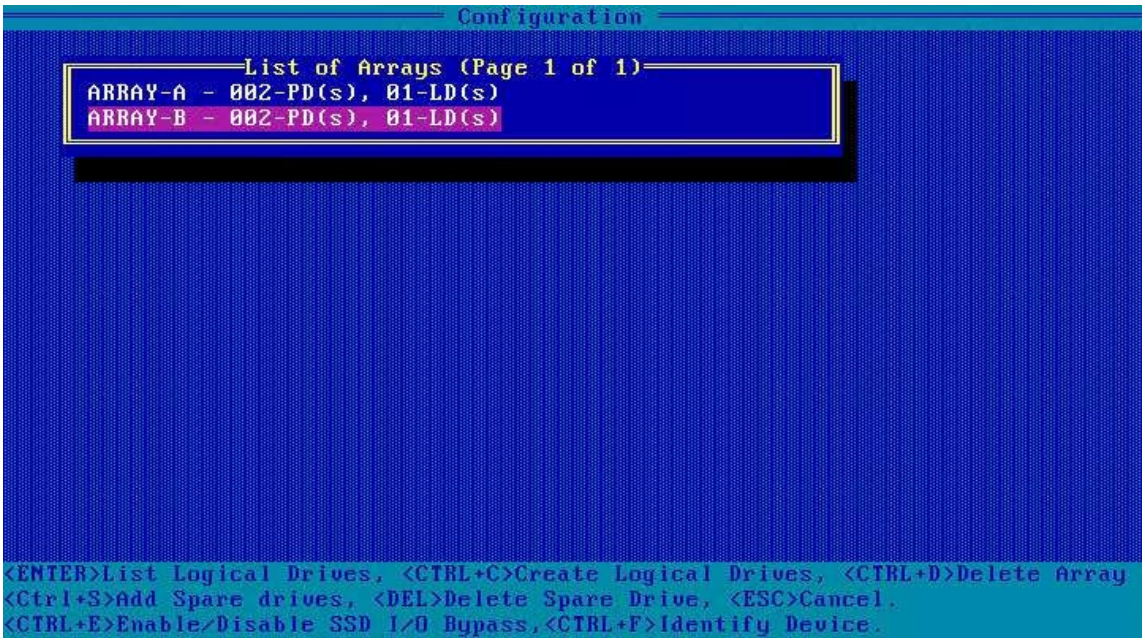

2. Use the arrow keys to select the array where the RAID volume to be set as a boot device is located, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 4-10.

Figure 4-10 List of Logical Drives Screen

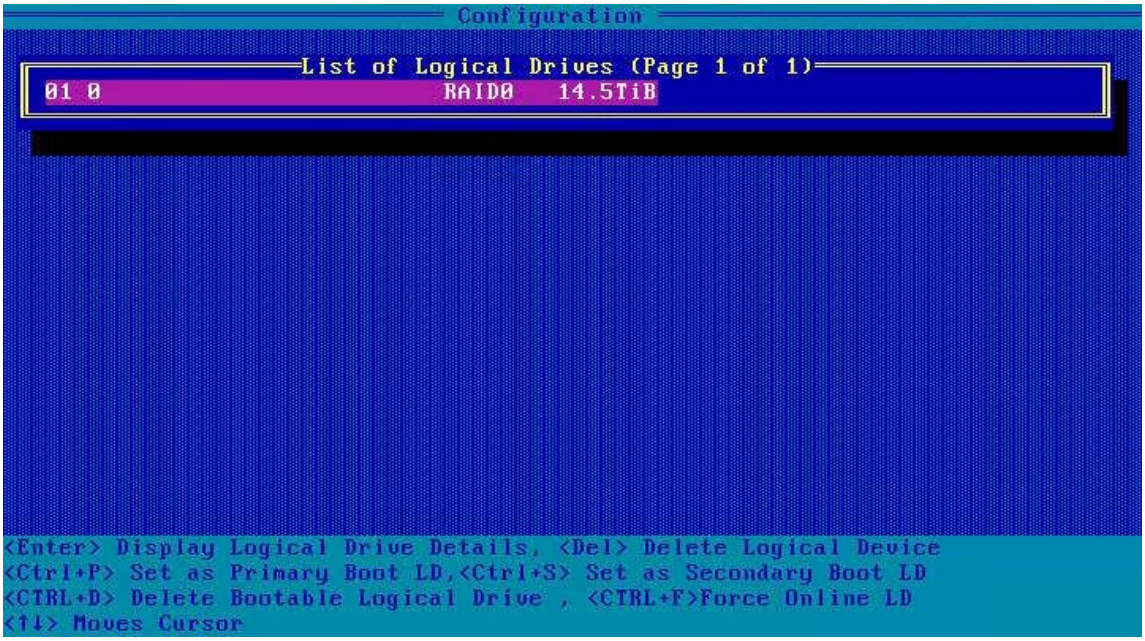

3. Select the RAID volume to be booted first, and then press **Ctrl+P** to configure it as the first boot device, see Figure 4-11.

Figure 4-11 Configuring a Boot Device

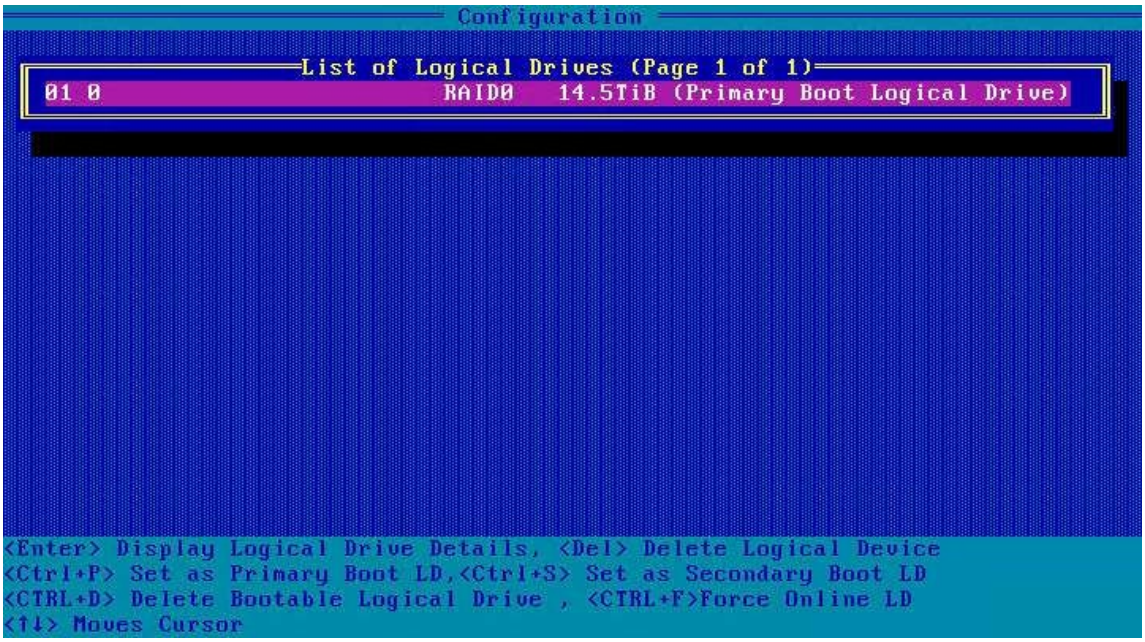

4.3 Initial Configuration (UEFI Mode)

Figure 4-12 shows the initial configuration flow of a VT SmartROC 3100 RAID controller card.

Figure 4-12 Initial Configuration Flow of a VT SmartROC 3100 RAID Controller Card

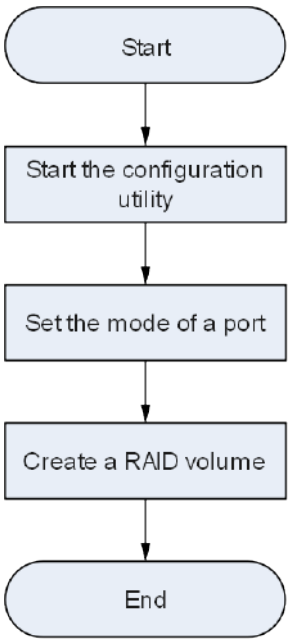

Note

- The SmartROC 3100 RAID controller card is in UEFI mode, and it currently does not support the boot disk settings.
- The SmartROC 3100 RAID controller card does not support the co-existence of configurations in both UEFI and legacy modes. If the mode is switched from UEFI to legacy, the configuration in UEFI mode must be cleared. Otherwise, the normal operation of the RAID controller card is affected.

4.3.1 Starting the Configuration Utility

Abstract

This procedure describes how to start the BIOS configuration utility of a VT SmartROC 3100 RAID controller card to log in to the management screen and complete the subsequent initial and common configurations.

Prerequisite

The boot mode is already set to UEFI in BIOS. For details, refer to "5.3.2 Setting the Boot Mode to UEFI".

Steps

1. Start the server system.
2. During the POST process, press **F2**/**DEL**. The **Aptio Setup** screen is displayed, see Figure 4-13.

Figure 4-13 Aptio Setup Screen

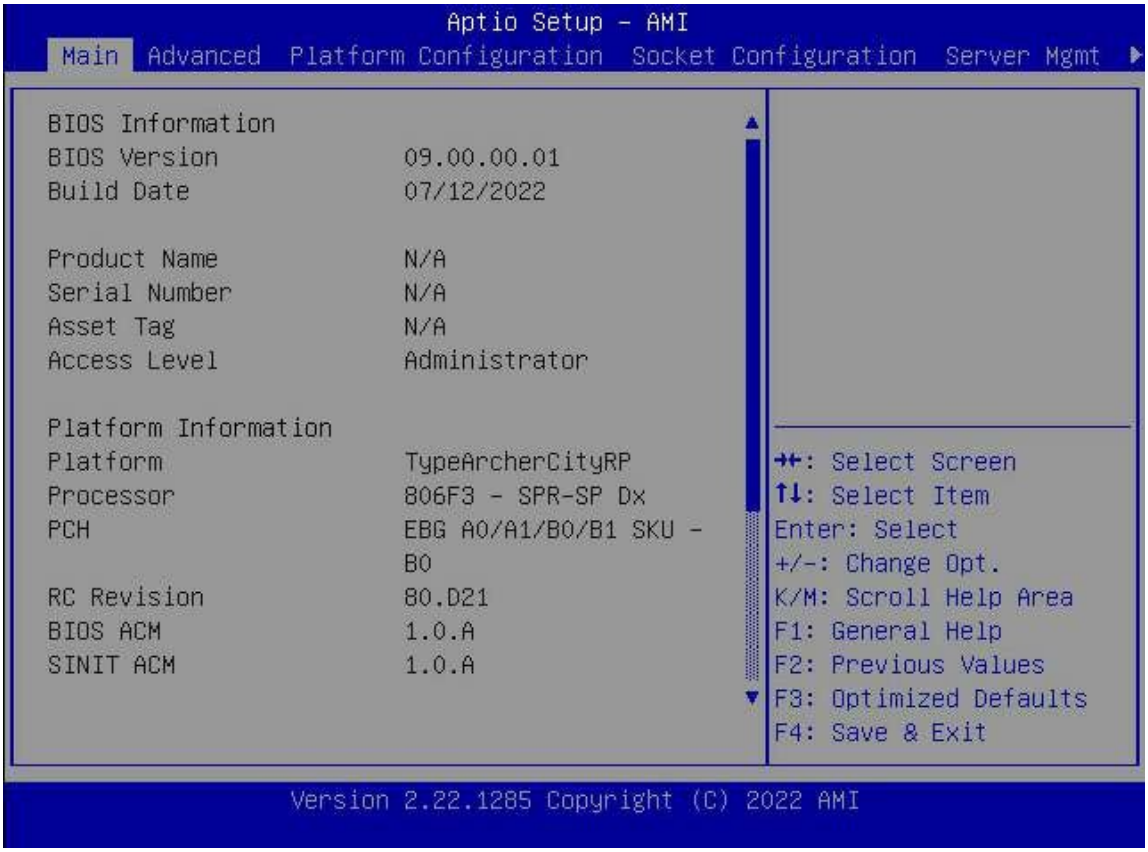

3. Use the arrow keys to select **Advanced**, and then press **Enter**. The **Advanced** screen is displayed, see Figure 4-14.

Figure 4-14 Advanced Screen

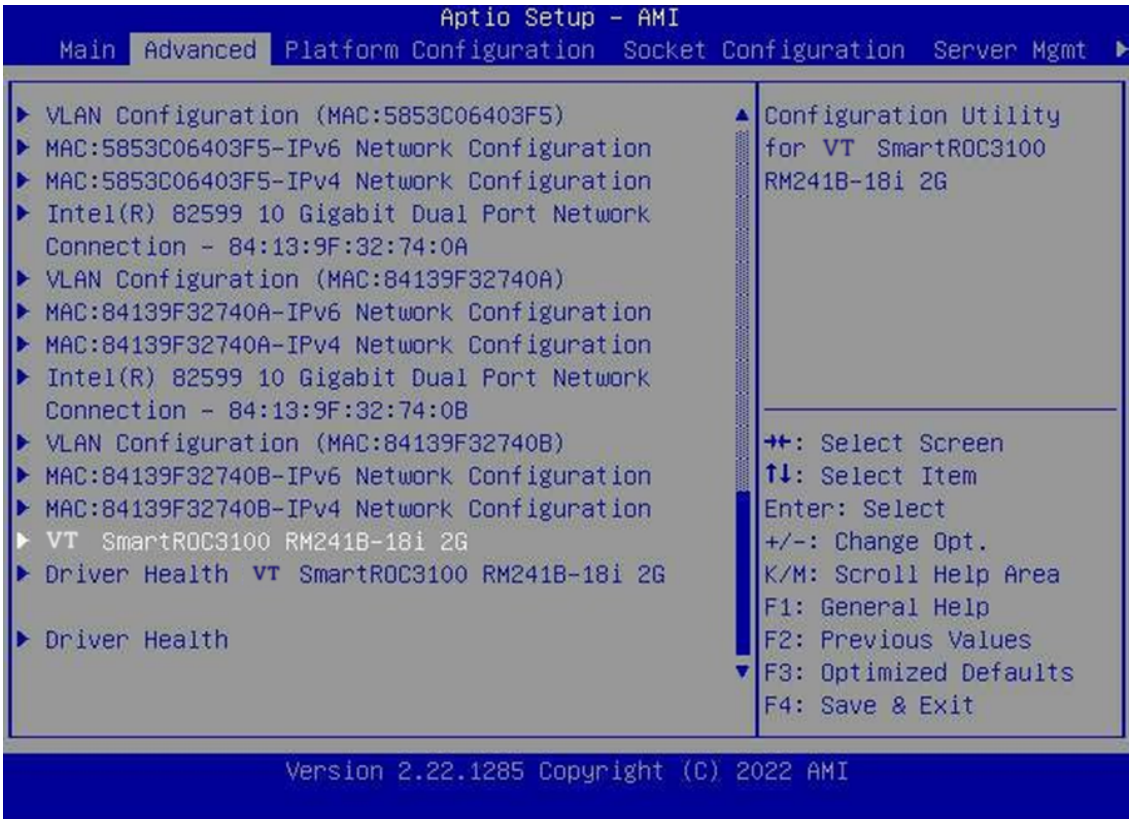

4. Use the arrow keys to select **VT SmartROC3100 RM242-18i 4G**, and press **Enter**. The controller management screen is displayed, see Figure 4-15.

Figure 4-15 Controller Management Screen

For a description of the functions of the menus on the controller management screen, refer to Table 4-4.

Table 4-4 Functions of Menus on the Controller Management Screen

Menu	Function Description
Controller Information	Displays the basic information, firmware, current temperature, and port configuration of the controller.
Configure Controller Settings	Provides advanced configuration options for the controller.
Array Configuration	Creates an array or RAID.
Disk Utilities	Displays the list of disk devices mounted under the controller as well as the basic disk information. It allows you to turn on the disk location indicator, erase disk data and upgrade the firmware.
Set Bootable Device(s) for Legacy Boot Mode	Configures, or clears the primary and secondary boot disks.
Administration	Allows the controller administrator to perform operations, such as upgrading the firmware and restoring factory defaults.

4.3.2 Creating a RAID Volume

Abstract

You can create RAID volumes at different levels as required.
The operations for creating RAID volumes at different levels are similar. This procedure uses creating a RAID 1 volume in UEFI mode as an example.

Prerequisite

- Sufficient SATA and SAS disks are installed on the server.
- The port mode for the disks to be connected is already set. For details, refer to "4.5.1 Setting the Mode of a Port".

Context

For a description of the number of disks required to create a RAID volume, refer to Table 4-5.

Table 4-5 Number of Disks Required for Creating a RAID Volume

RAID Level	Description
RAID 0	RAID 0 requires at least one disks.
RAID 1	RAID 1 requires at least two disks. Disks with different capacities can be used in a RAID 1 volume, but the logical capacity of each member disk depends on the space of the disk with the smallest capacity.
RAID 5	RAID 5 requires at least three disks.
RAID 6	RAID 6 requires at least four disks.
RAID 1+0	RAID 1+0 requires at least four disks. A RAID 1+0 volume consists of at least two RAID 1 volumes. For example, if there are four hard disks to be stored in the RAID 1+0 mode, you need to add them to two "Drive Group," and each "Drive Group" mounts two hard disks. The storage mode is RAID 1.
RAID 50	RAID 50 requires at least six disks. A RAID 50 volume consists of at least two RAID 5 volumes. For example, if six hard disks need to be stored in the RAID 50 mode, it is necessary to add six hard disks to two "Drive Group," each of which is mounted with three hard disks, and the storage mode is RAID 5.
RAID 60	RAID 60 requires at least eight disks. A RAID 60 volume consists of at least two RAID 6 volumes. For example, if there are eight disks to be used in RAID 60 mode, you need to add them to two drive groups, each of which is mounted with four disks in RAID 6 mode.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-16.

Figure 4-16 Array Configuration Screen

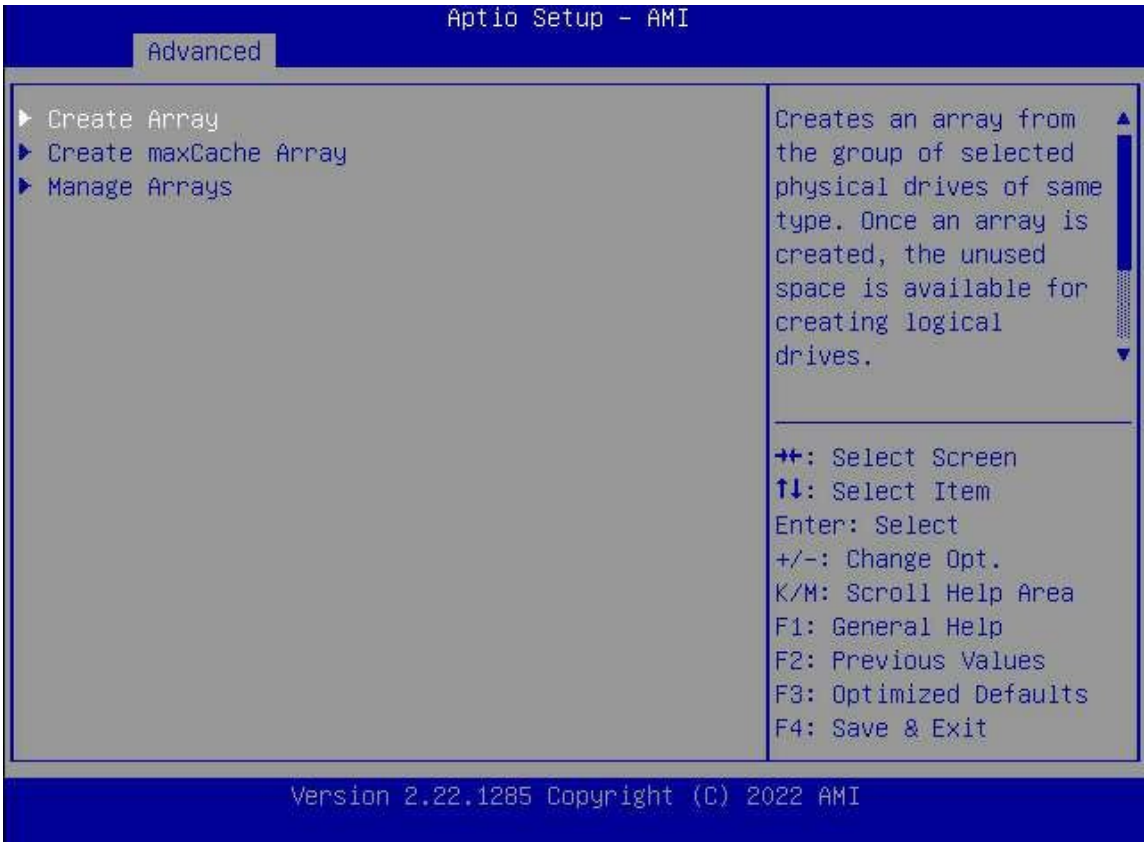

2. Use the arrow keys to select **Create Array**, and then press **Enter**. In the displayed disk list, all the disks that can be used to create a RAID volume are displayed, see Figure 4-17.

Figure 4-17 Viewing the Disk List

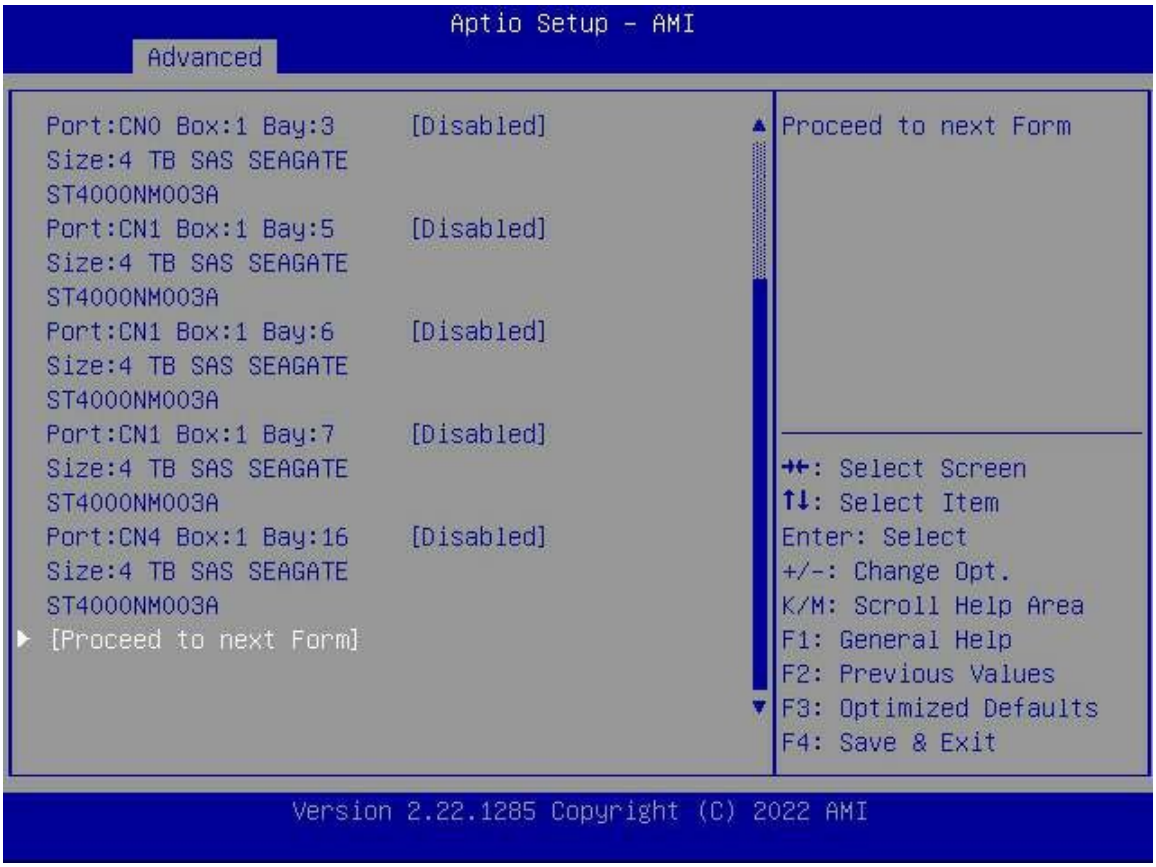

3. Select the disk to be added to the array, and then press **Enter** to set the disk port to **Enabled** status, see Figure 4-18.

Figure 4-18 Confirming the Configuration

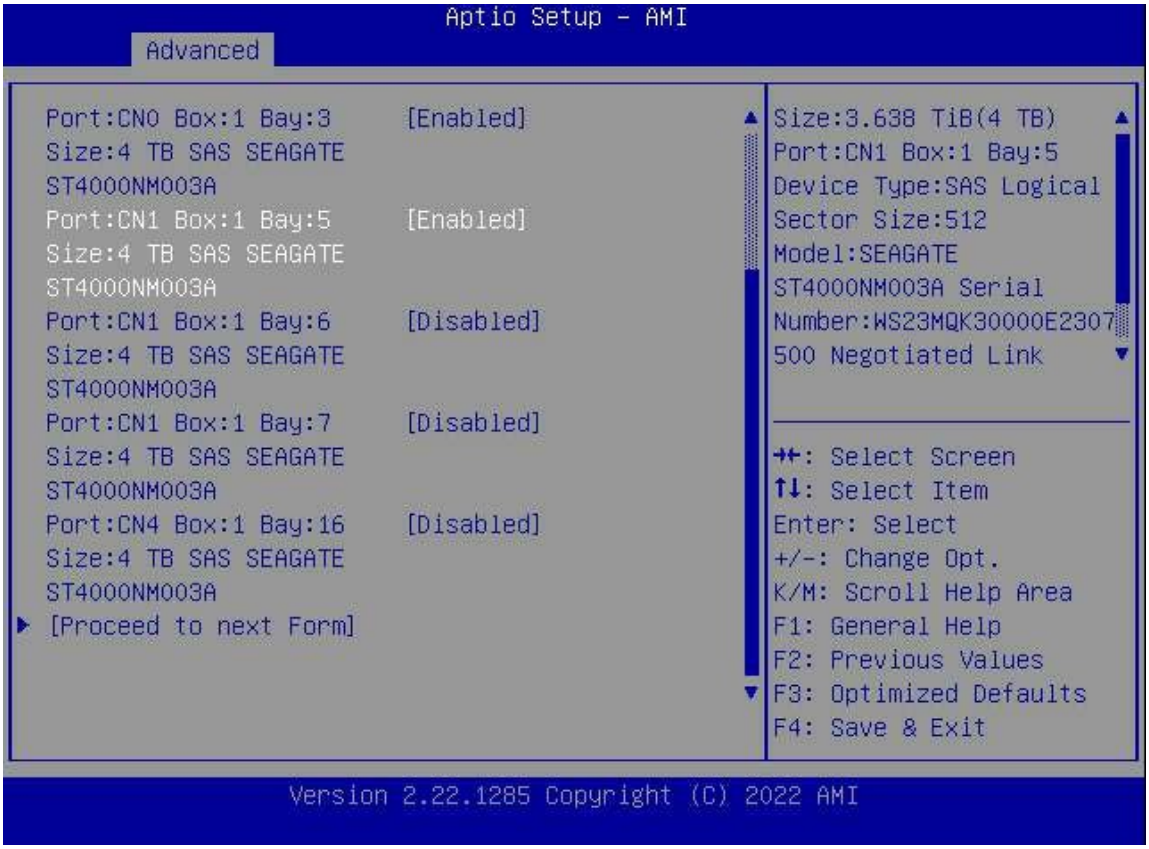

Note

The disks for creating a RAID volume must be of the same type. It is forbidden to select disks with interface types such as SATA and SAS at the same time.

4. Use the arrow keys to select **Proceed to Next Form**, and then press **Enter**. The screen for creating RAID is displayed, see Figure 4-19.

Figure 4-19 Creating RAID

5. Use the arrow keys to select **RAID Level**, and then press **Enter**. From the displayed short-cut menu, select the desired RAID level, see Figure 4-20.

Figure 4-20 Selecting a RAID Level

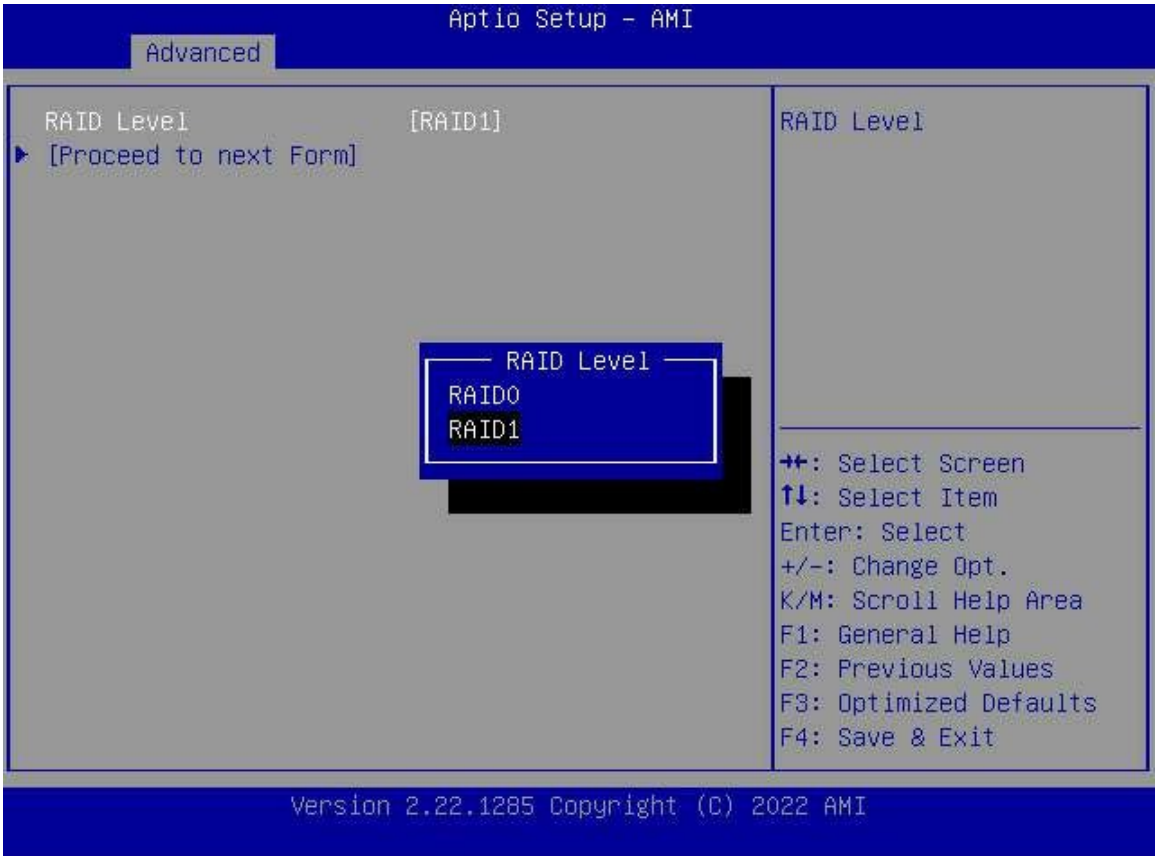

6. Use the arrow keys to select **Proceed to Next Form**, and press **Enter**. The screen for setting RAID logical volumes is displayed, see Figure 4-21.

Figure 4-21 Setting RAID Logical Volumes

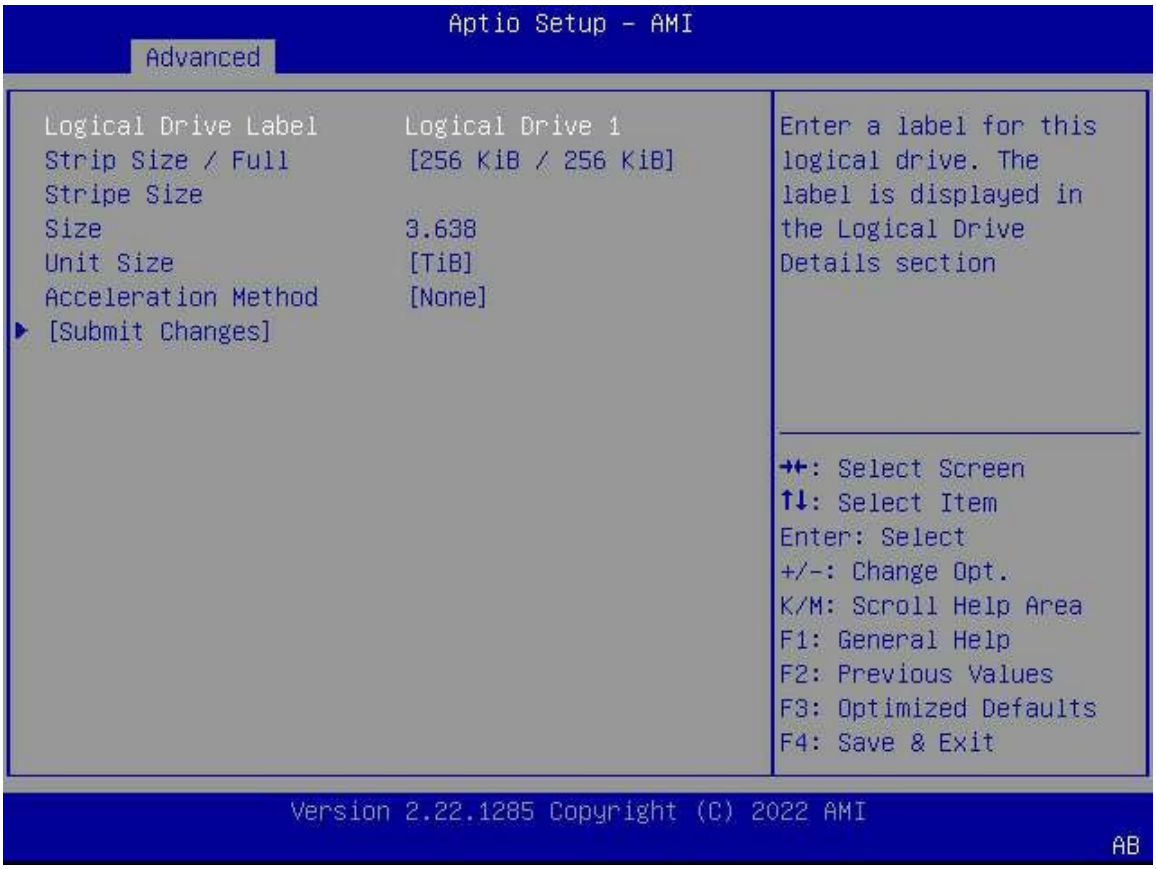

7. Use the arrow keys to select the parameters that you want to modify. In the displayed operation boxes, configure the related parameters. For a description of the parameters for setting a logical volume, refer to Table 4-6.

Table 4-6 Parameter Descriptions for Logical Volume Configuration

Parameter	Description
Logical Drive Label	Sets the name of the RAID logical volume, for example, "Logical Drive 1".
Strip Size/Full Strip Size	● Strip Size indicates the size of the current stripe. The stripe size should be equal to the size of average disk IO requests generated by server applications. In the optimum status, only one IO operation is executed for each IO request. The size can be 16 KiB, 32 KiB, 64 KiB, 128KiB, 512 KiB, or 1024 KiB. The default value is 256 KiB. The recommended stripe size configurations are as follows: → For a Web server, 8 KiB is recommended. → For a groupware server (such as an email server), 16 KiB is recommended. → For a database server, 16 KiB or 32 KiB is recommended. → For a file server, 32 KiB or 64 KiB is recommended. → For a video file server, 64 KiB, 128 KiB, or 256 KiB is recommended.

Parameter	Description
	● Full Strip Size indicates the total size of all stripes. When you set **Strip Size**, the system automatically calculates **Full Stripe Size**.
Size	Displays the storage size of the RAID volume in accordance with the total storage space of the disks added to the RAID volume. By default, all available space is used to create a RAID logical volume. To create multiple RAID logical volumes, you can define the size of the volumes.
Unit Size	Select the unit (MiB/GiB/TiB) of the logical drive.
Acceleration Method	Sets the caching mode for the RAID volume. ● IO Bypass: This option is valid only when the RAID logical volume is formed by SSDs. ● Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time. ● None: disables the controller cache. Neither **IO Bypass** nor **Controller Cache** is used. Keep the default configuration unless otherwise specified.

8. Use the arrow keys to select **Submit Changes**, and press **Enter**. The RAID volume is created successfully, see Figure 4-22.

Figure 4-22 RAID Volume Created Successfully

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.4 Common Configurations (Legacy Mode)

By using the BIOS configuration utility, you can configure and maintain a created RAID volume. For a description of the common operations on a SmartROC 3100 RAID controller card in legacy mode, refer to Table 4-7.

Table 4-7 Common Operations on a SmartROC 3100 RAID Controller Card

Common Operation	Description
Querying RAID volume information	Refer to "4.4.1 Querying RAID Volume Information".
Creating a hot spare disk	Refer to "4.4.2 Creating a Hot Spare Disk".
Deleting a RAID volume	Refer to "4.4.3 Deleting a RAID Volume".
Deleting a hot spare disk	Refer to "4.4.4 Deleting a Hot Spare Disk".
Locating a disk	Refer to "4.4.5 Locating a Disk".

Common Operation	Description
Configuring a pass-through disk	Refer to "4.4.6 Configuring a Pass-Through Disk".
Enabling the caching function	Refer to "4.4.7 Enabling the Caching Function".

4.4.1 Querying RAID Volume Information

Abstract

This procedure describes how to query the RAID information created on a SmartROC 3100 RAID controller card, such as the RAID volume status and member disk status.

Prerequisite

A RAID volume is created successfully. For details, refer to “4.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 4-23.

Figure 4-23 List of Arrays Screen

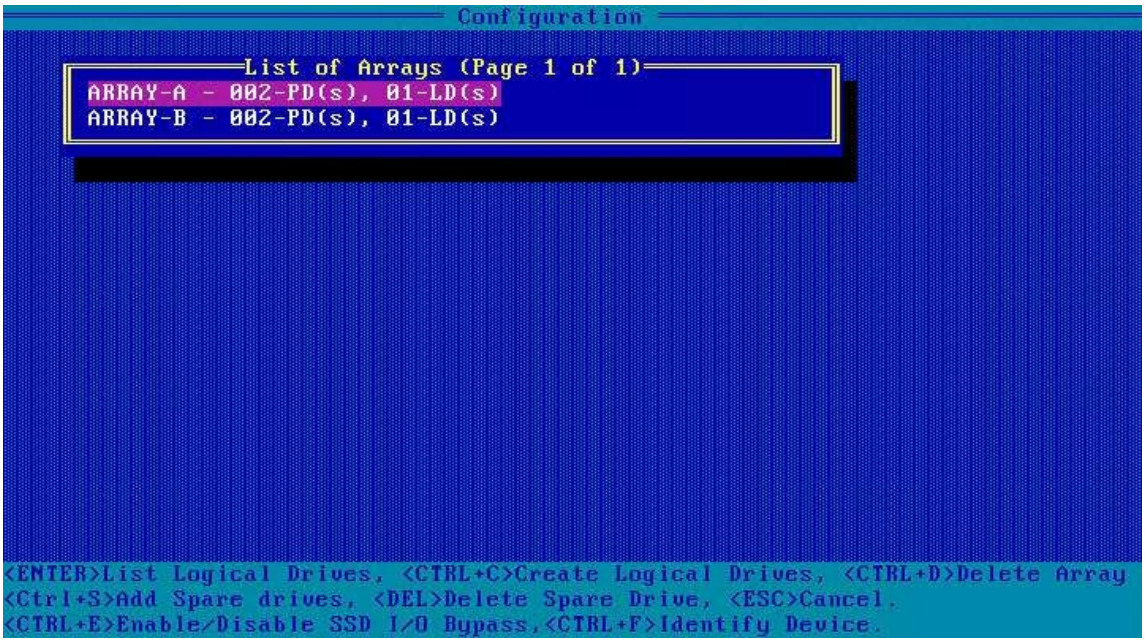

2. Select the RAID volume whose properties you want to view, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 4-24.

Figure 4-24 List of Logical Drives Screen

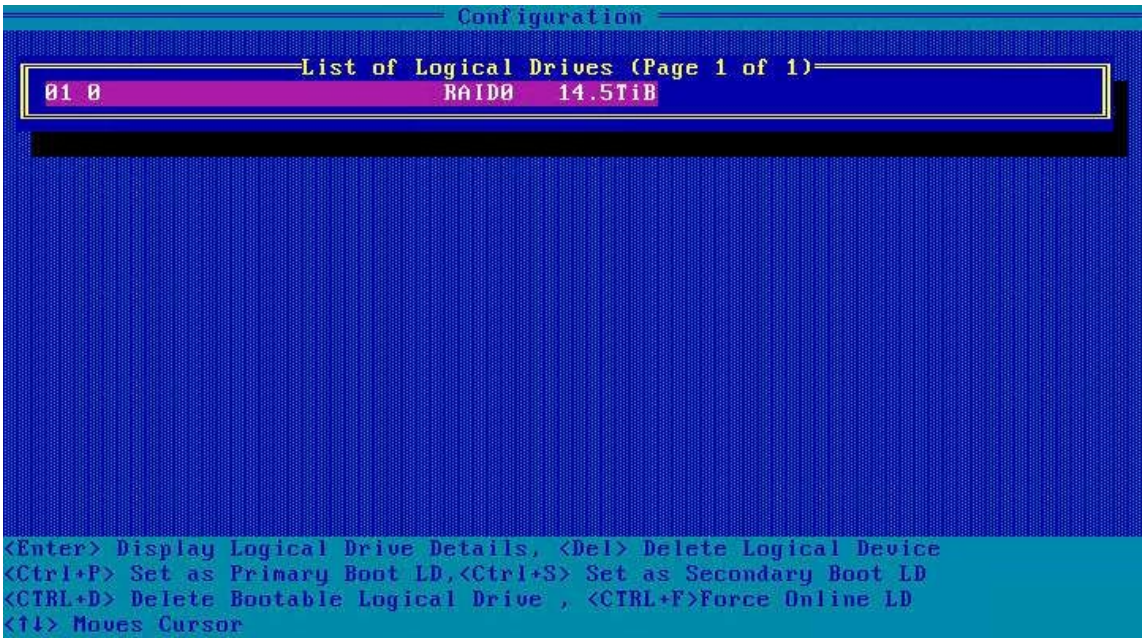

3. Press **Enter**. In the displayed **Logical Drive Details** dialog box, view the property information about the RAID volume, see Figure 4-25.

Figure 4-25 Logical Drive Details Dialog Box

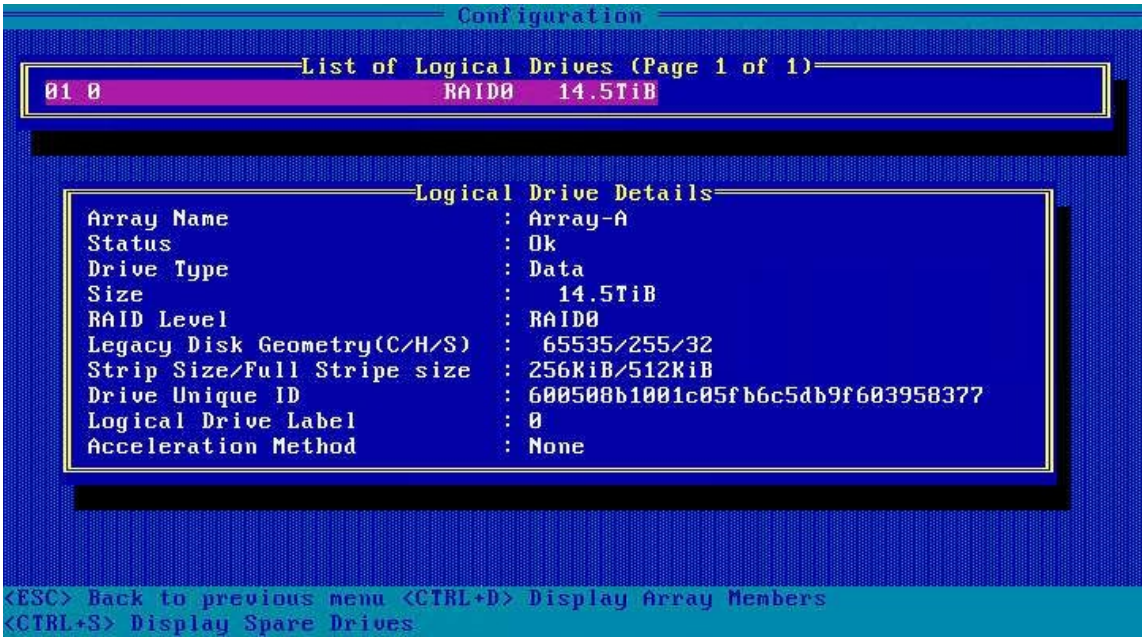

4. (Optional) To view the member disk information, press **Ctrl+D**. The **Array Member Drives** dialog box is displayed, see Figure 4-26.

Figure 4-26 Array Member Drives Dialog Box

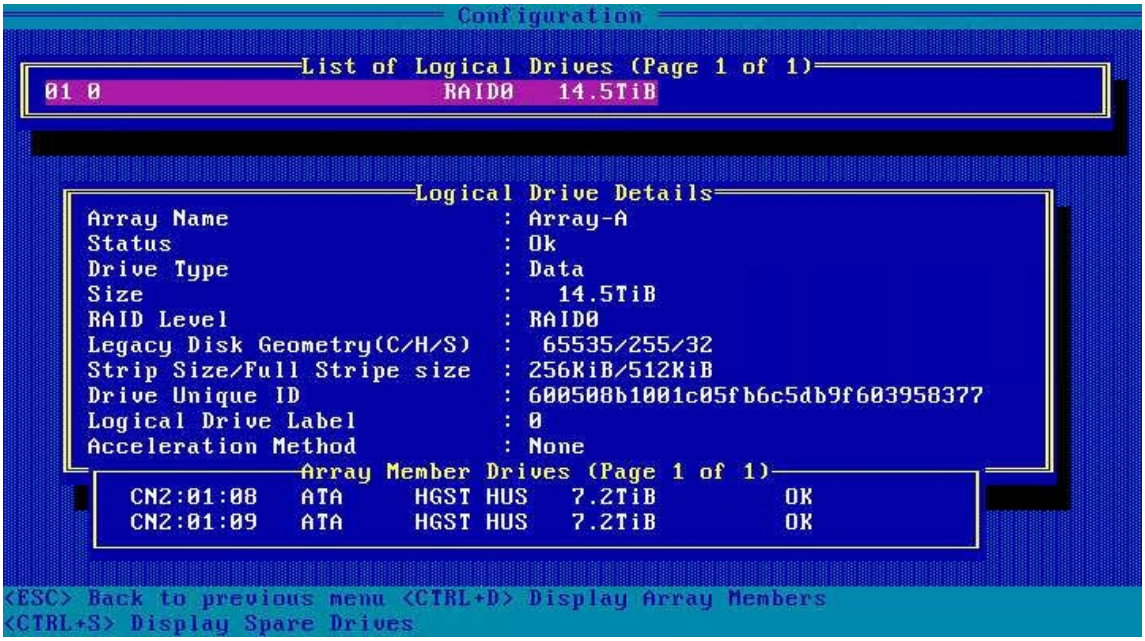

4.4.2 Creating a Hot Spare Disk

Abstract

A hot spare disk improves the data security of a RAID array. For a description of the hot spare disk types supported by a SmartROC 3100 RAID controller card, refer to Table 4-8.

Table 4-8 Hot Spare Disk Types

Type	Description
Dedicated	● This type of hot spare disks is exclusive to the specified one or more disk groups of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a dedicated hot spare disk temporarily takes over the faulty disk.
Auto Replace	● This type of hot spare disks provides the hot standby function for a disk group of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a hot spare disk of this type automatically replaces the faulty disk.

Prerequisite

There are sufficient idle disks on the server.

Context

When creating a hot spare disk, pay attention to the following points:

- Multiple hot spare disks can be created for a disk group, but only one type of hot spare disk can be set at a time. That is, either **Dedicated** or **Auto Replace** is specified.
- An idle disk can be set as a hot spare disk. The disk that has been used to create a RAID volume cannot be set as a hot spare disk.
- The hot spare disk must be of the same type as that of any member disk in the corresponding disk group. That is, all of them are SATA disks or SAS disks, and the hot spare disk's capacity must not be less than the maximum capacity of the member disks.
- Disk groups at all levels except RAID 0 support hot spare disks.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 4-27.

Figure 4-27 List of Arrays Screen

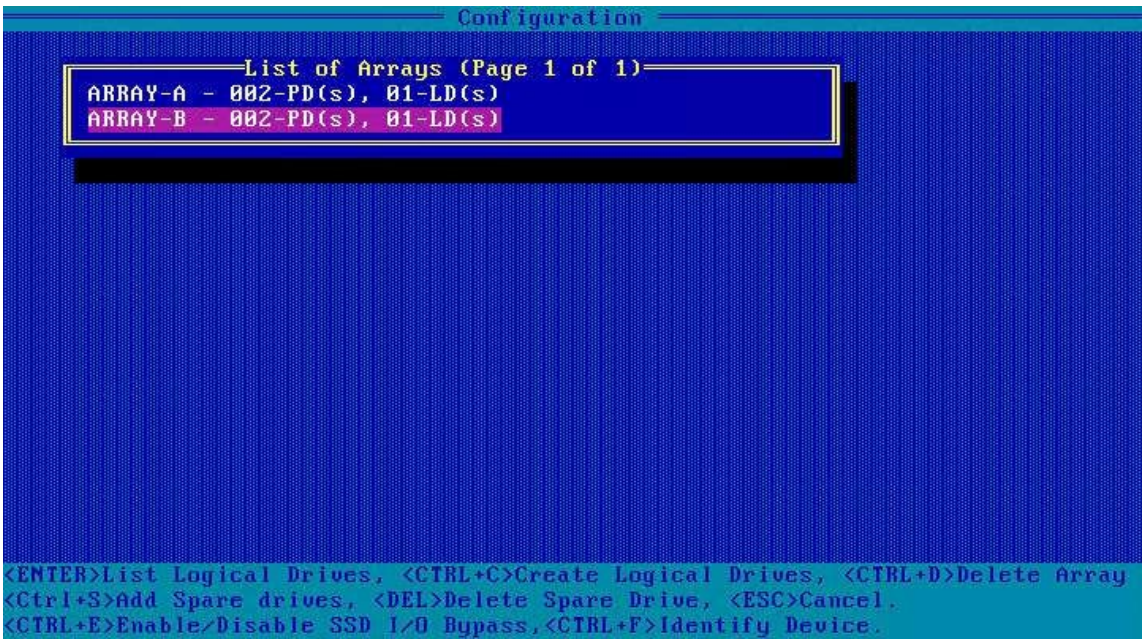

2. Use the arrow keys to select the RAID volume for which you want to configure a hot spare disk, and press **Ctrl+S**. The **Select Hotspare Drives** screen is displayed, see Figure 4-28.

Figure 4-28 Select Hotspare Drives Screen

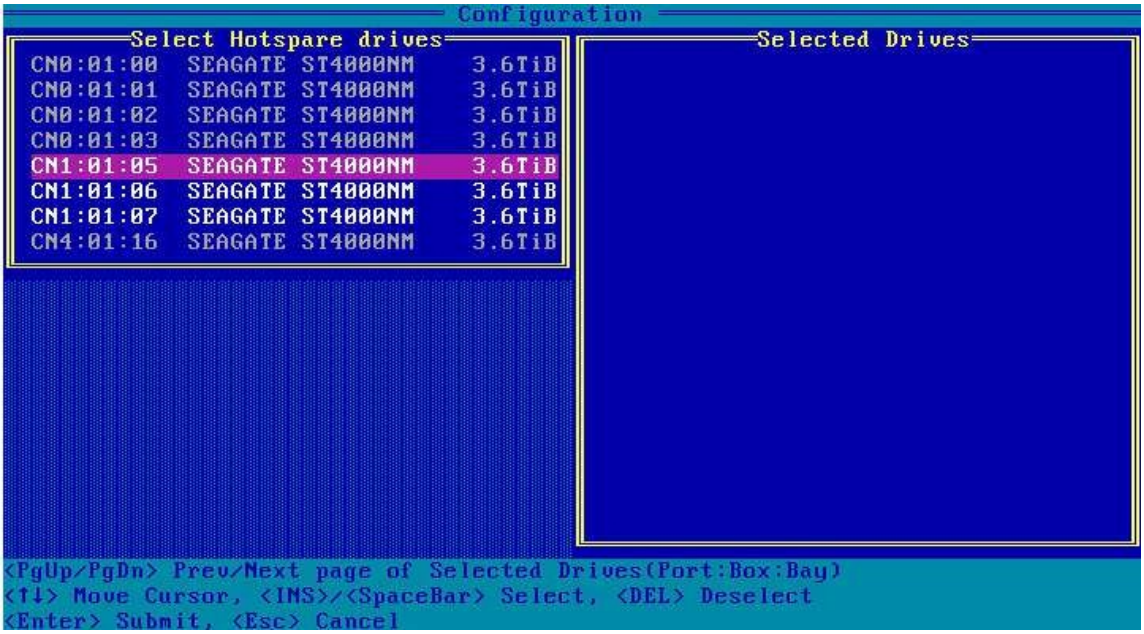

3. Use the arrow keys to select the idle disk to be set as a hot spare disk, and then press **Insert** to add the disk to the **Selected Drives** list, see Figure 4-29.

Figure 4-29 Selected Drives Screen

Note

Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

4. Press **Enter**. A confirmation dialog box is displayed, see Figure 4-30.

Figure 4-30 Confirmation Dialog Box

5. At the cursor in the confirmation dialog box, enter **Y**. The **Select Spare Type** screen is displayed, see Figure 4-31.

Figure 4-31 Select Spare Type Screen

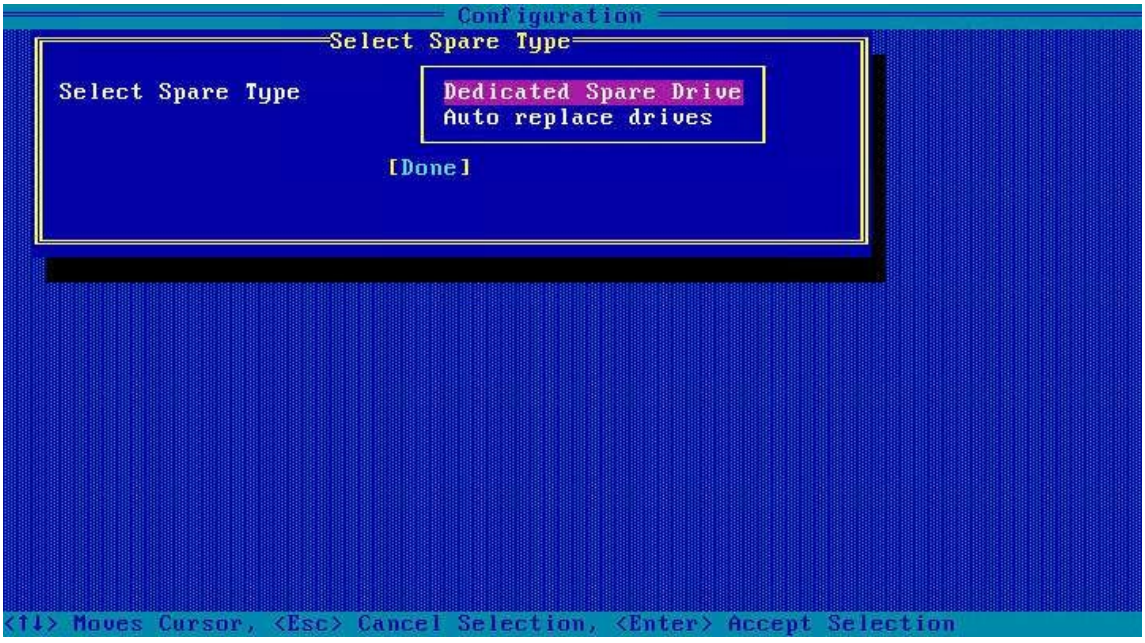

6. Use the arrow keys to select the type of hot spare disk to be created, and then press **Enter** for confirmation, see Figure 4-32.

Figure 4-32 Select Spare Type Screen

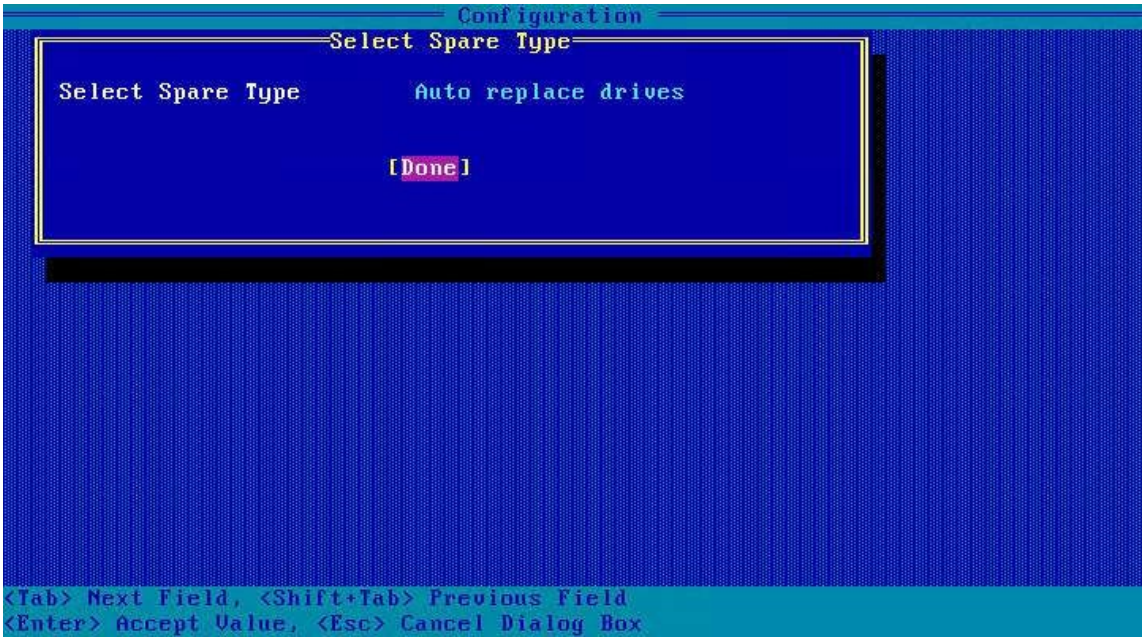

7. Use **Tab** to select **Done**, and then press **Enter** to create the hot spare disk, see Figure 4-33.

Figure 4-33 Creating a Hot Spare Disk

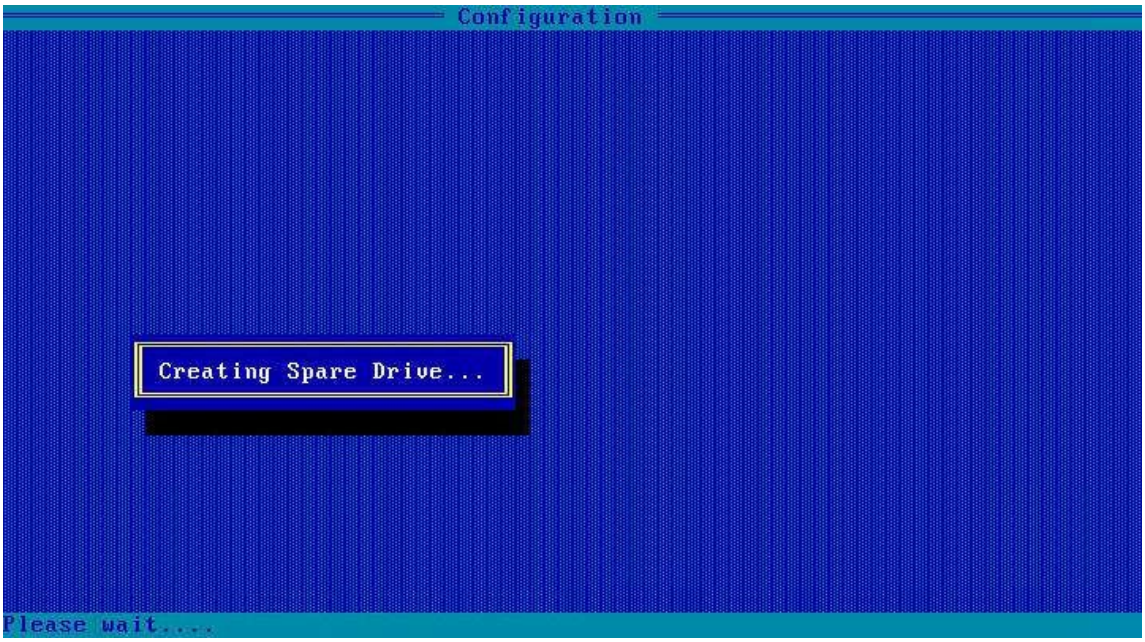

4.4.3 Deleting a RAID Volume

Abstract

When a server no longer needs a RAID volume, you can delete the RAID volume to release the disk space.

Notice

The data that is lost during deletion of the RAID volume cannot be restored. Therefore, you must make sure that you have backed up important data before deleting the volume.

Prerequisite

A RAID volume is created successfully. For details, refer to “4.2.2 Creating a RAID Volume”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 4-34.

Figure 4-34 List of Arrays Screen

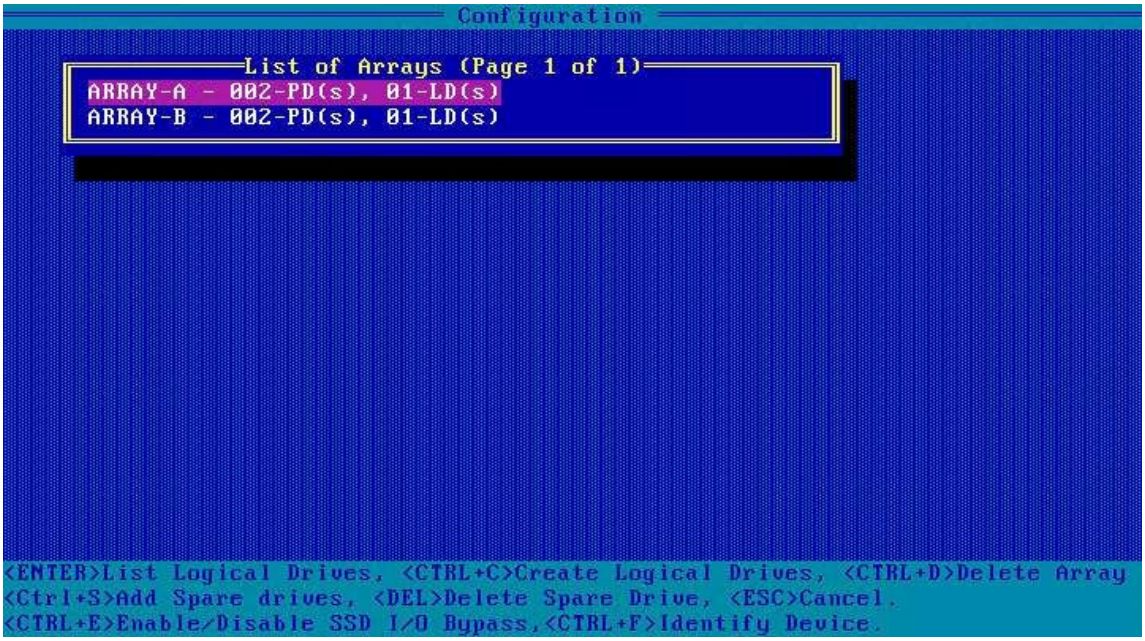

2. Use the arrow keys to select the array where the RAID volume to be deleted is located, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 4-35.

Figure 4-35 List of Logical Drives Screen

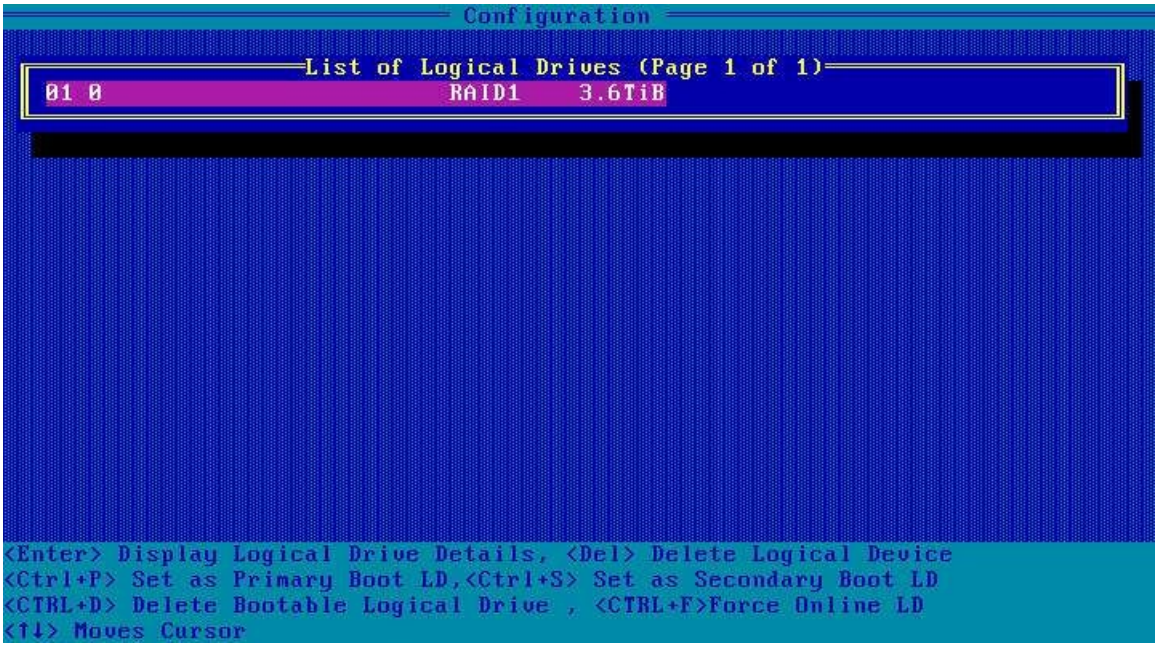

3. Use the arrow keys to select the RAID volume to be deleted, and press **Delete**. A warning dialog box is displayed, see Figure 4-36.

Figure 4-36 Warning Dialog Box

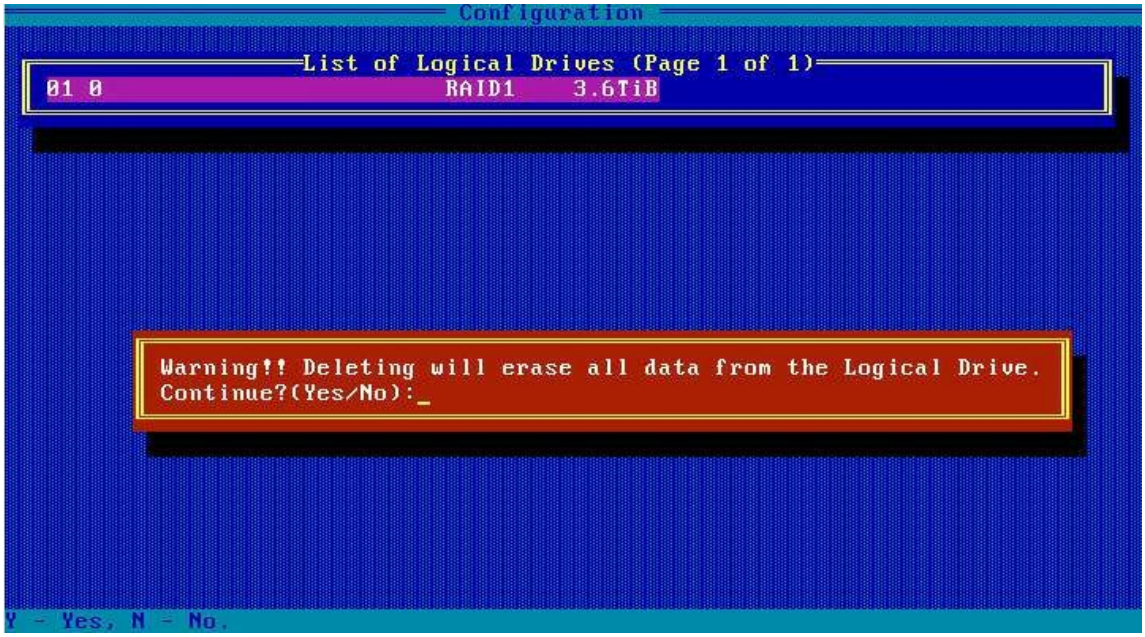

4. At the cursor in the warning dialog box, enter **Y** to delete the selected hot spare disk, see Figure 4-37.

Figure 4-37 Deleting a RAID Volume

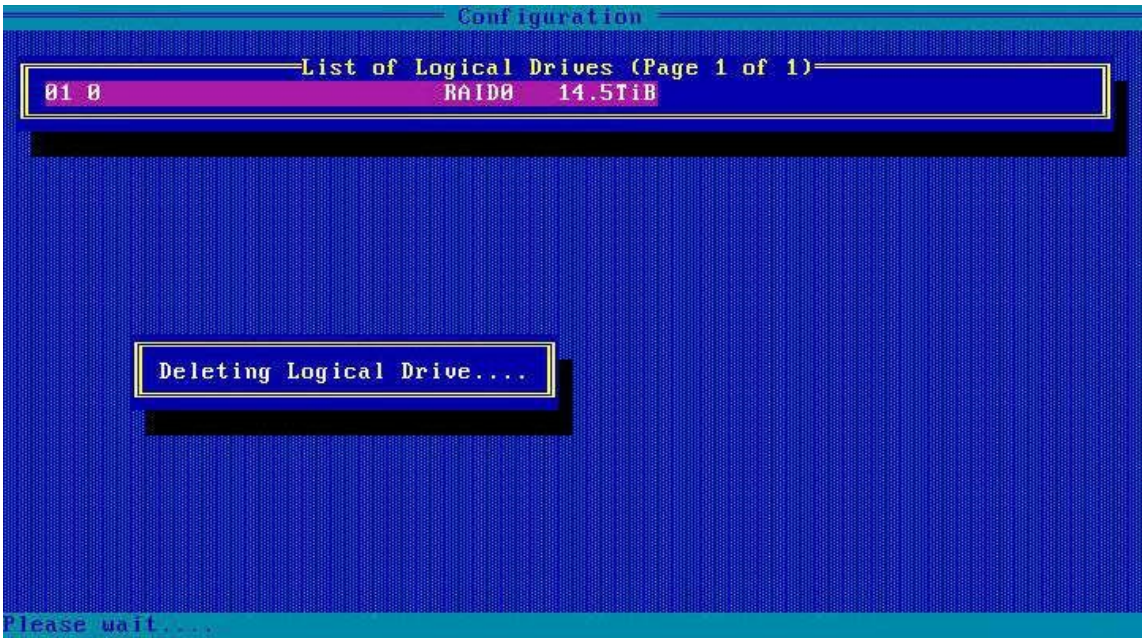

4.4.4 Deleting a Hot Spare Disk

Abstract

When the number of disks of a server cannot meet the requirements, you can delete an existing hot spare disk and restore it to a common disk.

Prerequisite

A hot spare disk is already created. For details, refer to “4.4.2 Creating a Hot Spare Disk”.

Steps

1. On the **Configuration Menu** screen, use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 4-38.

Figure 4-38 List of Arrays Screen

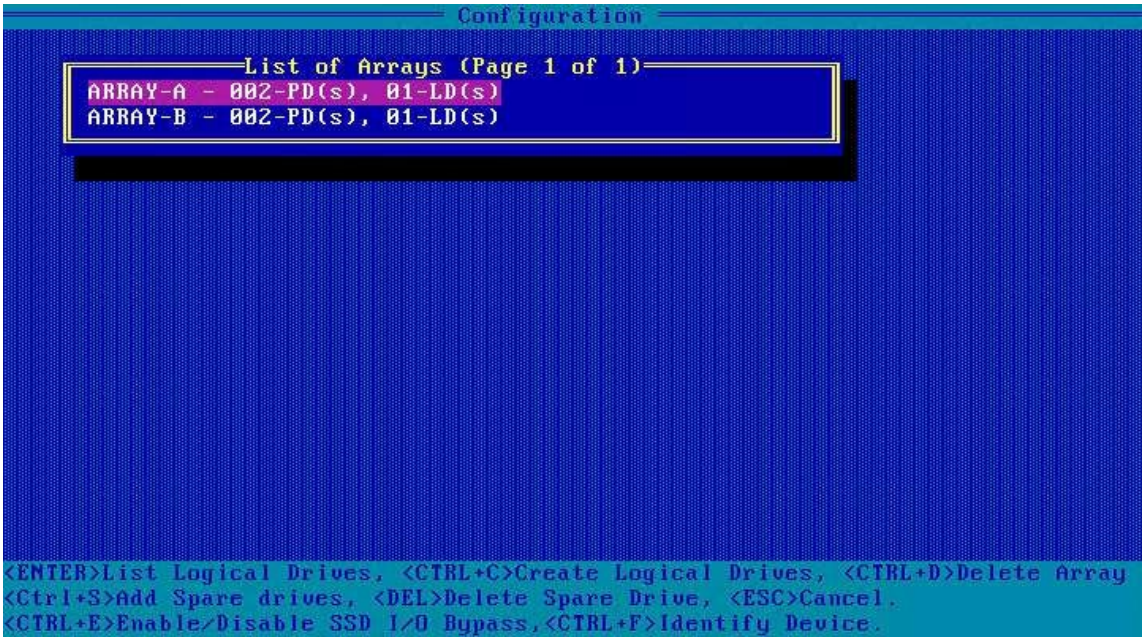

2. Use the arrow keys to select the array where the hot spare disk to be deleted is located, and press **Delete**. The **Select Hotspare drives** screen is displayed, see Figure 4-39.

Figure 4-39 Select Hotspare Drives Screen

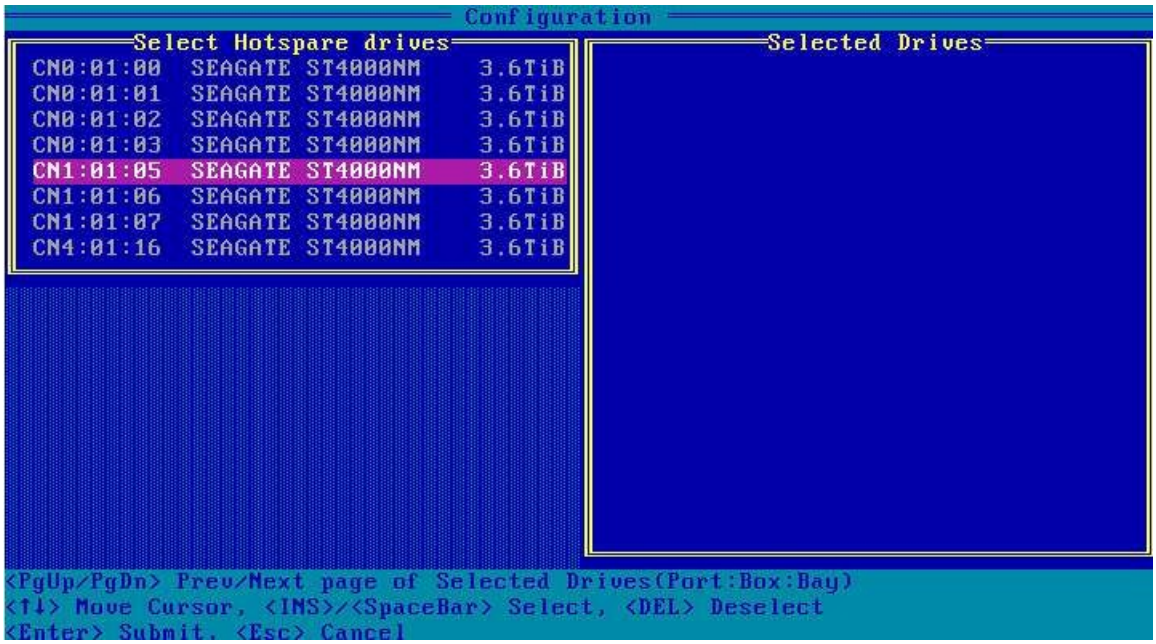

3. Use the arrow keys to select the hot spare disk to be deleted, and then press **Insert** to add the disk to the **Selected Drives** list, see Figure 4-40.

Figure 4-40 Selected Drives List

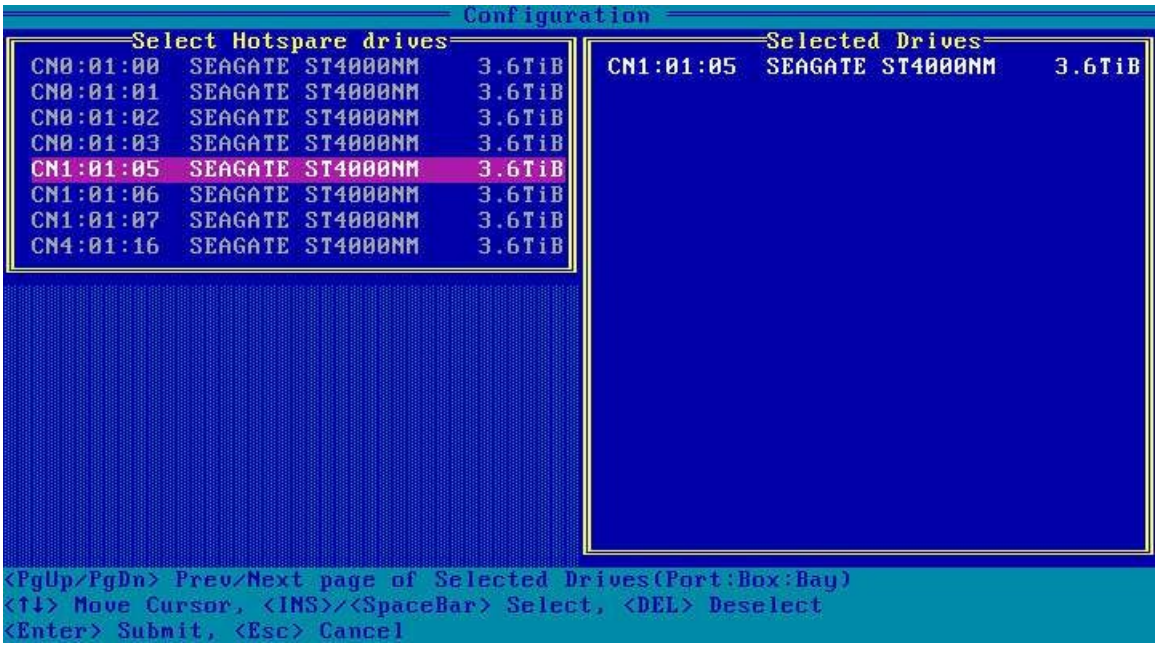

Note

Press the **Delete** key to delete the selected disk from the **Selected Drives** list.

4. Press **Enter**. A confirmation dialog box is displayed, see Figure 4-41.

Figure 4-41 Confirmation Dialog Box

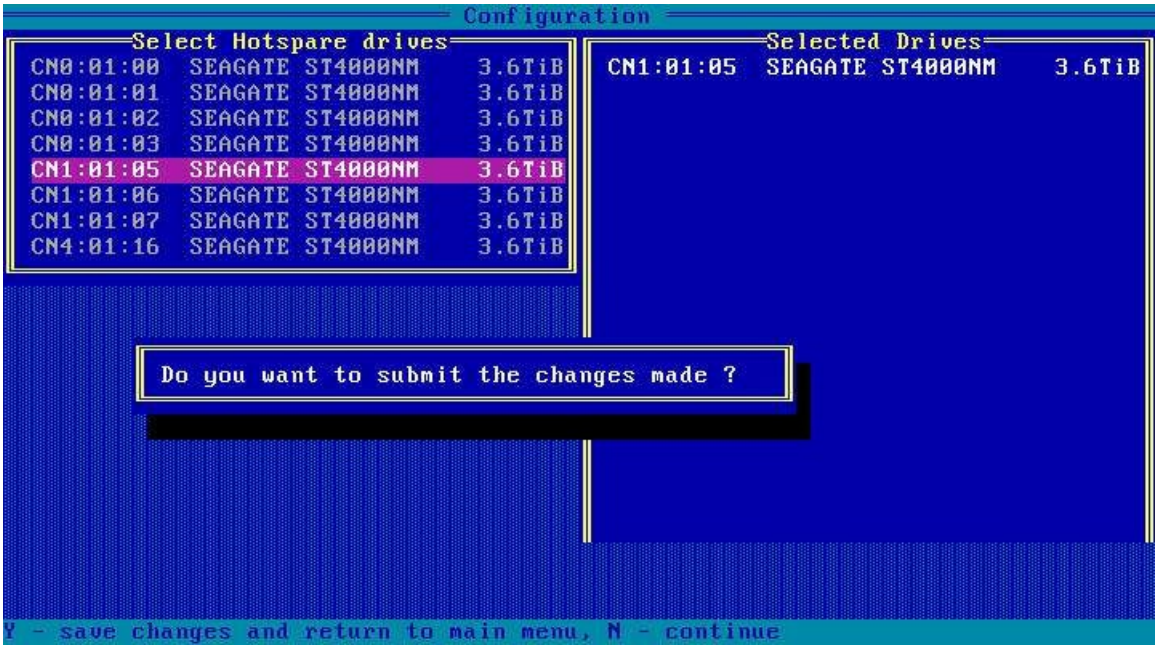

5. At the cursor in the confirmation dialog box, enter **Y** to delete the selected hot spare disk, see Figure 4-42.

Figure 4-42 Deleting a Hot Spare Disk

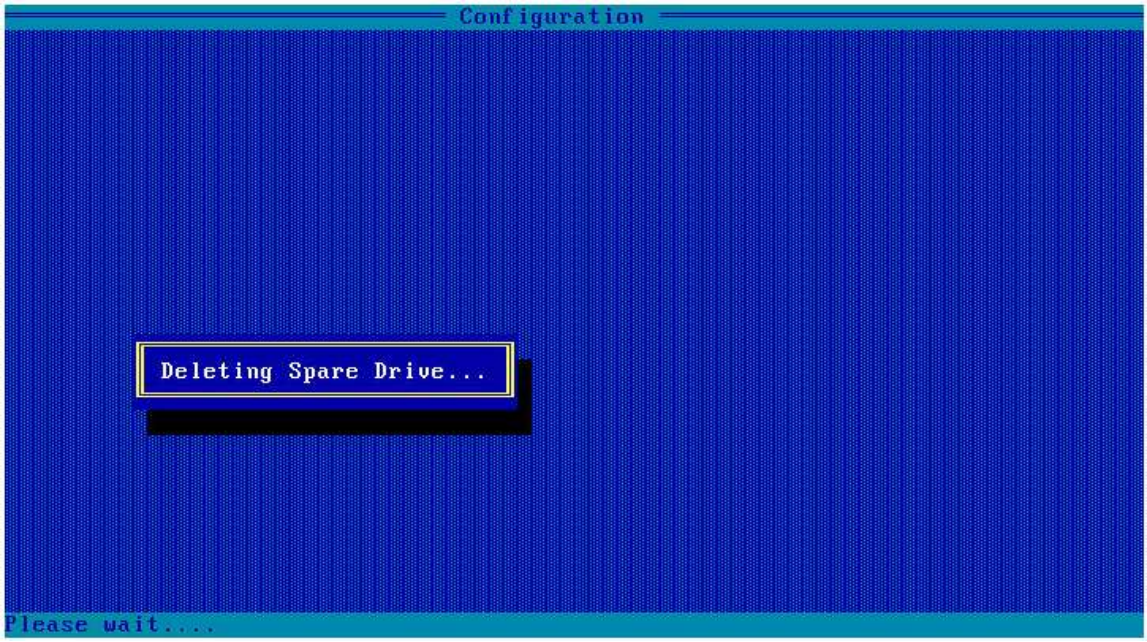

4.4.5 Locating a Disk

Abstract

After the indicator of a disk is lit, you can locate the disk so that you can easily replace or maintain it.

Steps

1. In the Options area on the BIOS configuration utility screen, use the arrow keys to select **Disk Utilities**, and then press **Enter**. The **Select Disks and press <Enter>** screen is displayed, see Figure 4-43.

Figure 4-43 Select Disks and Press <Enter> Screen

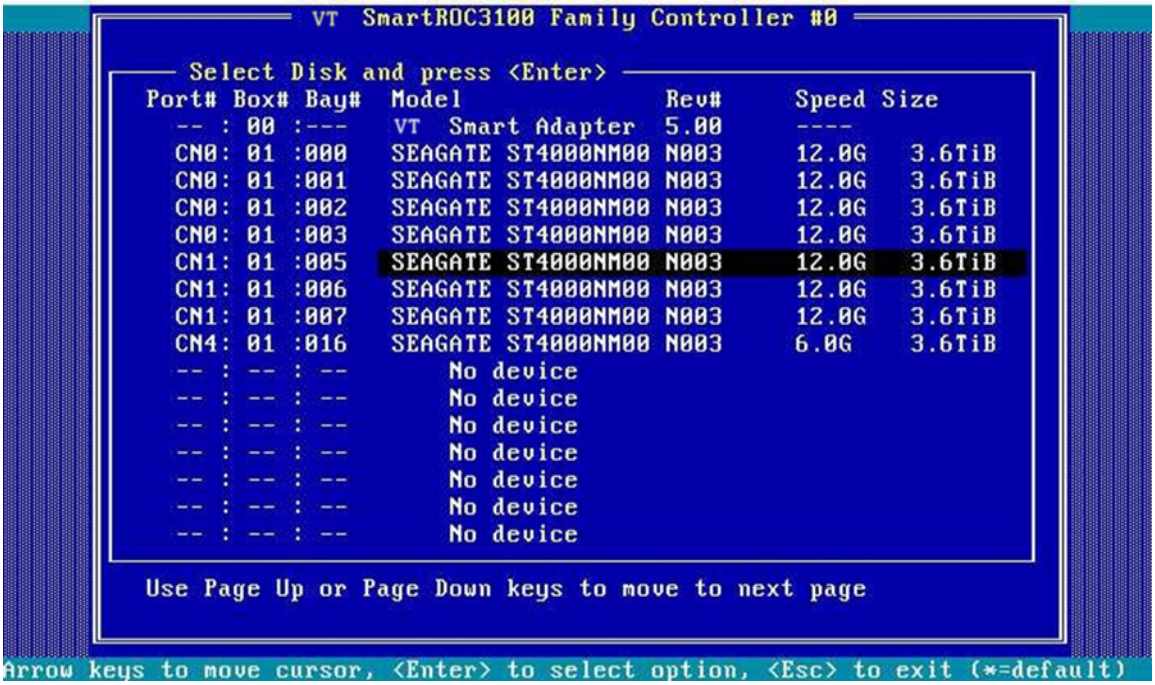

2. Use the arrow keys to select the disk to be located, and then press **Enter**. A function menu is displayed, see Figure 4-44.

Figure 4-44 Function Menu

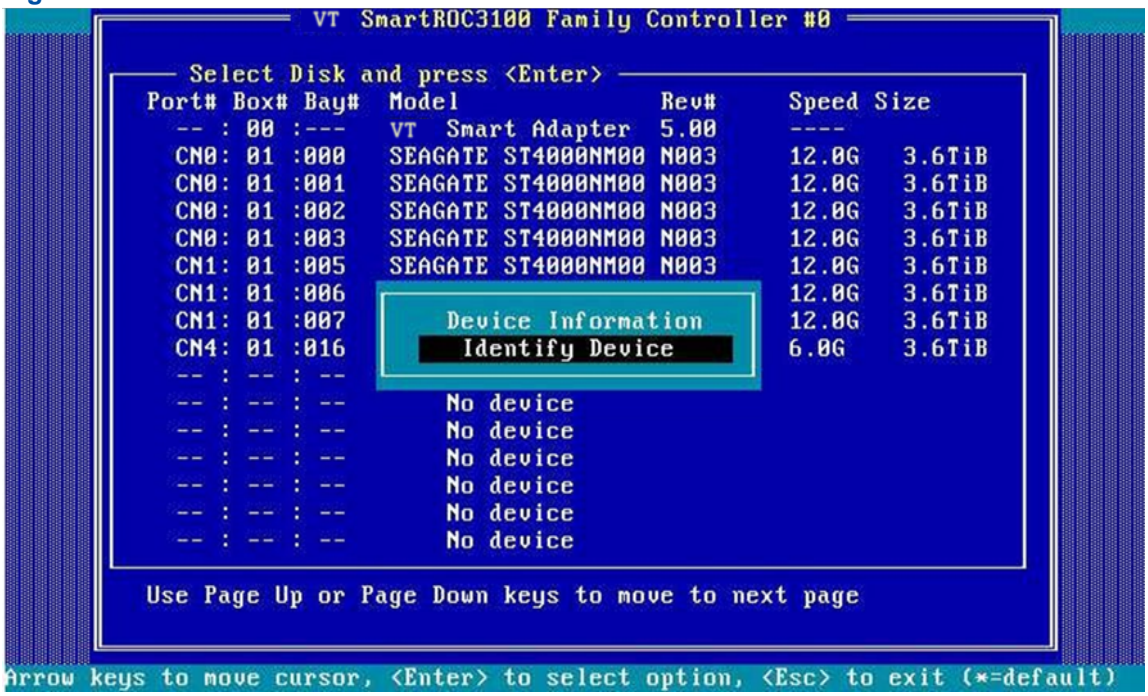

3. Use the arrow keys to select **Identify Device**, and then press **Enter**. A prompt message is displayed, see Figure 4-45. At this time, the red indicator of the corresponding disk is lit and flashes continuously.

Figure 4-45 Disk Indicator Lit

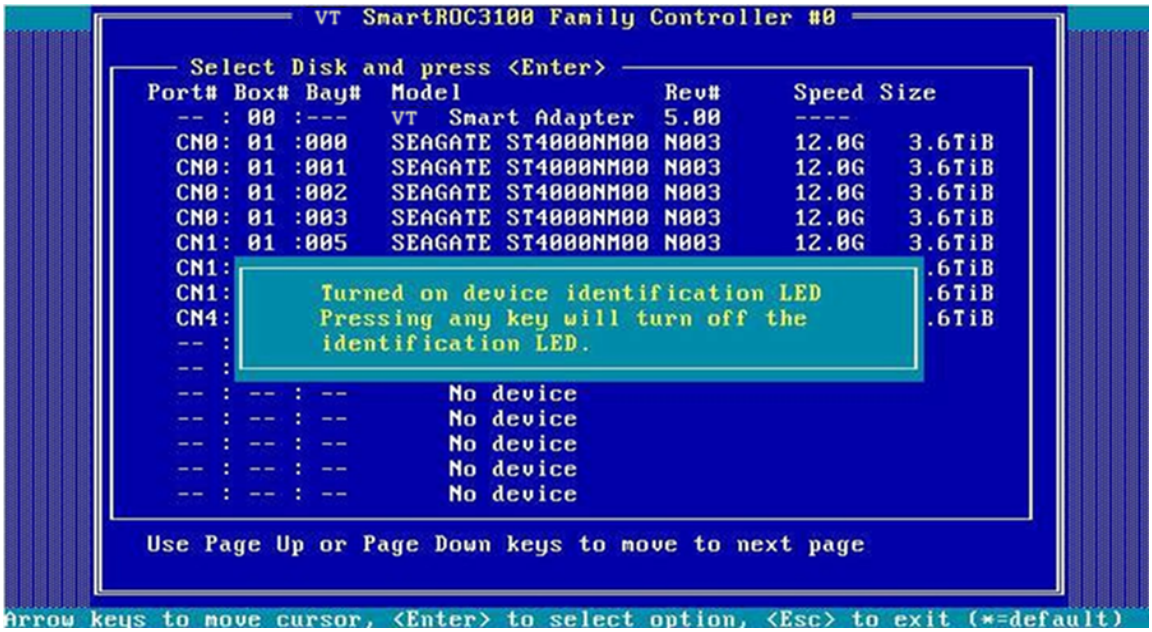

4. Press any key to go out the disk indicator. The disk locating ends.

4.4.6 Configuring a Pass-Through Disk

Abstract

In accordance with the actual RAID configuration, you can configure a pass-through disk in either of the following ways:

- Setting the mode of all ports in batches:
  If the disks connected to all ports are not used to create a RAID volume, the mode of these ports can be set to HBA in a unified manner.
- Setting the mode of ports separately:
  If the disks connected to some ports have been used to create a RAID volume, the mode of these ports can be set to Mixed and that of other ports can be set to HBA.

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.
- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

1. In the Options area on the BIOS configuration utility screen, use the arrow keys to select **Controller Details**, and then press **Enter**. The controller details are displayed. The value of **Controller Mode** is **MIXED**, see Figure 4-46.

Figure 4-46 Controller Details

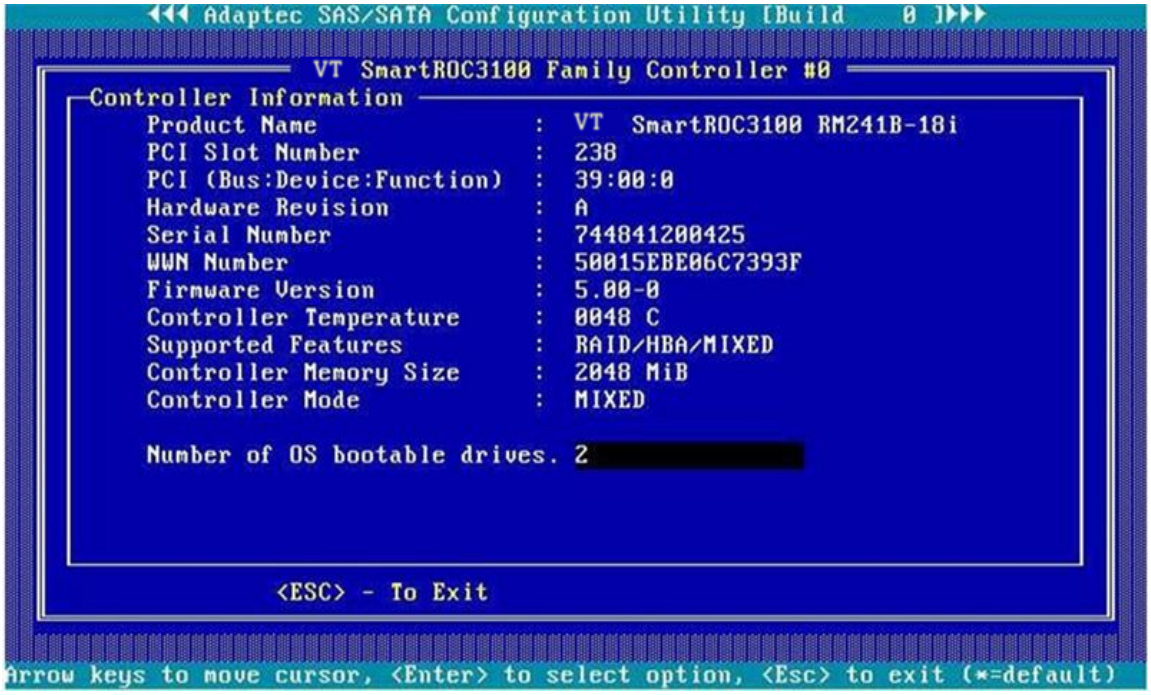

2. Press **ESC** to return to the BIOS configuration screen. Use the arrow keys to select **Configure Controller Settings** and press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-47.

Figure 4-47 Configure Controller Settings Screen

3. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**. The **Configure Controller Port Mode** screen is displayed, see Figure 4-48.

Figure 4-48 Configure Controller Port Mode Screen

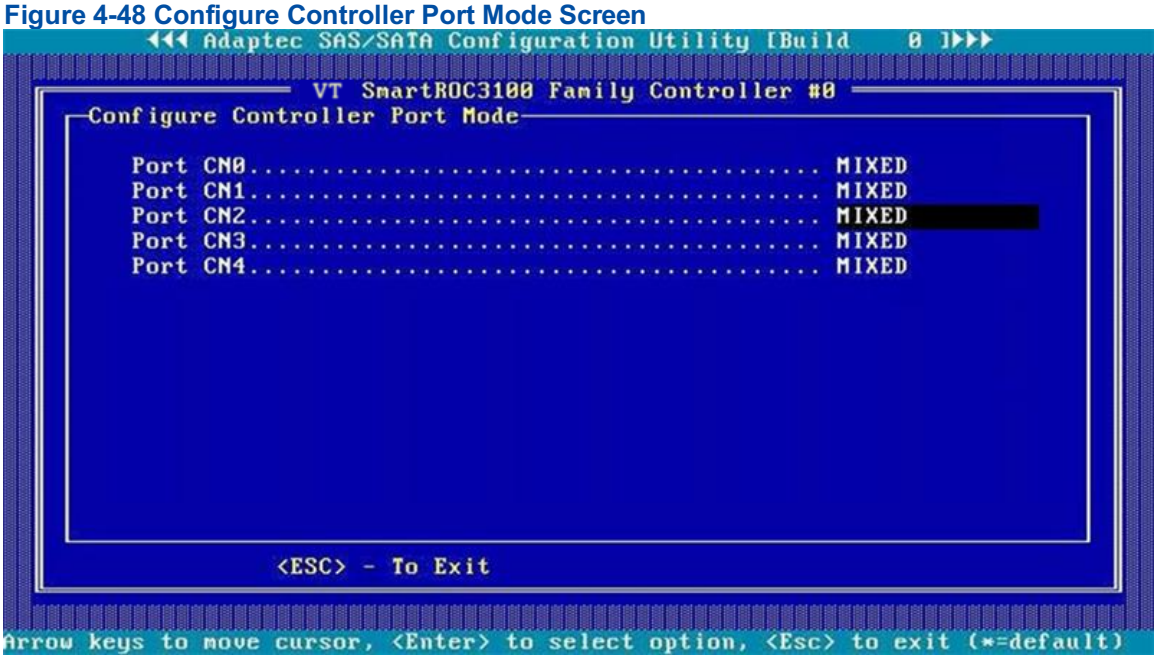

4. Use the arrow keys to select the port whose connected disk is not used to create a RAID volume, and press **Enter**. In the displayed dialog box, set the port mode to **HBA**, see Figure 4-49.

Figure 4-49 Setting the Port Mode to HBA

5. (Optional) If the connected disk is already used to create a RAID volume, set the port mode to **MIXED**, see Figure 4-50.

Figure 4-50 Setting the Port Mode to MIXED

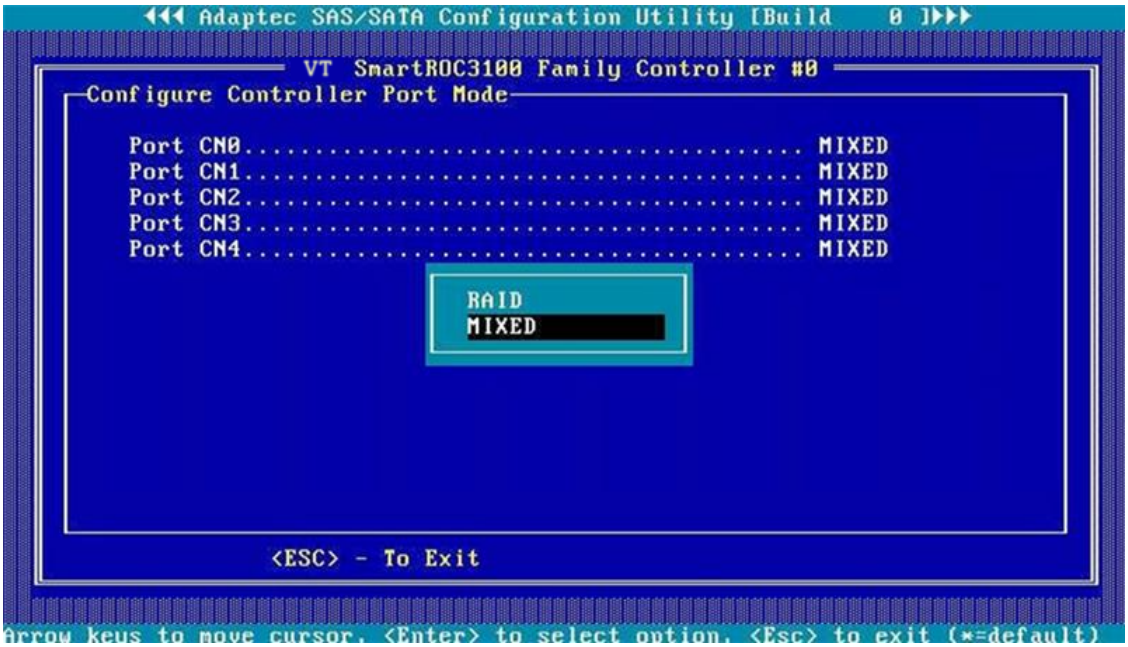

Note

The mode of the ports whose connected disks are already used to create a RAID volume cannot be set to HBA.

6. Set the port mode as required, and then press **Esc** to exit. In the displayed dialog box, select **Yes** to save the configuration, see Figure 4-51.

Figure 4-51 Saving the Configuration

7. Wait until a configuration success message is displayed, see Figure 4-52.

Figure 4-52 Successful Configuration

8. Press **Esc** multiple times until the **Exit Utility** dialog box is displayed, see Figure 4-53.

Figure 4-53 Exit Utility Dialog Box

9. In the **Exit Utility** dialog box, select **NO**, and then press **Enter**. The BIOS configuration utility screen is displayed, see Figure 4-54.

Figure 4-54 BIOS Configuration Utility Screen

10. Use the arrow keys to select **Controller Details**, and then press **Enter**. The controller details are displayed, see Figure 4-55. The value of **Controller Mode** is **HBA**.

Figure 4-55 Controller Details

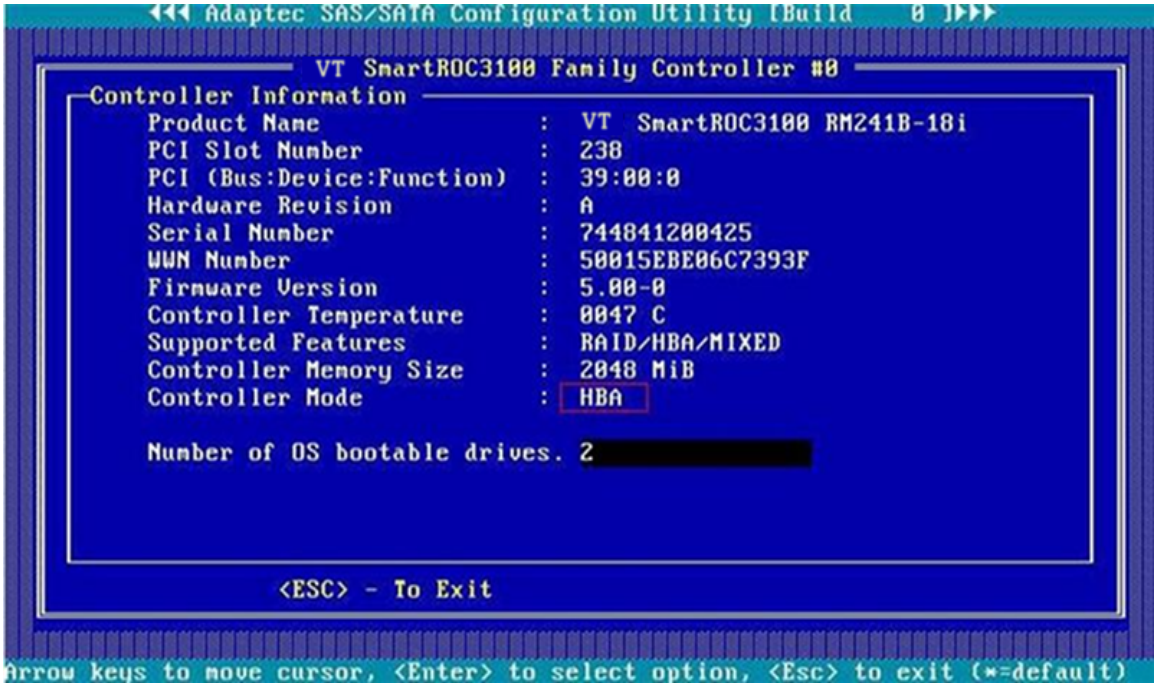

4.4.7 Enabling the Caching Function

Abstract

A VT SmartROC 3100 RAID controller card supports the caching function. In legacy mode, a SmartROC 3100 RAID supports the following three caching modes:

- IO Bypass: valid only when a RAID logical volume is formed by SSDs.
- Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time.
- None: disables the controller cache. Neither **IO Bypass** nor **Controller Cache** is used.

After you select **Controller Cache** mode, the caching function is enabled, which improves the data read/write speed.

Note

You can enable the caching function only by referring to this procedure. The caching function cannot be enabled on the Web portal of the BMC.

Prerequisite

A RAID volume is created successfully. For details, refer to "4.2.2 Creating a RAID Volume".

Context

Enabling the caching function improves the data read/write speed. The details are as follows:

- When a RAID controller card reads the data, if the data has been written into the Cache, the data can be directly read from the Cache to prevent the hard disk from searching for the data again, thus saving the response time and improving the data read speed.
- When a RAID controller card writes the data , the data is directly written into the Cache. The RAID controller card refreshes the data to the hard disk only when the written data is accumulated to a certain extent, achieving batch data write. In addition, the Cache is a fast read/write device, so the read/write speed of the Cache is higher than that of the hard disk, thus improving the data write speed.

Note

To ensure the data security in the Cache, you can configure a super capacitor for the RAID controller card. In case of unexpected power failure of the server, the super capacitor is used to supply power, and provides data security protection in the Cache.

Steps

1. Switch to UEFI mode. For details, refer to "5.3.2 Setting the Boot Mode to UEFI".
2. Perform the following operations as required.

If...	Then...
The super capacitor is configured	View the status of the backup power. The value of **Backup Power Status** is **charged**.
The super capacitor is not configured	a. Disable **Survival Mode**. i. Use the arrow keys to select **VT SmartROC3100 RS241-18i > Array Configuration > Manage Arrays**, and press **Enter**. The **Manage Arrays** screen is displayed. ii. Use the arrow keys to select the array for which you need to disable **Survival Mode**, and then press **Enter**. The screen for managing the selected array is displayed. iii. Use the arrow keys to select **I/O Bypass Settings**, and press **Enter**. In the displayed dialog box, set **I/O Bypass** to **Disabled**. b. Enable **No Battery Write Cache**. i. Use the arrow keys to select **VT SmartROC3100 RS241-18i > Configure Controller Settings > Modify Cache Settings**, and press **Enter**. The **Modify Cache Settings** screen is displayed. ii. Use the arrow keys to select **No Battery Write Cache**, and press **Enter**. In the displayed dialog box, set **No Battery Write Cache** to **Enabled**.

Note

When the super capacitor is not configured and the **No battery Write Cache** is disabled, there is no cache setting option in Legacy mode.

3. Return to Legacy mode. For details, refer to "5.3.1 Setting the Boot Mode to Legacy".
4. In the **Options** area, use the arrow keys to select **Configuration Menu > Manage Arrays**, and then press **Enter**. The **List of Arrays** screen is displayed, see Figure 4-56.

Figure 4-56 List of Arrays Screen

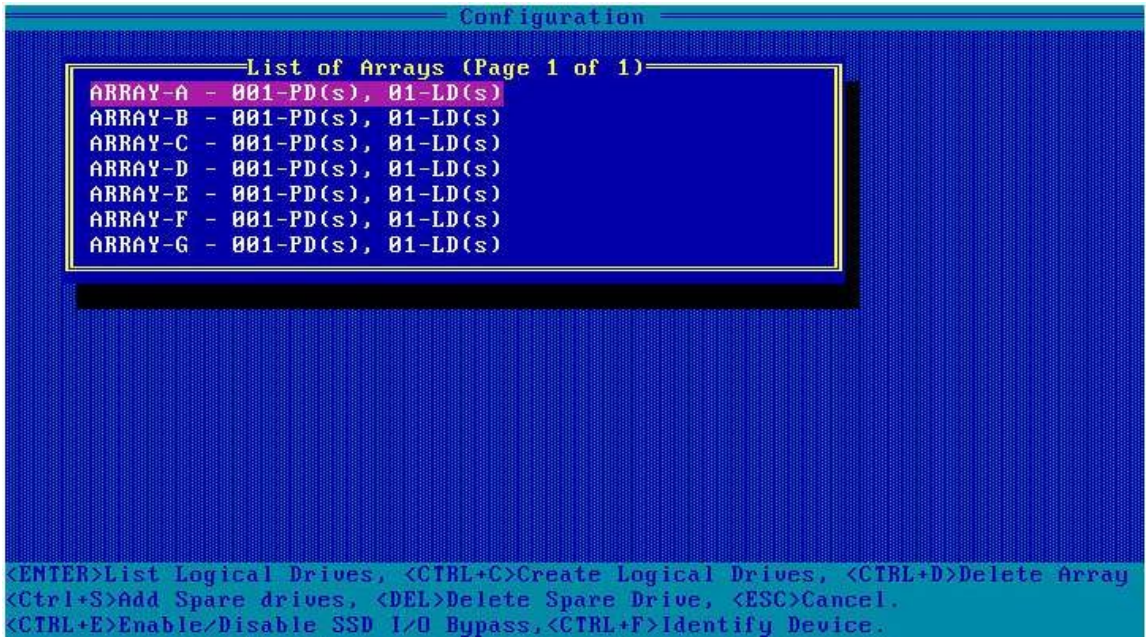

5. Use the arrow keys to select the array for which the caching function needs to be enabled, and then press **Enter**. The **List of Logical Drives** screen is displayed, see Figure 4-57.

Figure 4-57 List of Logical Drives Screen

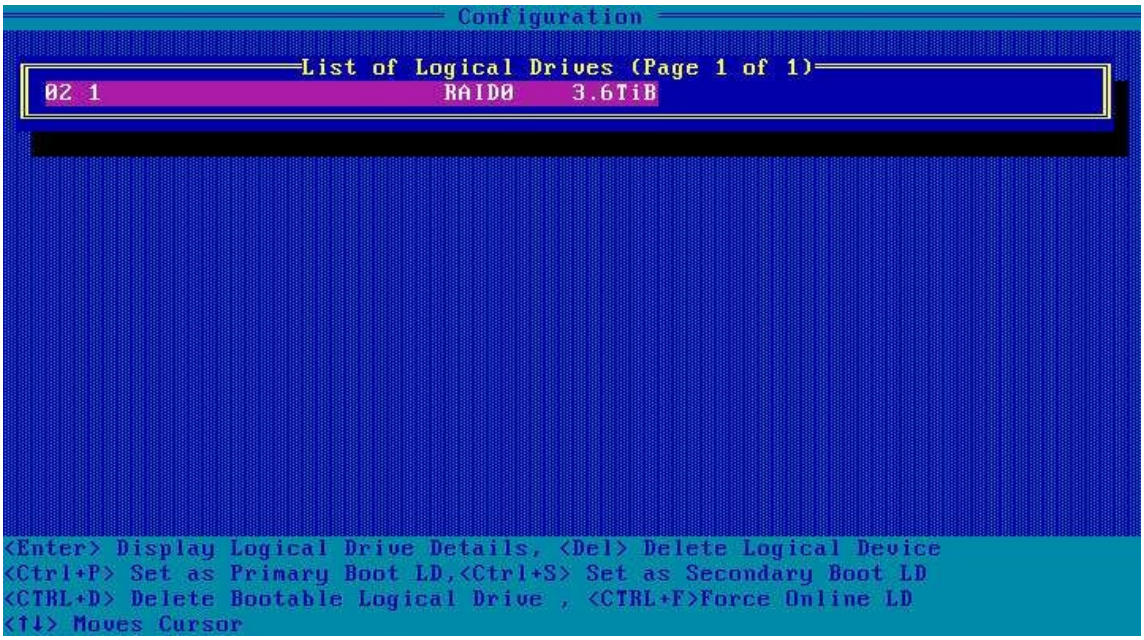

6. Press **Enter**. In the displayed **Logical Drive Details** dialog box, view the property information about the RAID volume, see Figure 4-58. The value of **Acceleration Method** is **None**, indicating that the caching function is disabled.

Figure 4-58 Logical Drive Details Dialog Box

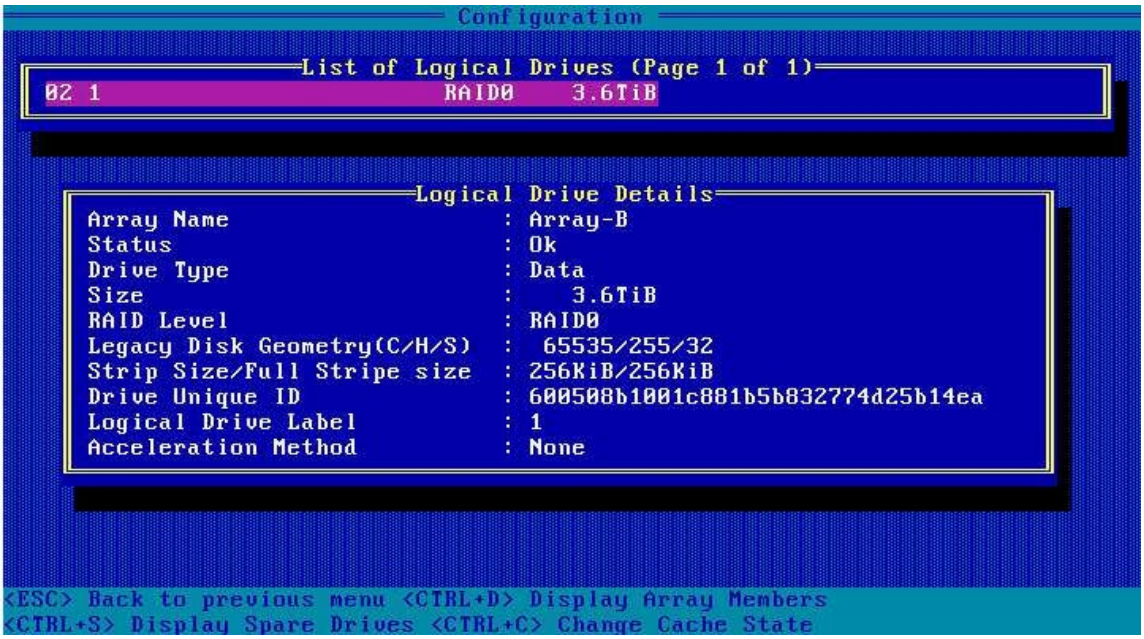

7. Press **Ctrl+C** to switch the caching mode. On the screen as shown in Figure 4-59, the value of **Acceleration Method** becomes **Controller Cache**, indicating that the caching function is enabled.

Figure 4-59 Enabling the Caching Function

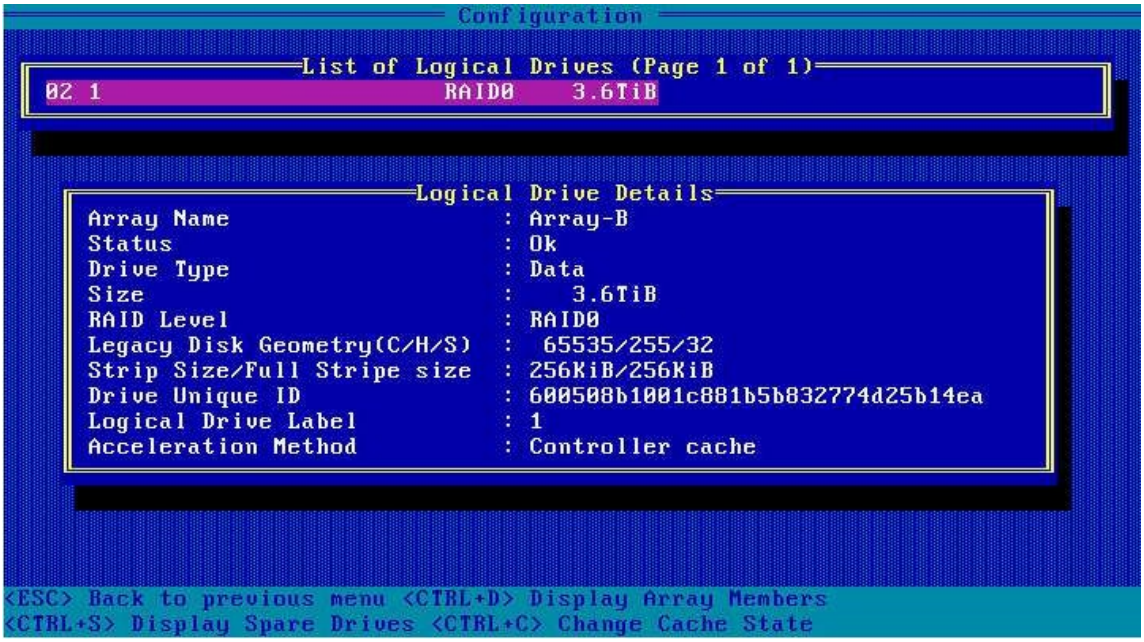

4.5 Common Configurations (UEFI Mode)

By using the BIOS configuration utility, you can configure and maintain a created RAID volume. For a description of the common operations on a SmartROC 3100 RAID controller card in UEFI mode, refer to Table 4-9.

Table 4-9 Common Operations on a SmartROC 3100 RAID Controller Card

Common Operation	Description
Setting a port mode	Refer to "4.5.1 Setting the Mode of a Port".
Locating a disk	Refer to "4.5.2 Locating a Disk".
Creating a hot spare disk	Refer to "4.5.3 Creating a Hot Spare Disk".
Changing a hot spare disk	Refer to "4.5.4 Changing a Hot Spare Disk".
Deleting a hot spare disk	Refer to "4.5.5 Deleting a Hot Spare Disk".
Configuring a power mode	Refer to "4.5.6 Configuring the Performance or Power Mode".
Deleting a RAID volume	Refer to "4.5.7 Deleting a RAID Volume".
Deleting a disk group	Refer to "4.5.8 Deleting a Disk Group".
Clearing RAID configuration information	Refer to "4.5.9 Clearing RAID Configuration Information".
Configuring a pass-through disk	Refer to "4.5.10 Configuring a Pass-Through Disk".
Enabling the caching function	Refer to "4.5.11 Enabling the Caching Function".

4.5.1 Setting the Mode of a Port

Abstract

The ports of a SmartROC 3100 RAID controller card, namely, the ports connected with a disk backplane and disk cables, support three modes: RAID, HBA and Mixed. Before adding the disk corresponding to a port to a RAID logical volume, you must check the port mode.
The SmartROC 3100 RAID controller card supports setting the port mode in the following two ways:

- Setting the mode of all ports in batches
- Setting the mode of a port separately

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.

- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

- Setting the Mode of All Ports in Batches
  1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-60.

Figure 4-60 Configure Controller Settings Screen

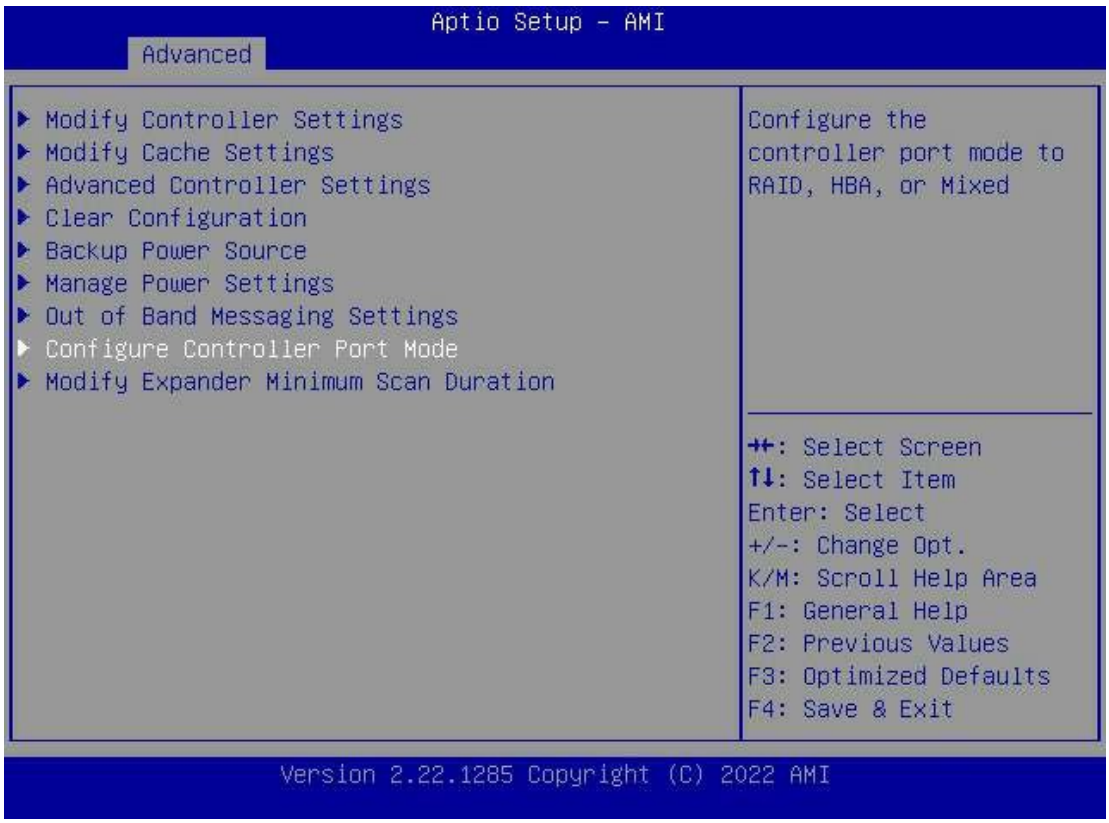

  2. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**. The **Configure Controller Port Mode** screen is displayed, see Figure 4-61.

Figure 4-61 Configure Controller Port Mode Screen

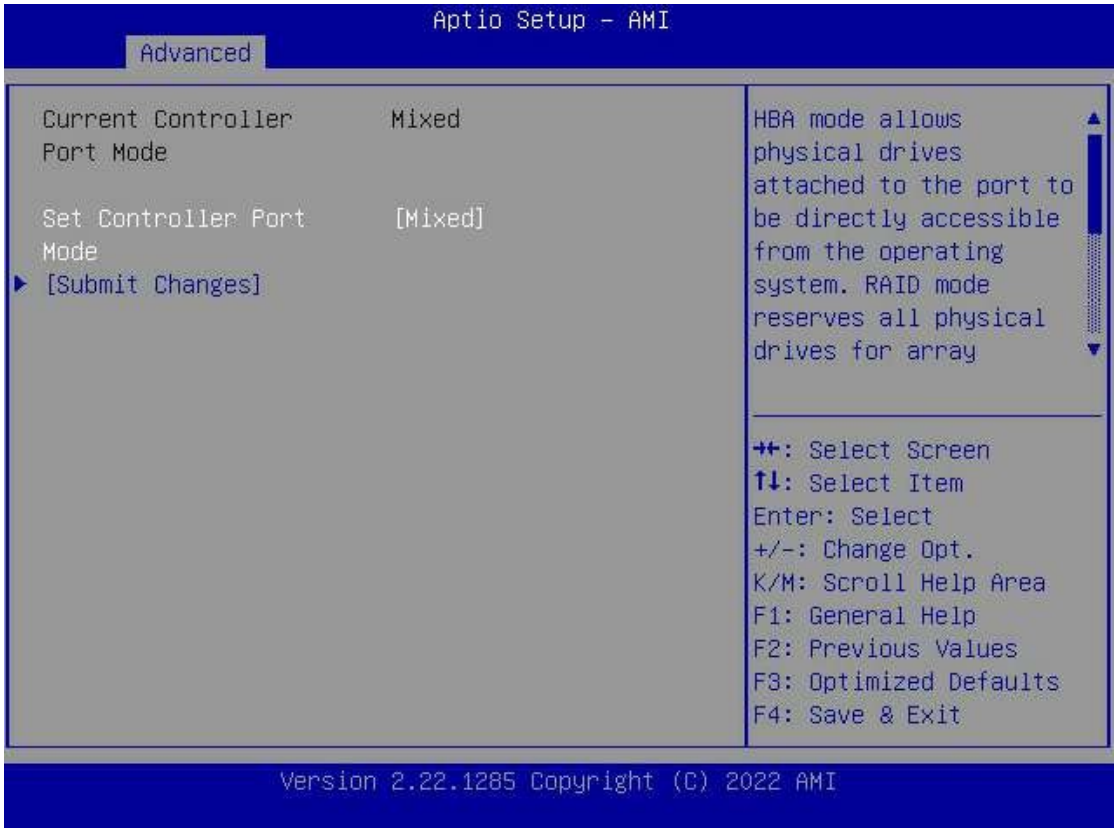

3. Use the arrow keys to select **Set Controller Port Mode**, and then press **Enter**. The **Set Controller Port Mode** dialog box is displayed, see Figure 4-62.

Figure 4-62 Set Controller Port Mode Dialog Box

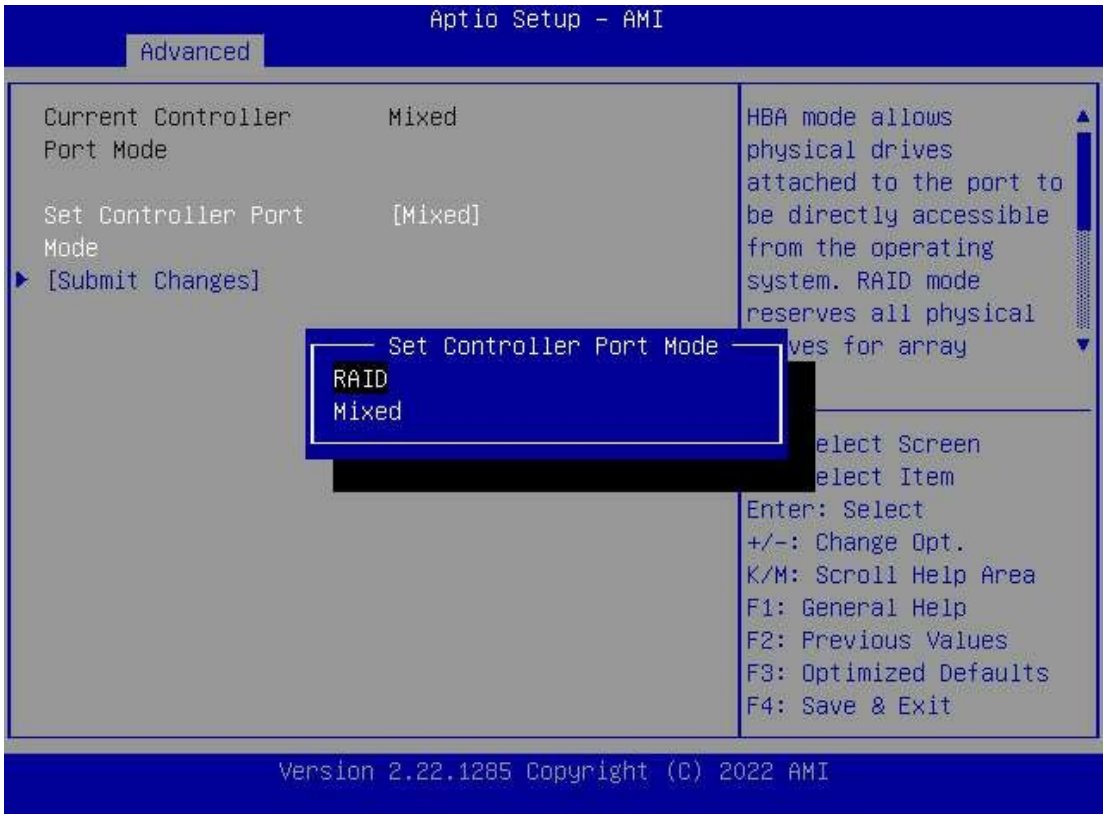

4. Use the arrow keys to select the port mode to be set, and then press **Enter**, see Figure 4-63.

Figure 4-63 Selecting a Port Mode

5. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 4-64.

Figure 4-64 Successful Setting

6. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

- Setting the Mode of a Port Separately

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-65.

Figure 4-65 Configure Controller Settings Screen

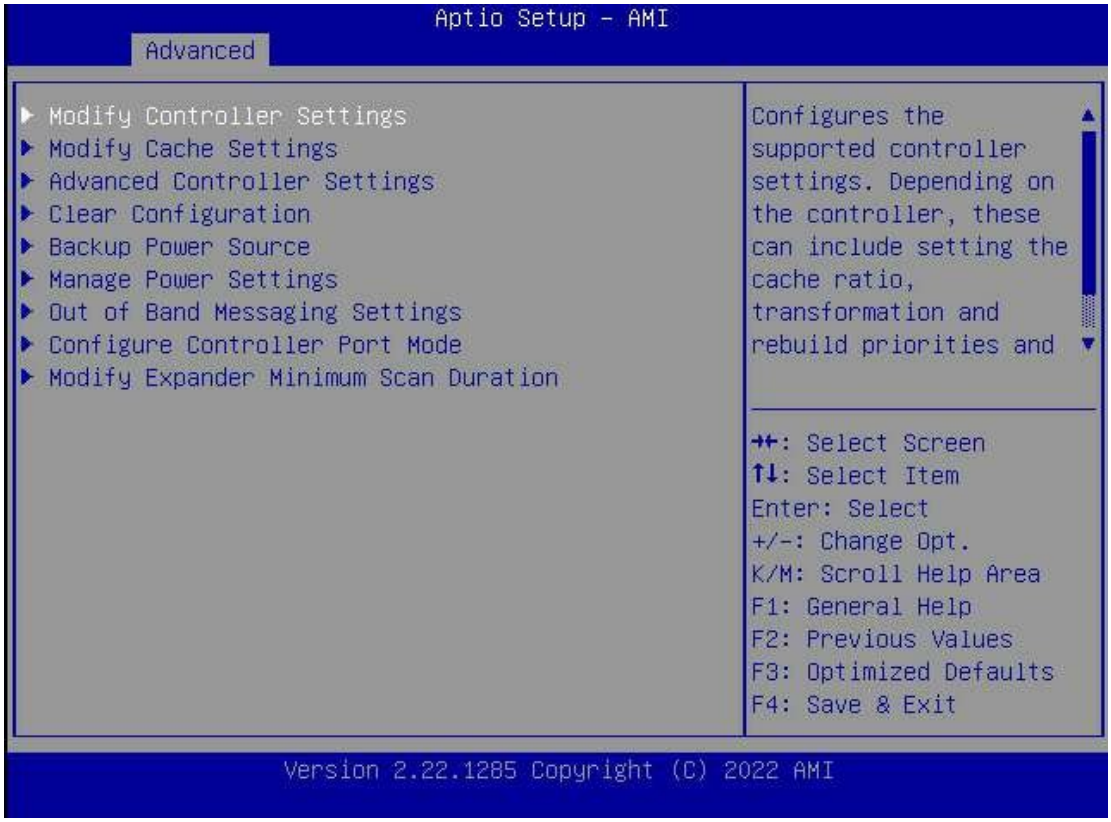

2. Use the arrow keys to select **Modify Controller Settings**, and then press **Enter**. The **Modify Controller Settings** screen is displayed, see Figure 4-66.

Figure 4-66 Modify Controller Settings Screen

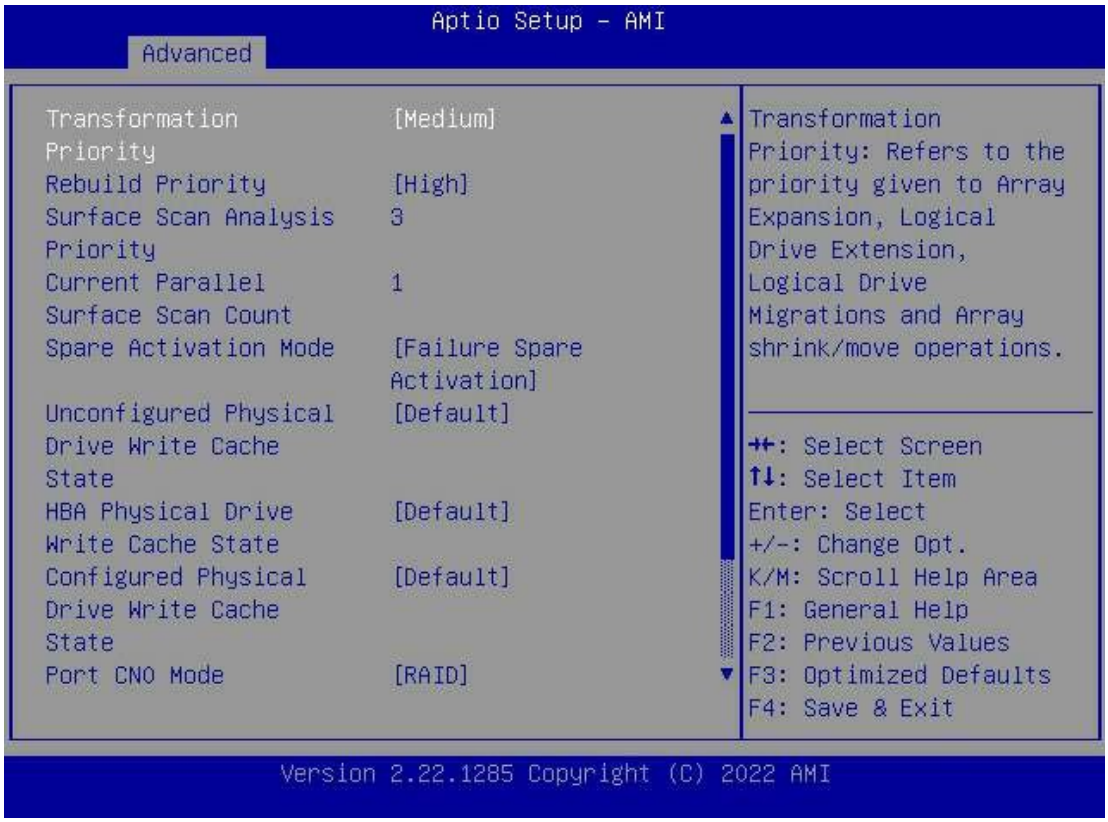

3. Use the arrow keys to select **Port CN0 Mode**, and then press **Enter**. The **Port CN0 Mode** dialog box is displayed, see Figure 4-67.

Figure 4-67 Port CNO Mode Dialog Box

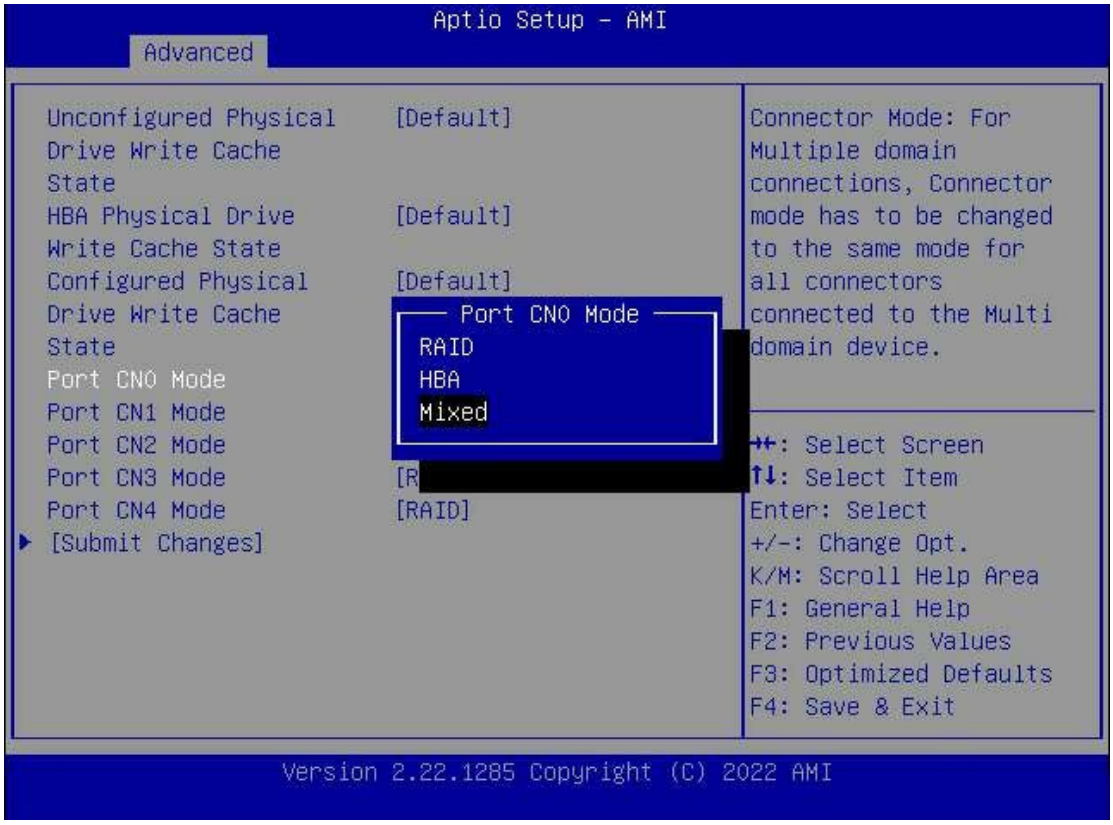

4. Use the arrow keys to select the port mode to be set, and then press **Enter**, see Figure 4-68.

Figure 4-68 Setting the Port Mode

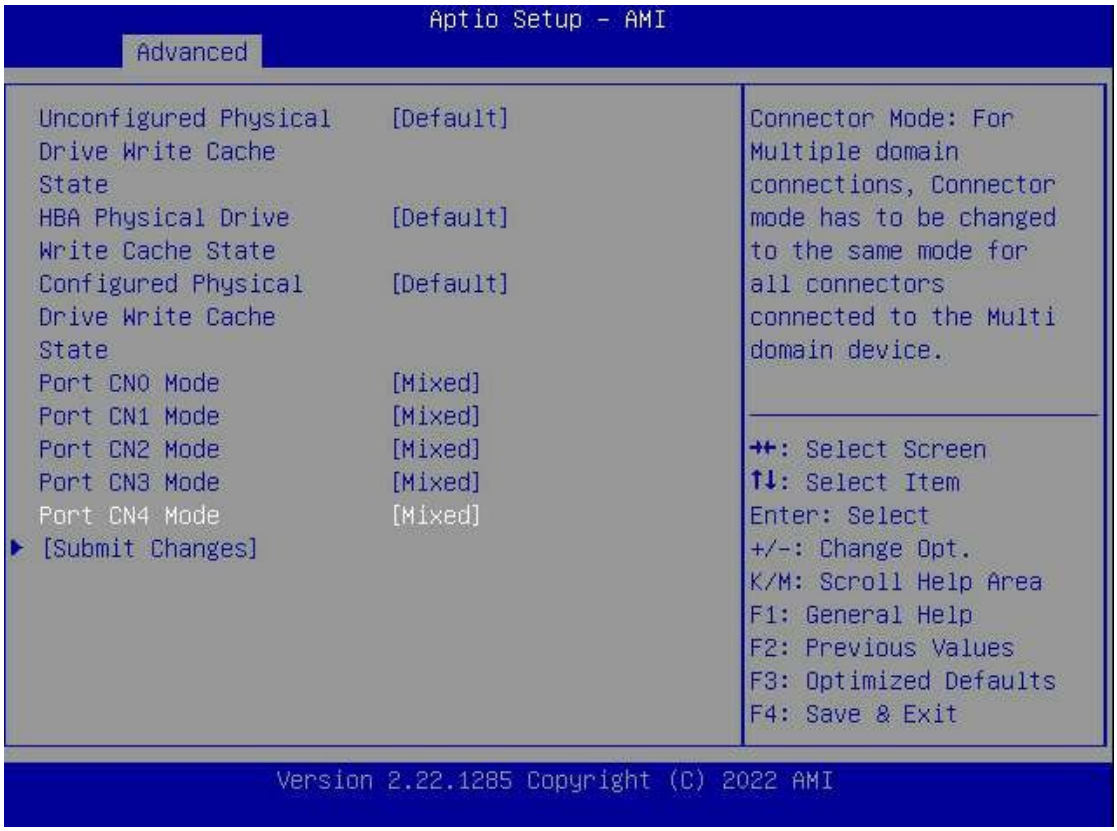

5. Repeat Step 3 and Step 4 to set the mode of another port, see Figure 4-69.

Figure 4-69 Setting the Port Mode

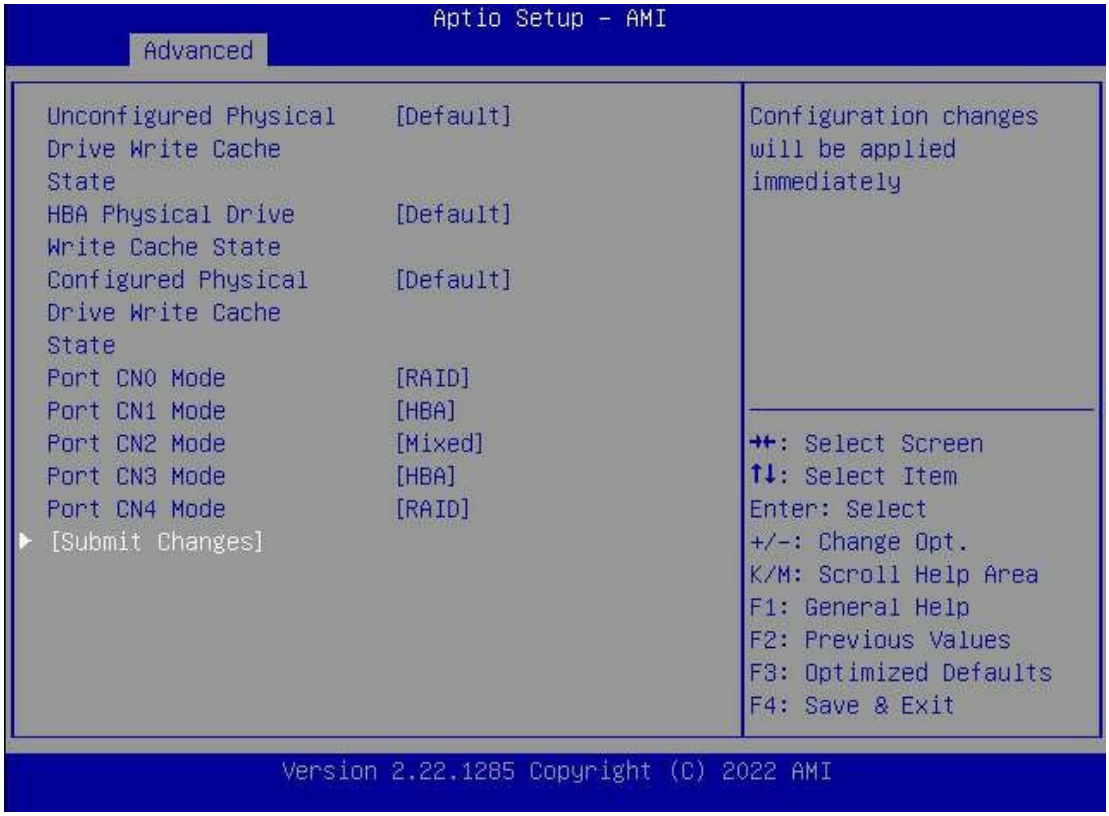

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 4-70.

Figure 4-70 Port Mode Set Successfully

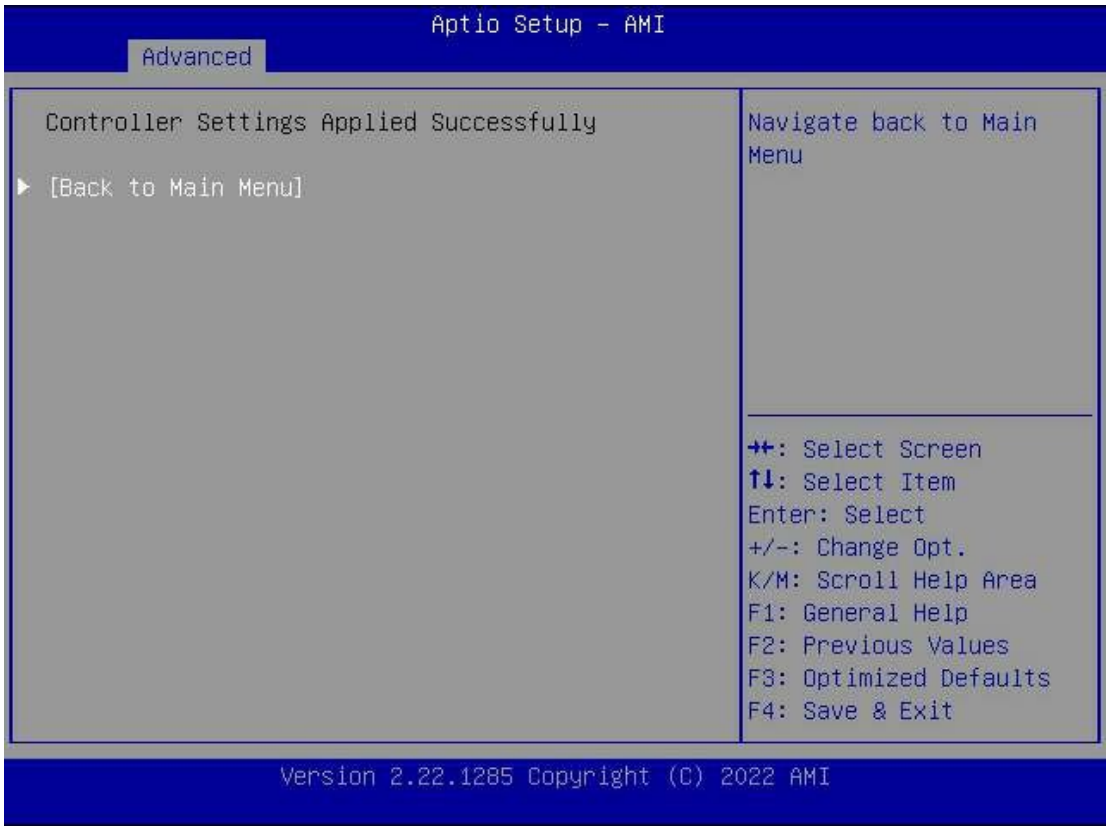

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.2 Locating a Disk

Abstract

After the indicator of a disk is lit, you can locate the disk so that you can easily replace or maintain it. You can locate a physical disk or multiple disks in a disk group.

Steps

- Locating a Single Physical Disk
  1. On the controller management screen, use the arrow keys to select **Disk Utilities**, and then press **Enter**. The screen for the list of physical disks mounted on the RAID controller card is displayed, see Figure 4-71.

Figure 4-71 Screen for the List of Disks Mounted on the RAID Controller Card

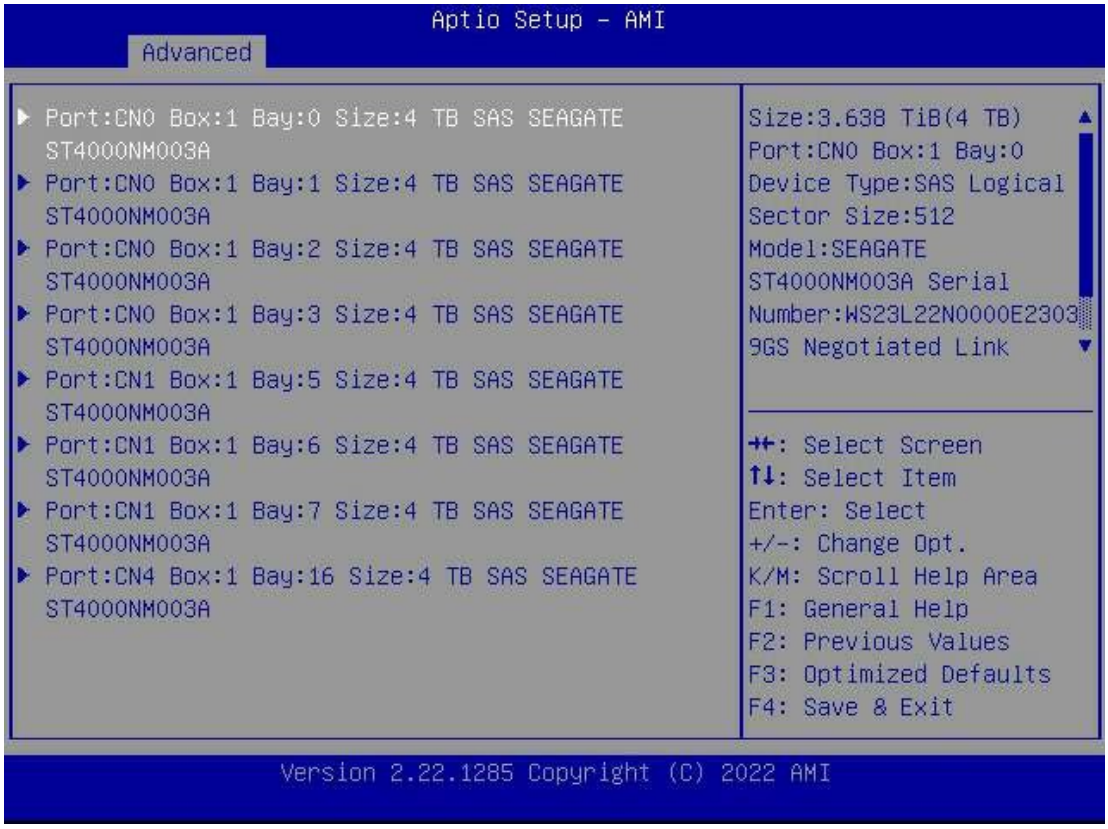

2. Use the arrow keys to select the disk to be located, and then press **Enter**. The screen for setting a disk is displayed, see Figure 4-72.

Figure 4-72 Setting a Disk

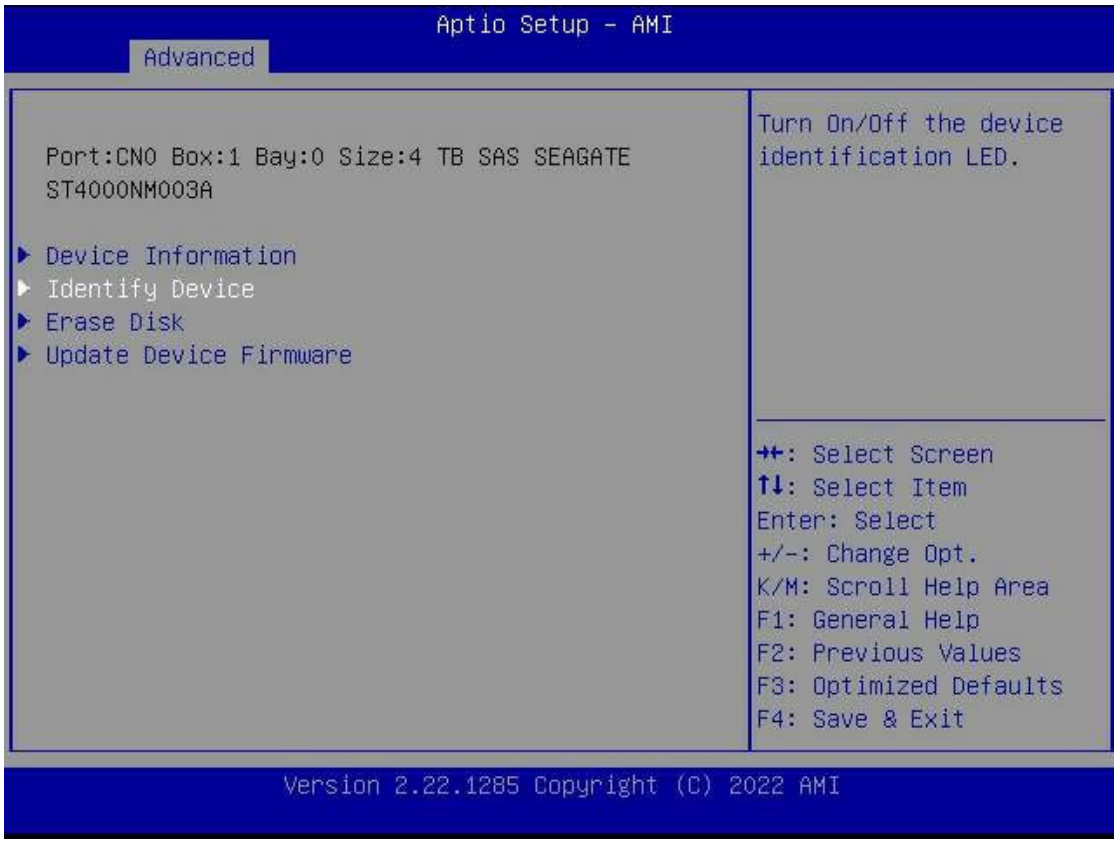

3. Use the arrow keys to select **Identify Device**, and then press **Enter**. The **Identify Device** screen is displayed, see Figure 4-73.

Figure 4-73 Identify Device Screen

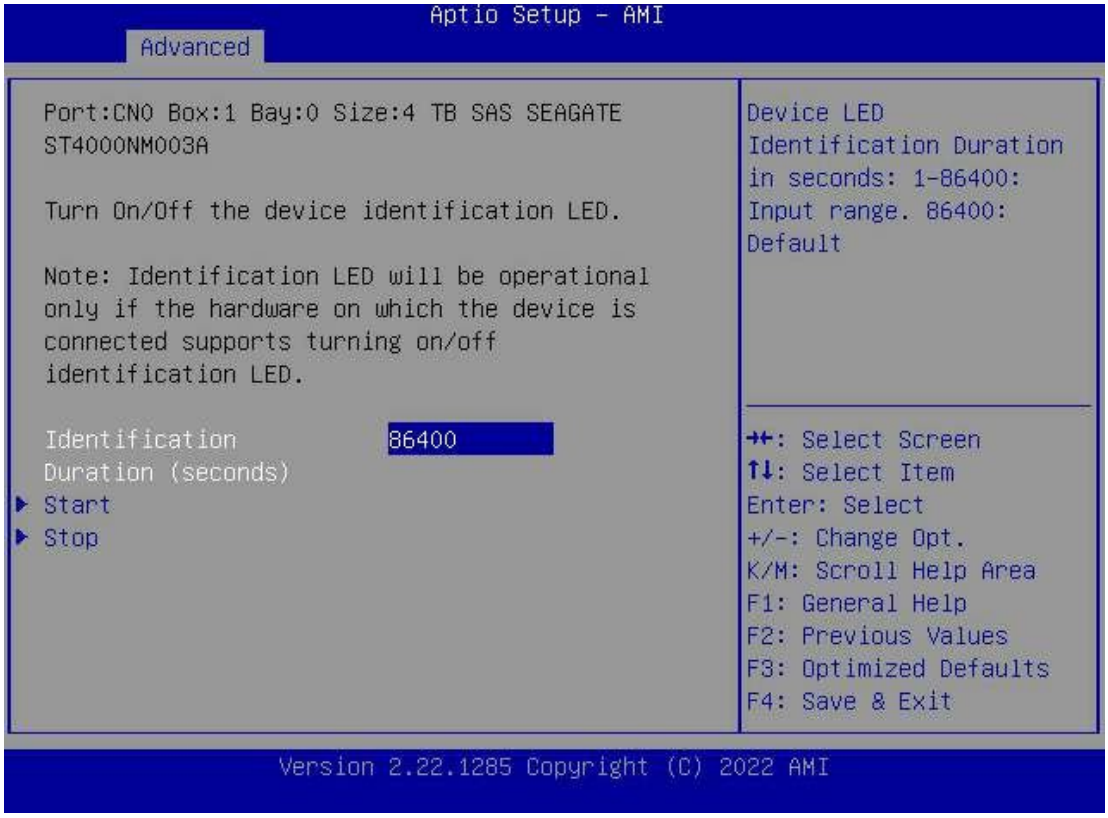

4. Next to **Identification Duration (seconds)**, enter the time of the lighting delay (unit: seconds), select **Yes**, and press **Enter** for confirmation, see Figure 4-74.

Figure 4-74 Lighting Delay

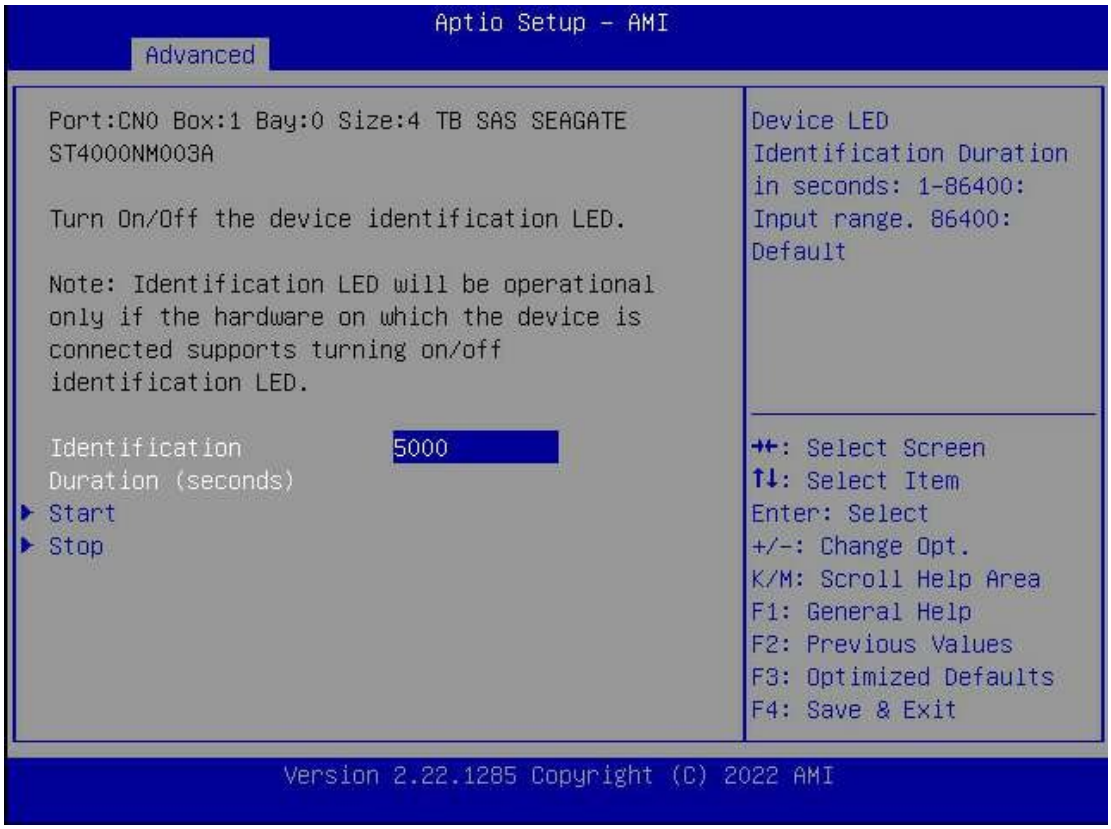

5. Use the arrow keys to select **Start**, and then press **Enter**. The red indicator of the corresponding disk is lit and continuously flashes, see Figure 4-75.

Figure 4-75 Starting Locating a Disk

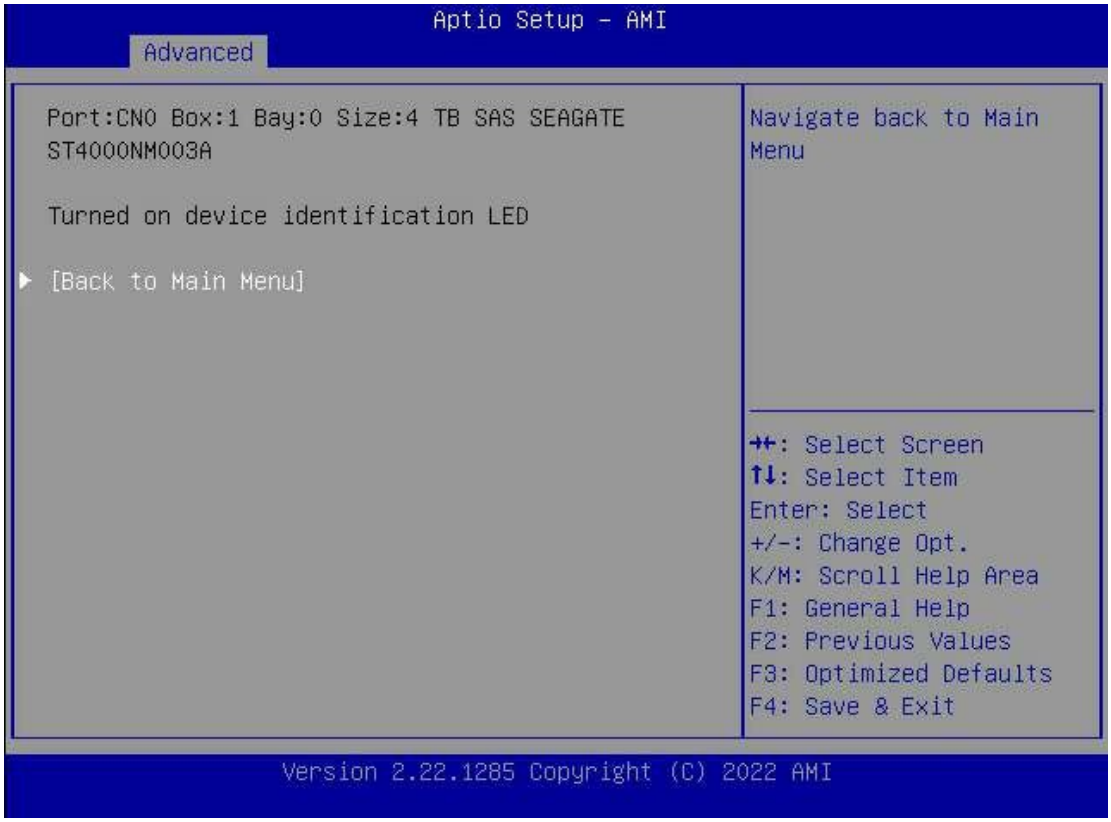

Note

The indicator flashing duration is the configured lighting delay. After the lighting delay is reached, the default value 86400 s (24 hours) is restored, and the indicator goes out.

6. (Optional) To go out the disk indicator and end the locating, press **Esc** to return to the locating screen, use the arrow keys to select **Stop**, and press **Enter**, see Figure 4-76.

Figure 4-76 Ending Disk Locating

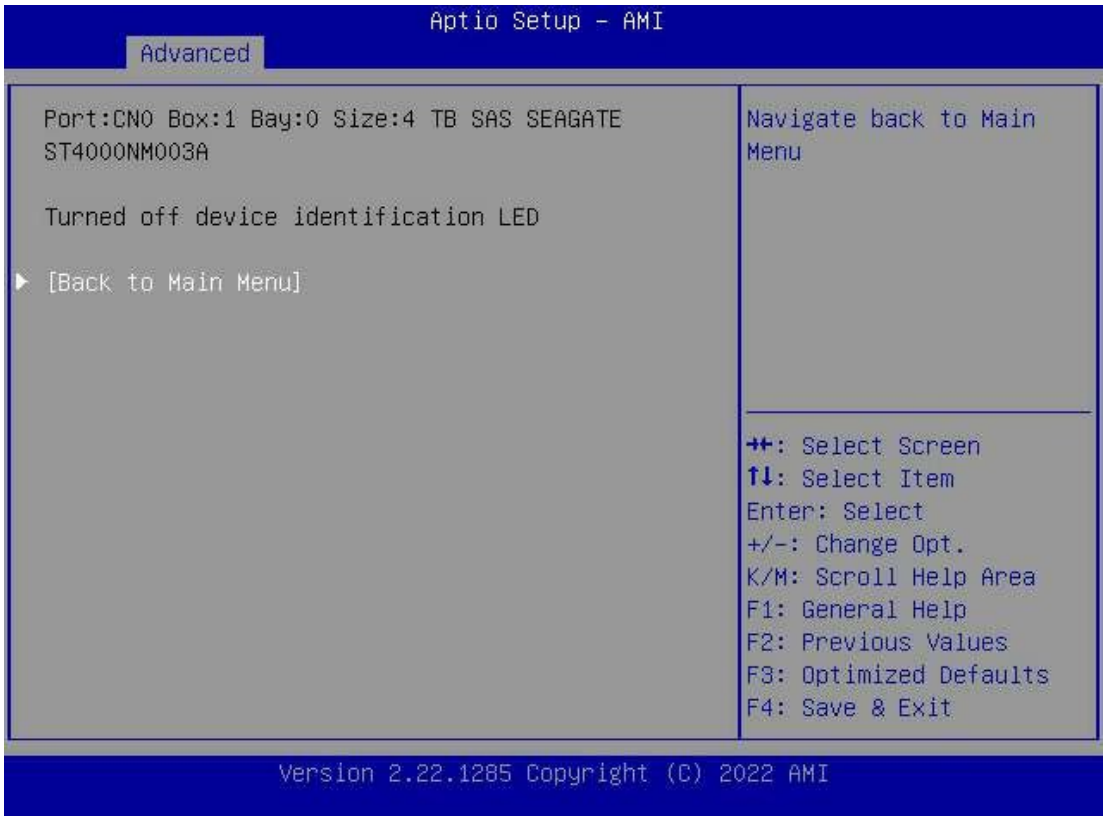

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

- Locating Multiple Disks in a Disk Group

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-77.

Figure 4-77 Array Configuration Screen

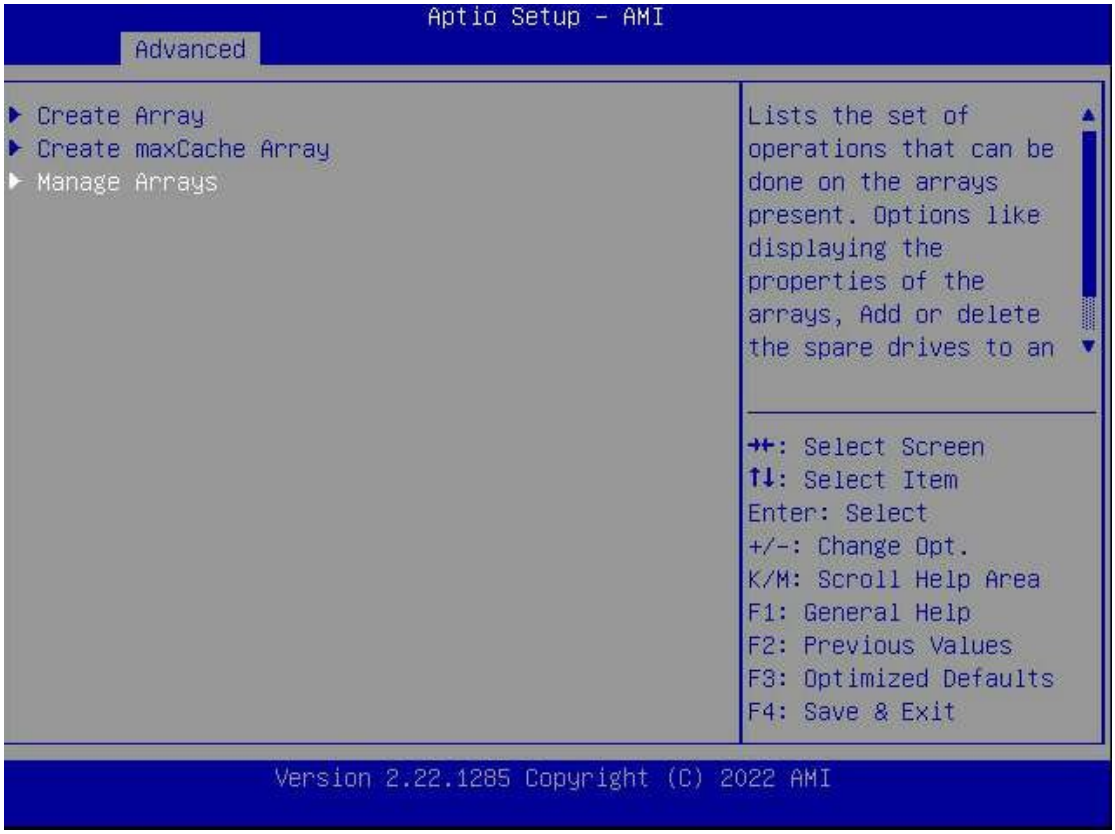

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-78.

Figure 4-78 Managing Arrays Screen

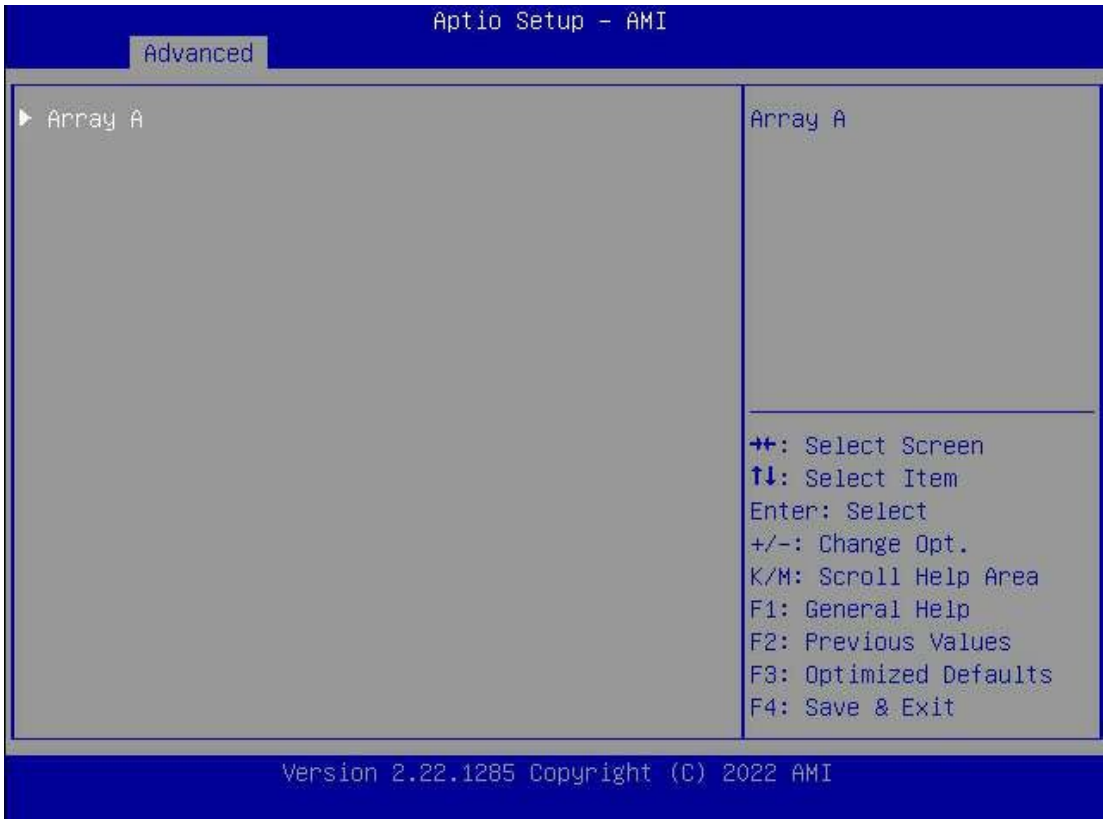

3. Use the arrow keys to select the array in which the disk to be located is placed, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-79.

Figure 4-79 Managing the Selected Array

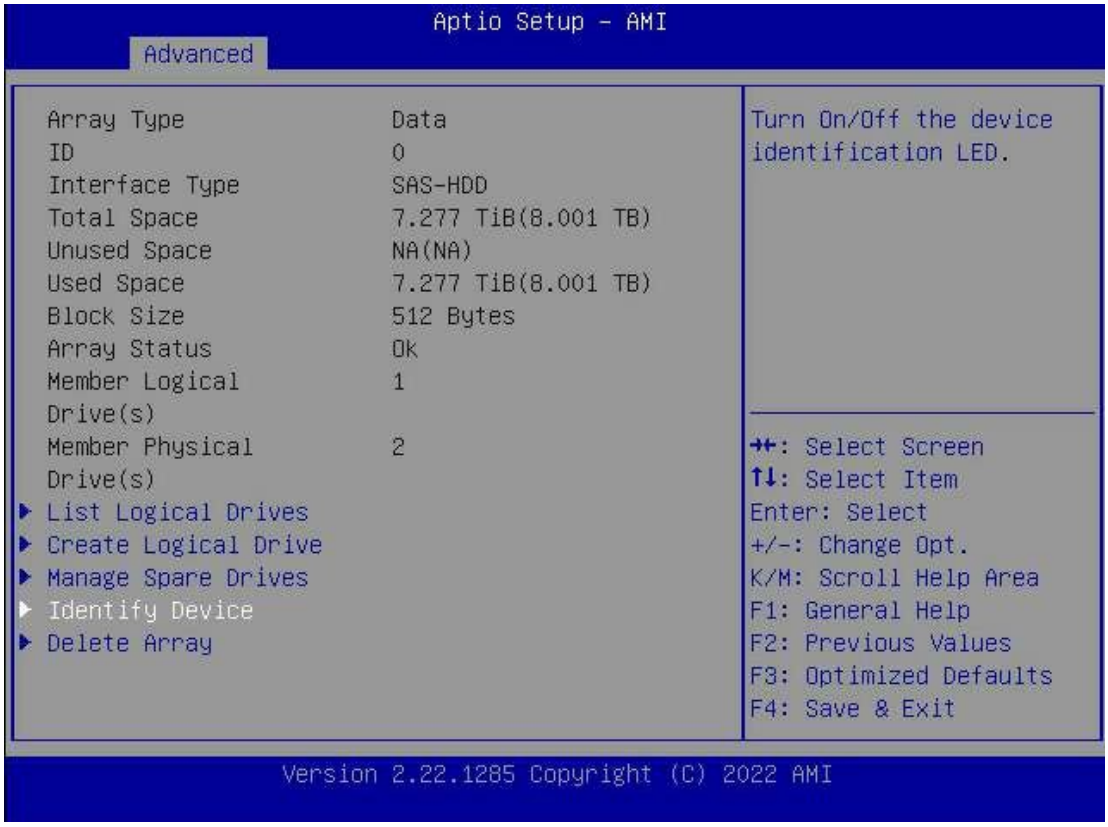

4. Use the arrow keys to select **Identify Device**, and then press **Enter**. The **Identify Device** screen is displayed, see Figure 4-80.

Figure 4-80 Identify Device Screen

5. Next to **Identification Duration (seconds)**, enter the time of the lighting delay (unit: seconds), and then press **Enter** for confirmation, see Figure 4-81.

Figure 4-81 Lighting Delay

6. Use the arrow keys to select **Start**, and then press **Enter**. The red indicators of all the disks that belong to the array are lit and continuously flash, see Figure 4-82.

Figure 4-82 Starting Locating a Disk

Note

The indicators of the hot spare disks belonging to the array are also lit and flash continuously at the same time. The indicator flashing duration is the configured lighting delay. After the lighting delay is reached, the default value 86400 s (24 hours) is restored, and the indicator goes out.

7. (Optional) To go out the disk indicator and end the locating, press **Esc** to return to the locating screen, and use the arrow keys to select **Stop**. The disk indicator goes out and the disk locating ends, see Figure 4-83.

Figure 4-83 Ending Disk Locating

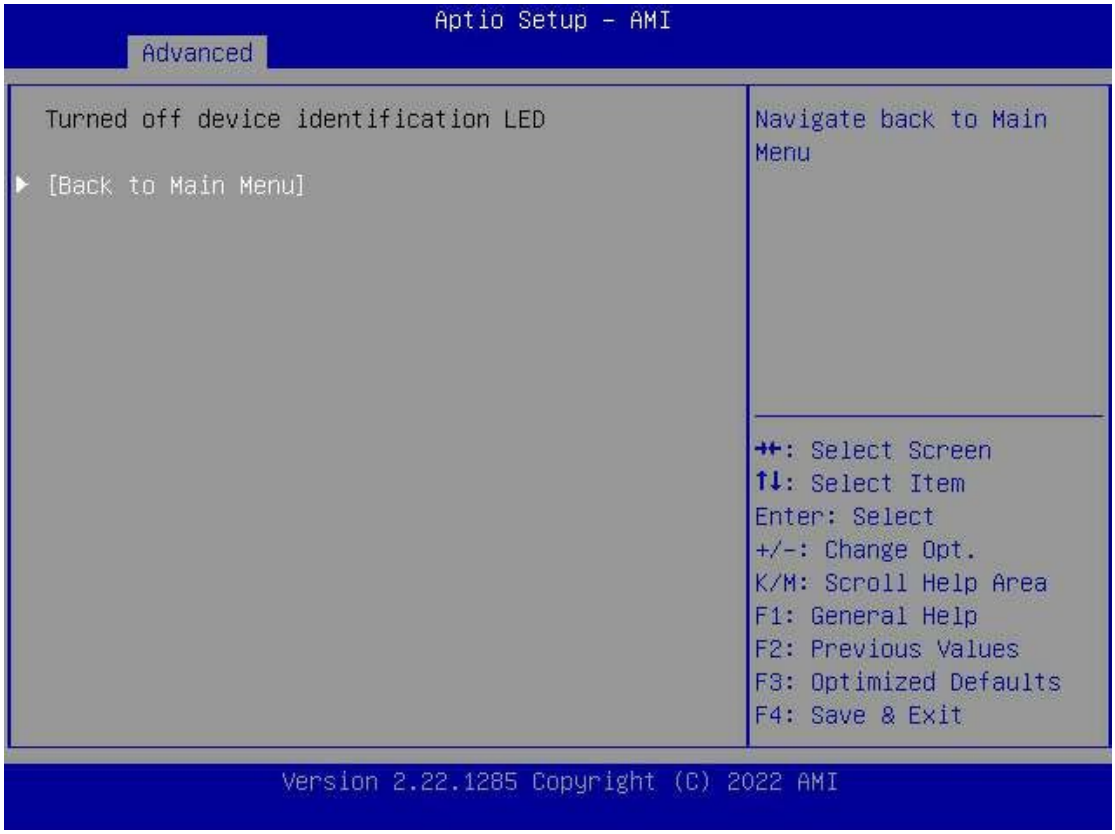

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.3 Creating a Hot Spare Disk

Abstract

A hot spare disk improves the data security of a RAID array. For a description of the hot spare disk types supported by a SmartROC 3100 RAID controller card, refer to Table 4-10.

Table 4-10 Hot Spare Disk Types

Type	Description
Dedicated	● This type of hot spare disks is exclusive to the specified one or more disk groups of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a dedicated hot spare disk temporarily takes over the faulty disk.
Auto Replace	● This type of hot spare disks provides the hot standby function for a disk group of a RAID controller card. One or more hot spare disks can be created for each disk group. ● When a disk in a disk group is faulty, a hot spare disk of this type automatically replaces the faulty disk.

Prerequisite

There are sufficient idle disks on the server.

Context

When creating a hot spare disk, pay attention to the following points:

- Multiple hot spare disks can be created for a disk group, but only one type of hot spare disk can be set at a time. That is, either **Dedicated** or **Auto Replace** is specified.
- An idle disk can be set as a hot spare disk. The disk that has been used to create a RAID volume cannot be set as a hot spare disk.
- The hot spare disk must be of the same type as that of any member disk in the corresponding disk group. That is, all of them are SATA disks or SAS disks, and the hot spare disk's capacity must not be less than the maximum capacity of the member disks.
- Disk groups at all levels except RAID 0 support hot spare disks.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-84.

Figure 4-84 Array Configuration Screen

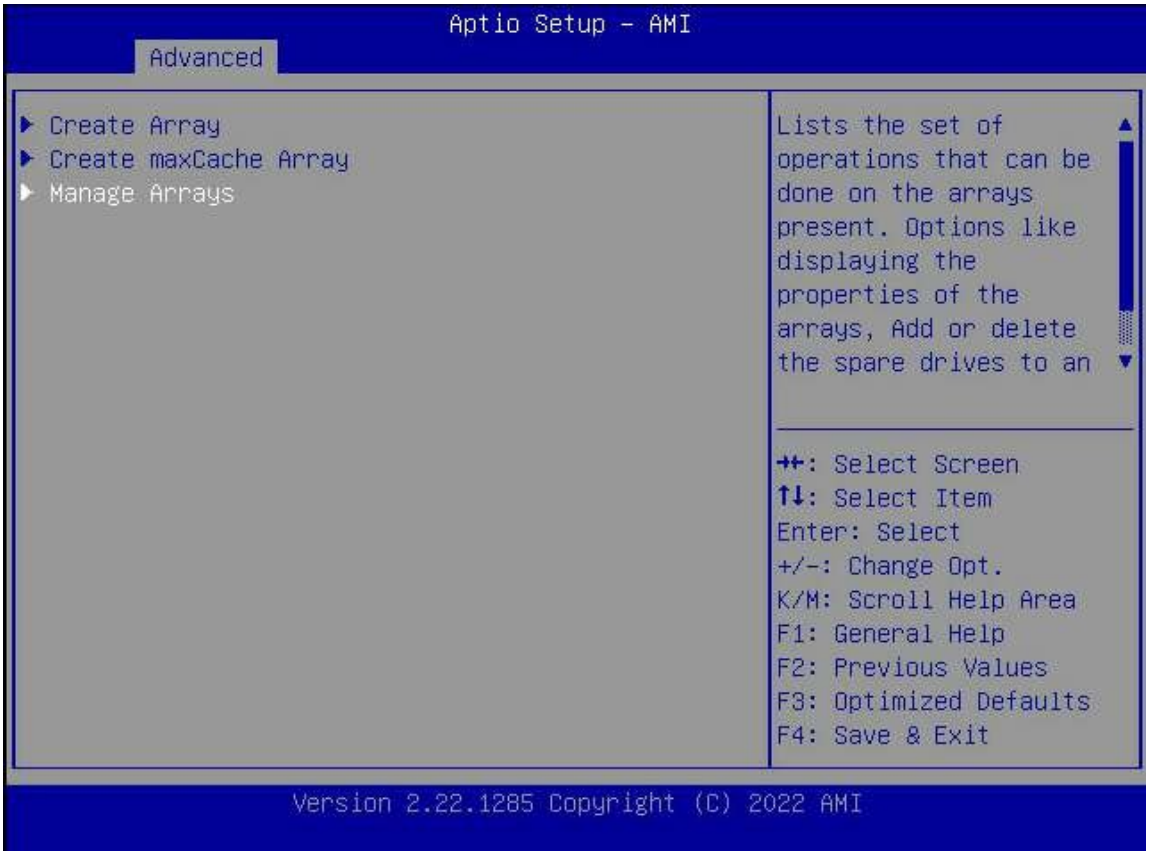

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-85.

Figure 4-85 Managing Arrays Screen

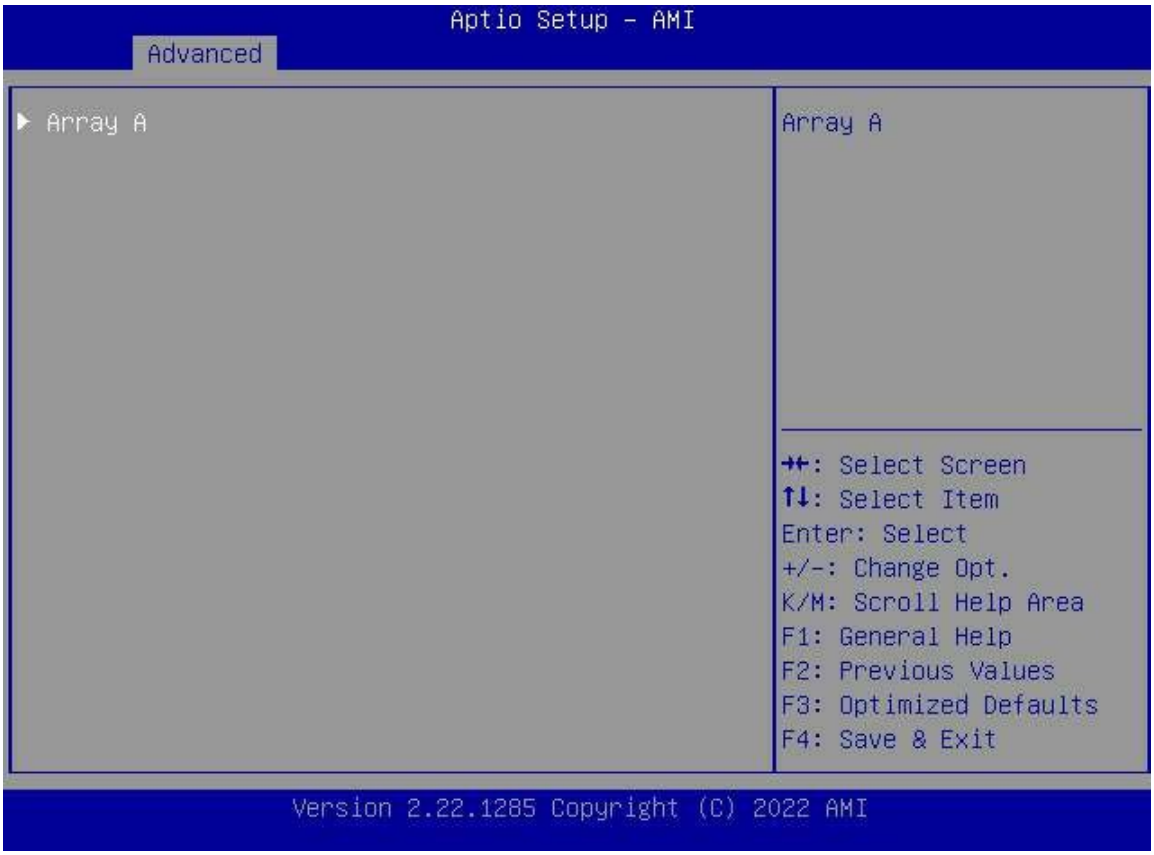

3. Use the arrow keys to select the array for which you need to create the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-86.

Figure 4-86 Managing the Selected Array

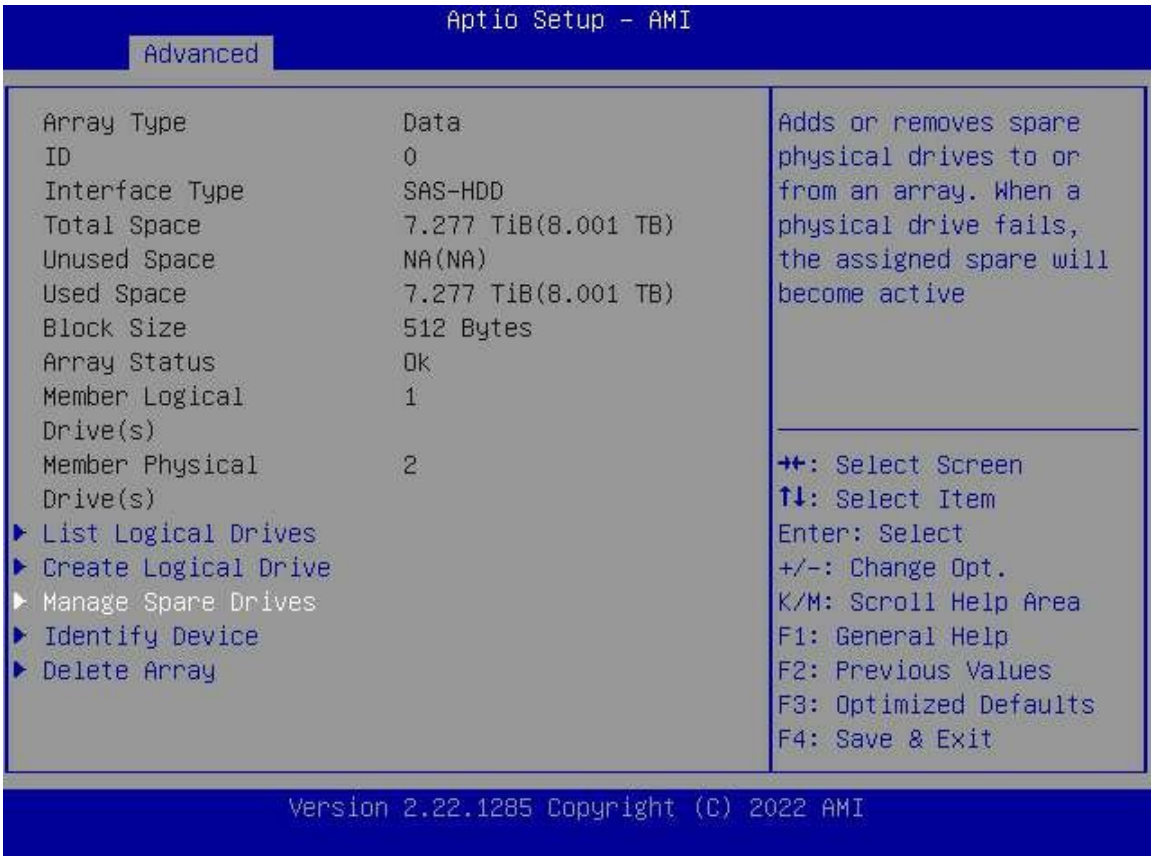

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 4-87.

Figure 4-87 Manage Spare Drives Screen

5. In accordance with your actual conditions, use the arrow keys to select the type of the hot spare disk to be created, and then press **Enter**. The screen for selecting a hot spare disk is displayed, see Figure 4-88.

Figure 4-88 Selecting a Hot Spare Disk

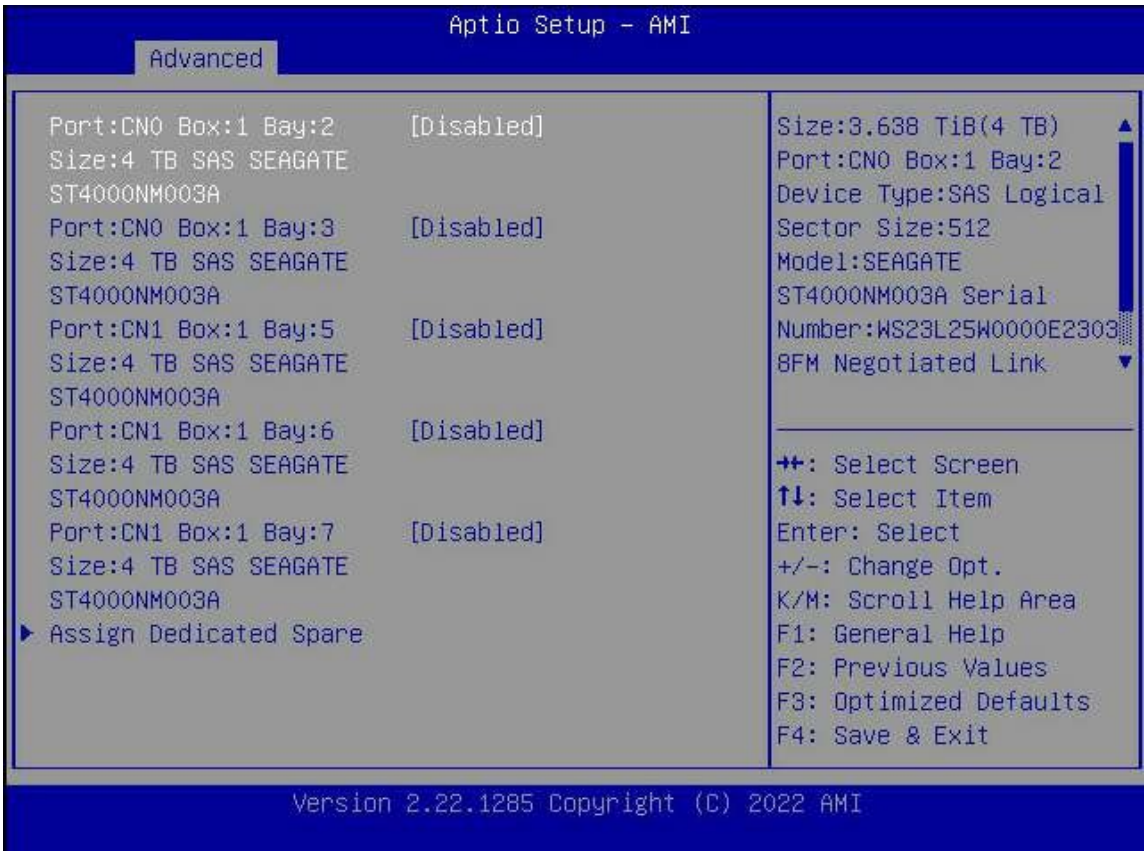

6. Use the arrow keys to select the disk to be set as a hot spare disk, press **Enter**, and then set the status of the disk to **Enabled**, see Figure 4-89.

Figure 4-89 Setting Disk Status

7. Use the arrow keys to select **Assign Dedicated Spare**, and then press **Enter**. The hot spare disk is successfully created, see Figure 4-90.

Figure 4-90 Hot Spare Disk Created Successfully

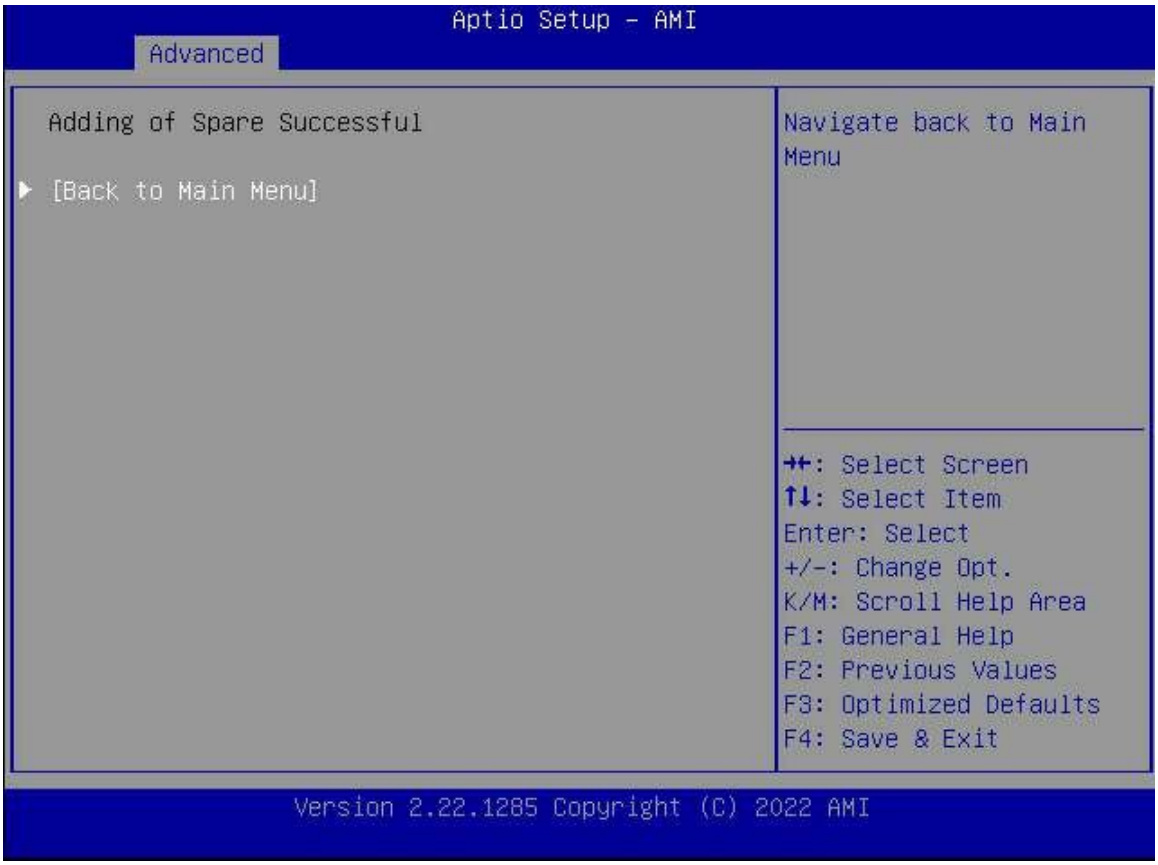

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.4 Changing a Hot Spare Disk

Abstract

A SmartROC 3100 RAID controller card supports modifying the type of a hot spare disk, namely, allowing type change between **Dedicated** and **Auto Replace**.

Only one type of hot spare disk can be set at a time. That is, **Dedicated** and **Auto Replace** cannot be specified at the same time. This procedure uses changing a hot spare disk of the **Dedicated** type to that of the **Auto Replace** type as an example to describe how to perform a type change.

Note

A hot spare disk of the **Auto Replace** type can be changed to that of the **Dedicated** type by referring to this procedure.

Prerequisite

A hot spare disk is already set to the **Dedicated** type. For details, refer to "4.5.3 Creating a Hot Spare Disk".

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-91.

Figure 4-91 Array Configuration Screen

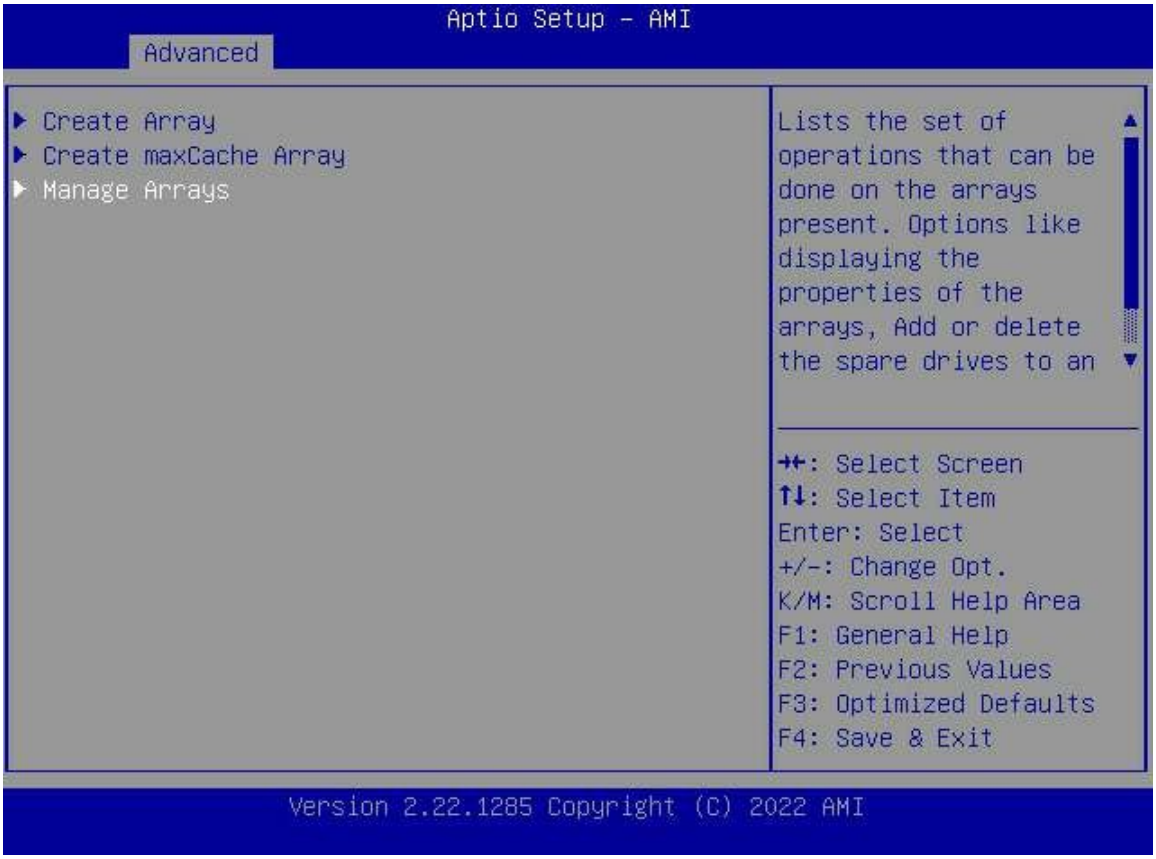

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-92.

Figure 4-92 Managing Arrays Screen

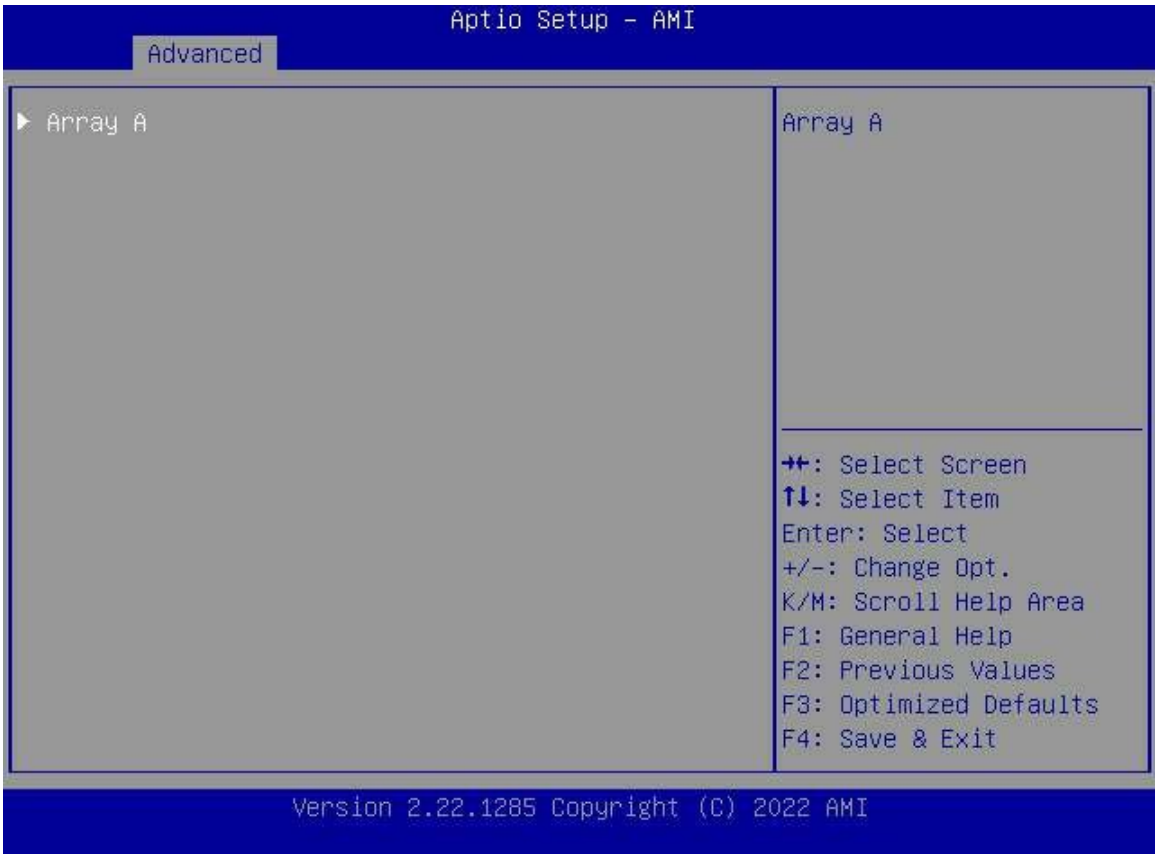

3. Use the arrow keys to select the array for which you need to modify the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-93.

Figure 4-93 Managing the Selected Array

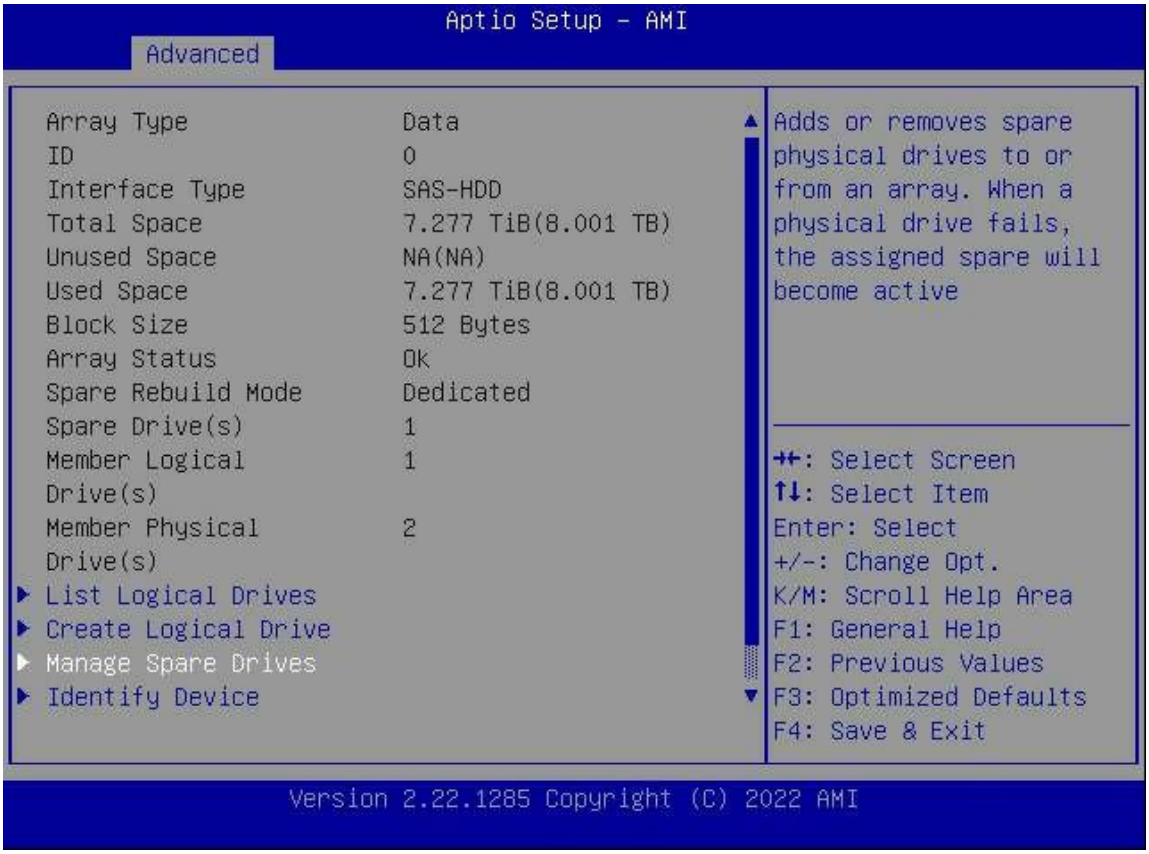

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 4-94.

Figure 4-94 Manage Spare Drives Screen

5. Use the arrow keys to select **Change Spare type to AutoReplace**, and then press **Enter**.
   The screen for hot spare disk type modification is displayed, see Figure 4-95.

Figure 4-95 Screen for Hot Spare Disk Type Modification

6. Select **Submit Changes**, and then press **Enter**. The type of hot spare disk is changed successfully, see Figure 4-96.

Figure 4-96 Hot Spare Disk Type Changed Successfully

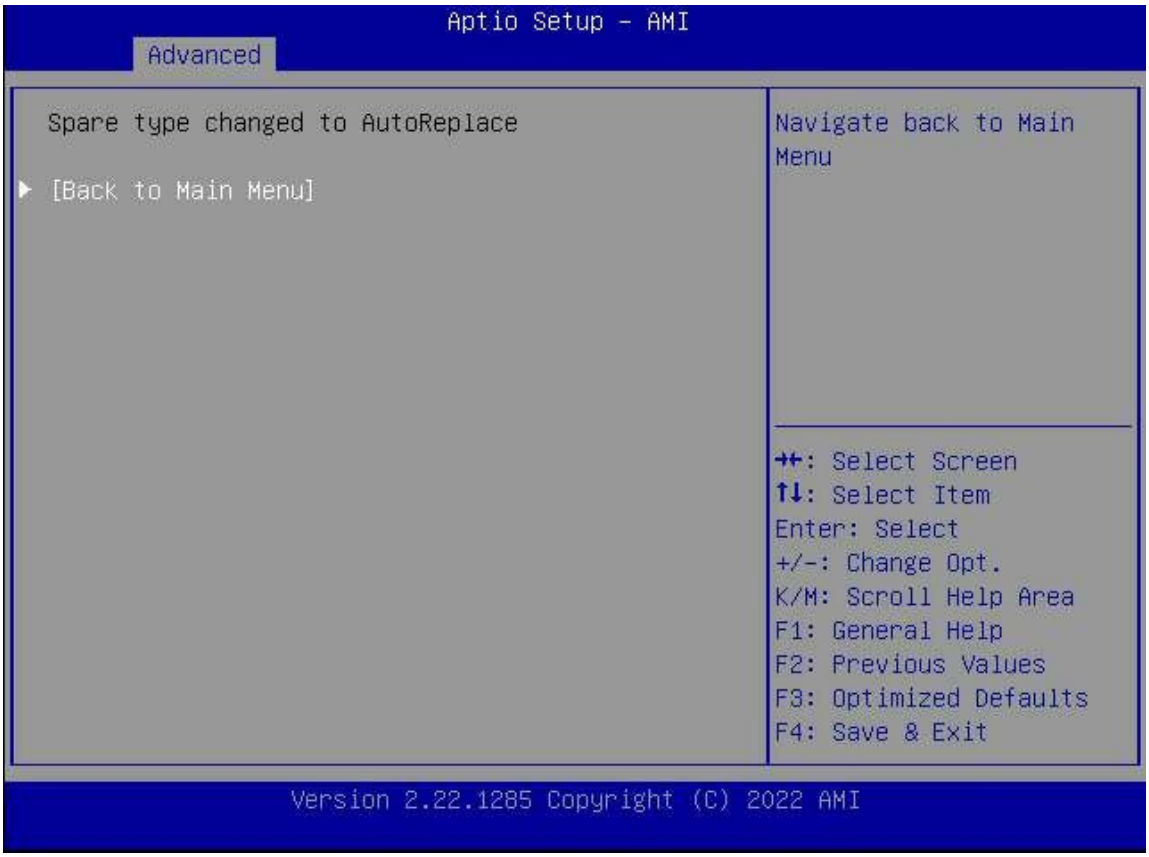

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.5 Deleting a Hot Spare Disk

Abstract

When the number of disks of a server cannot meet the requirements, you can delete an existing hot spare disk and restore it to a common disk.

Prerequisite

A hot spare disk is created successfully. For details, refer to "4.5.3 Creating a Hot Spare Disk".

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-97.

Figure 4-97 Array Configuration Screen

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-98.

Figure 4-98 Managing Arrays Screen

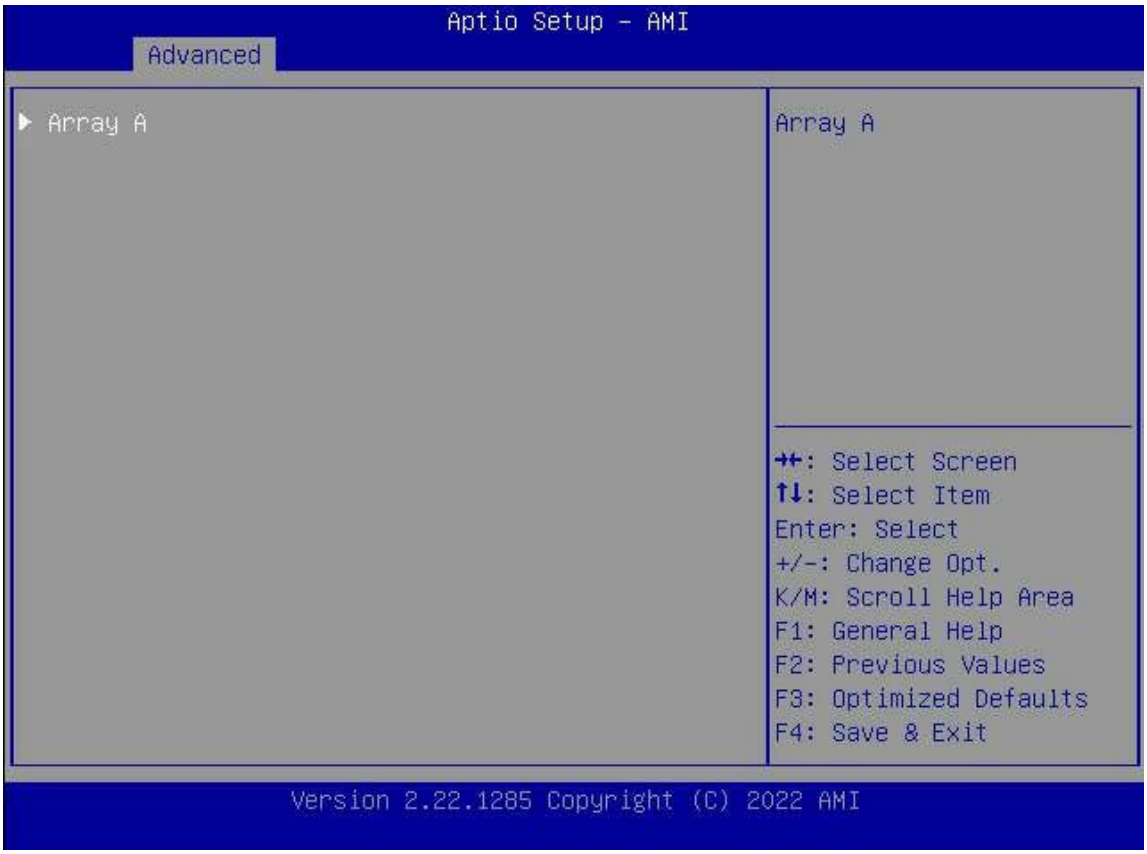

3. Use the arrow keys to select the array for which you need to delete the hot spare disk, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-99.

Figure 4-99 Managing the Selected Array

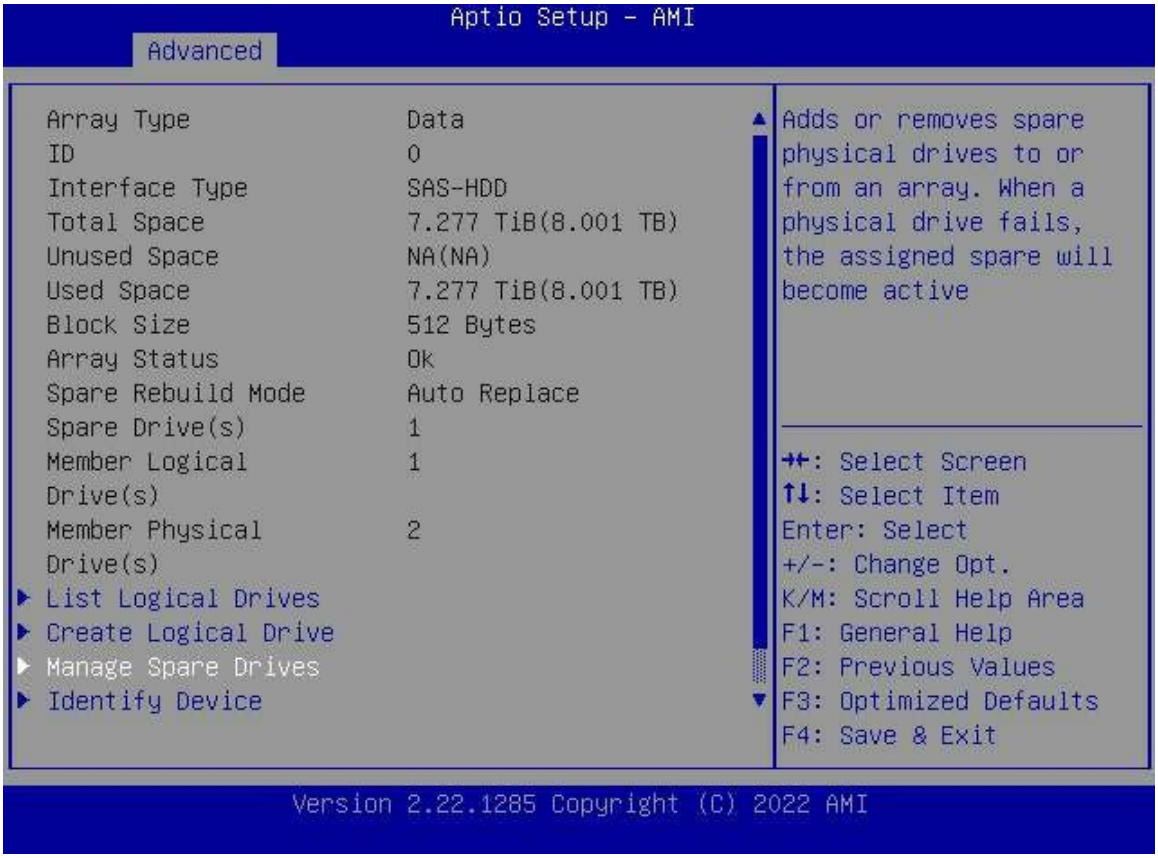

4. Use the arrow keys to select **Manage Spare Drives**, and then press **Enter**. The **Manage Spare Drives** screen is displayed, see Figure 4-100.

Figure 4-100 Manage Spare Drives Screen

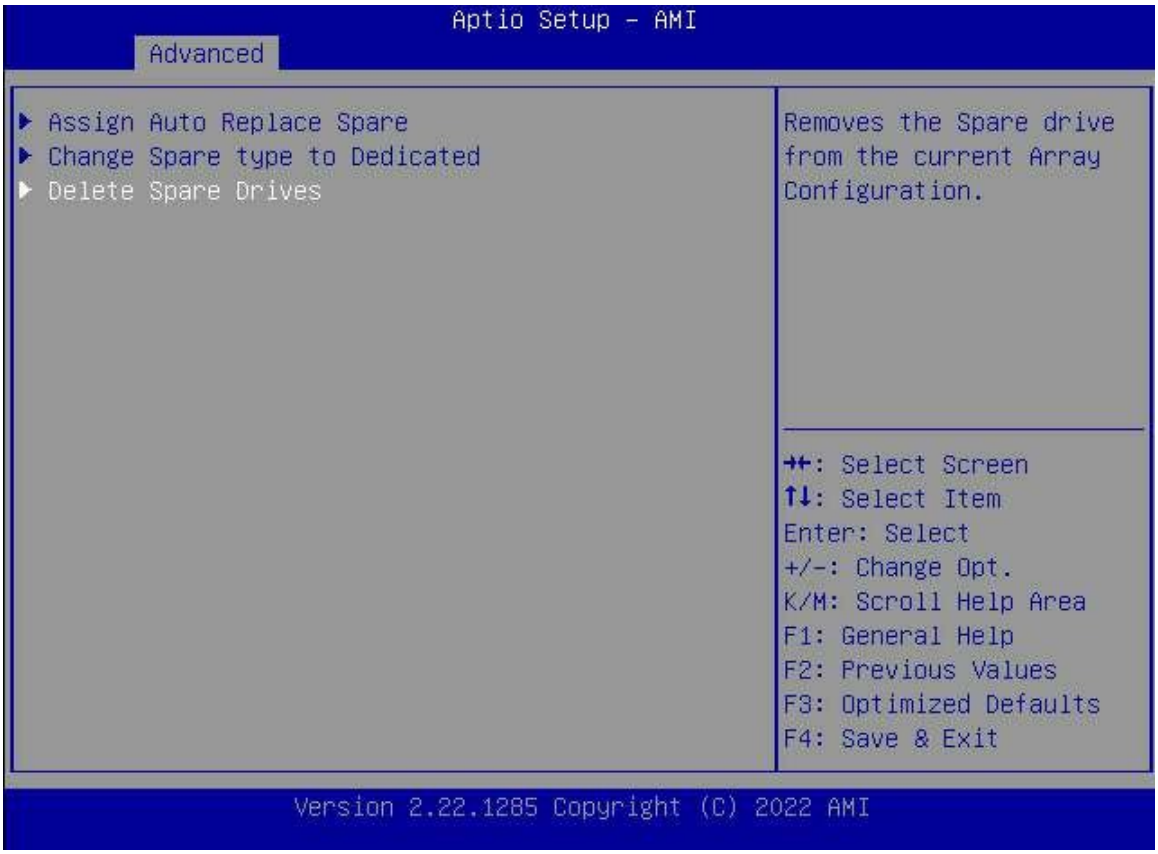

5. Use the arrow keys to select **Delete Spare Drives**, and then press **Enter**. The screen for selecting a hot spare disk is displayed, see Figure 4-101.

Figure 4-101 Selecting a Hot Spare Disk

6. Use the arrow keys to select the hot spare disk to be deleted, press **Enter**, and then set the status of the disk to **Enabled**, see Figure 4-102.

Figure 4-102 Setting Disk Status

7. Use the arrow keys to select **Delete Spare Drives** and press **Enter**. The hot spare disk is deleted successfully, see Figure 4-103.

Figure 4-103 Hot Spare Disk Deleted Successfully

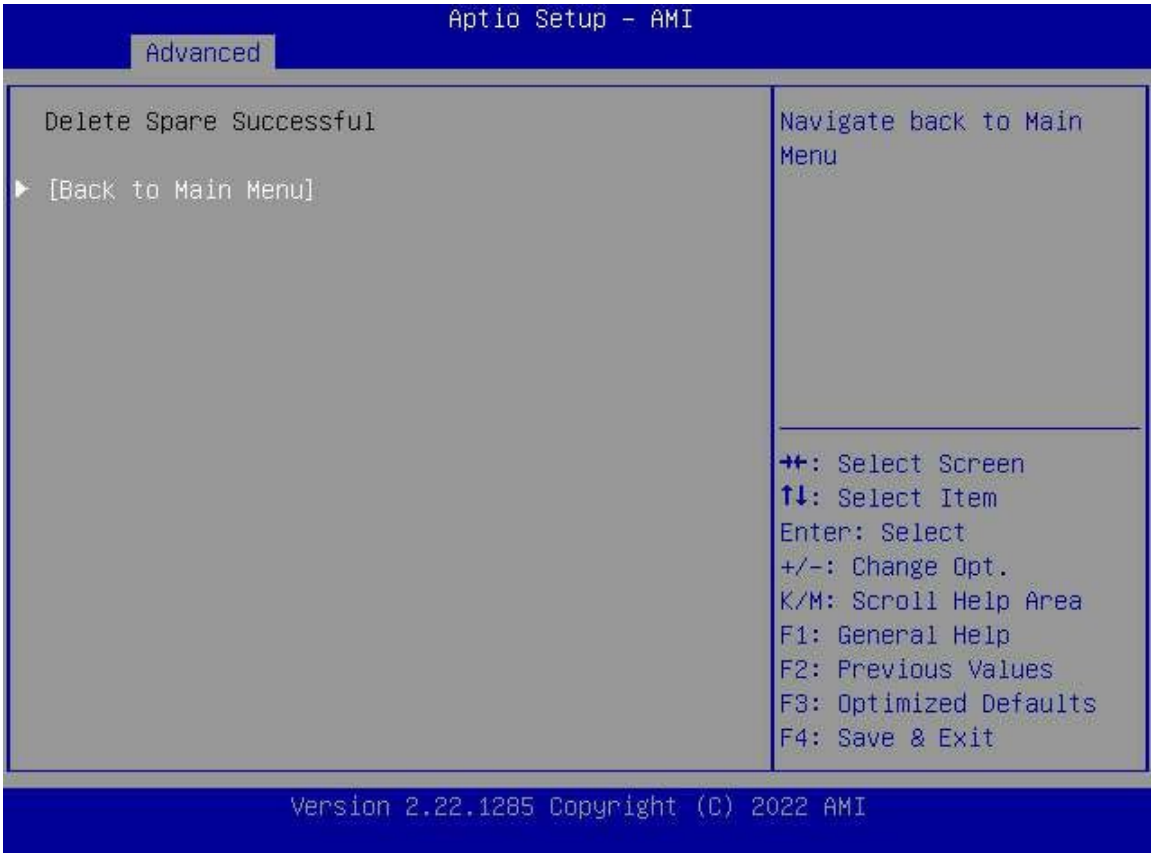

8. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.6 Configuring the Performance or Power Mode

Abstract

This procedure describes how to configure a power mode for a SmartROC 3100 RAID controller card. For a description of the power modes supported by the SmartROC 3100 RAID controller card, refer to Table 4-11.

Table 4-11 Power Supply Mode Descriptions

Power Mode	Description
Minimum Power	In this mode, the static settings of the power are adjusted to the possible lowest value, and the power is dynamically reduced based on the working load.
Balanced	In this mode, the static settings of the power are adjusted in accordance with the actual RAID configuration, and the power is dynamically reduced based on the working load.
Maximum Performance	In this mode, the static settings of the power are adjusted to the possible highest value, and the power is not dynamically reduced based on the working load.

Steps

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-104.

Figure 4-104 Configure Controller Settings Screen

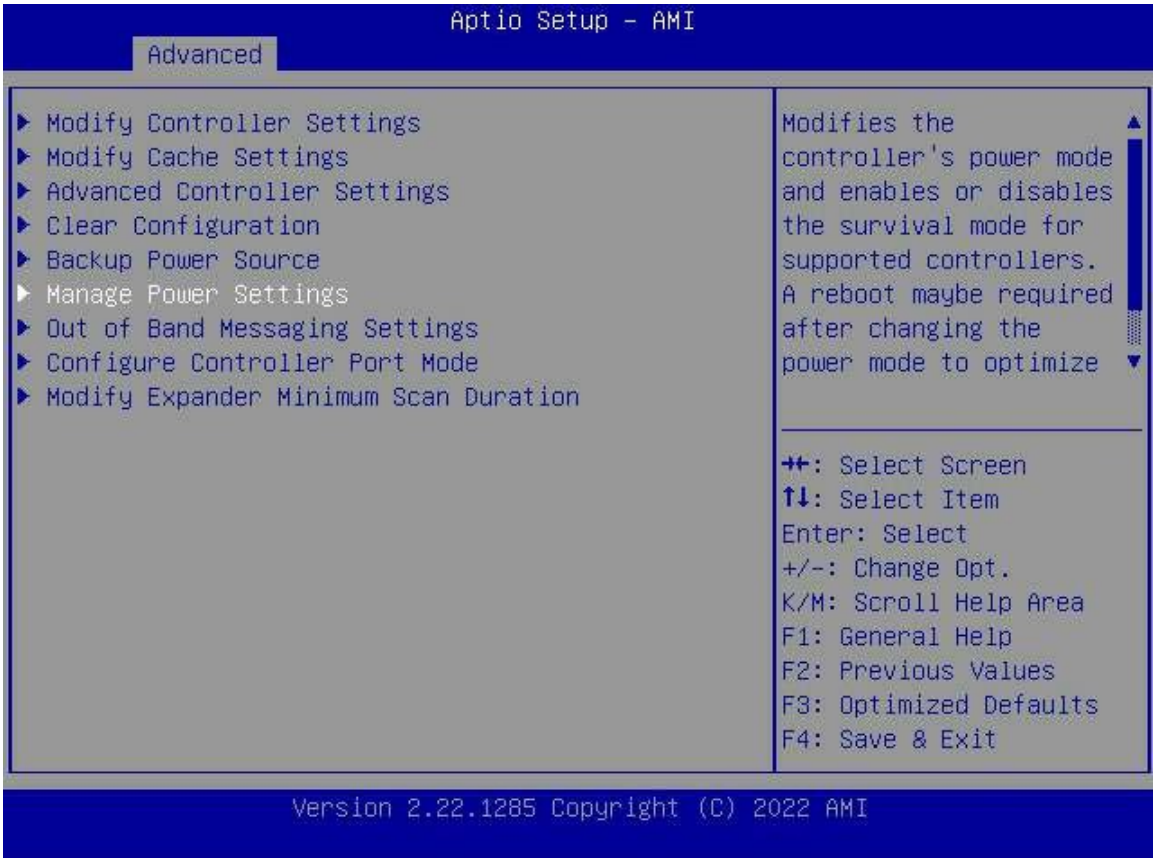

2. Use the arrow keys to select **Manage Power Settings**, and then press **Enter**. The **Manage Power Settings** screen is displayed, see Figure 4-105.

Figure 4-105 Manage Power Settings Screen

3. Use the arrow keys to select **Power Mode**, and then press **Enter**. The **Port Mode** menu is displayed, see Figure 4-106.

Figure 4-106 Power Mode Menu

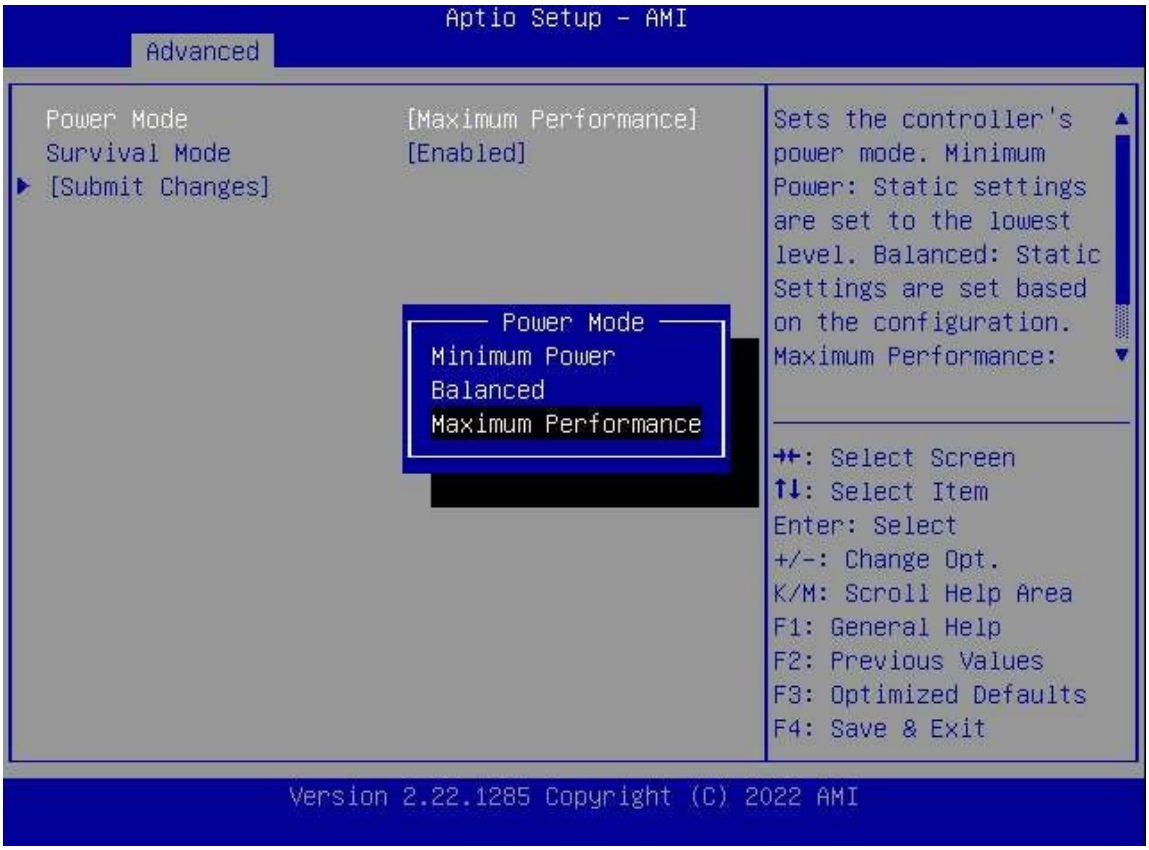

4. In accordance with your actual conditions, use the arrow keys to select the power mode to be applied, and then press **Enter** for confirmation, see Figure 4-107.

Figure 4-107 Selecting a Power Mode

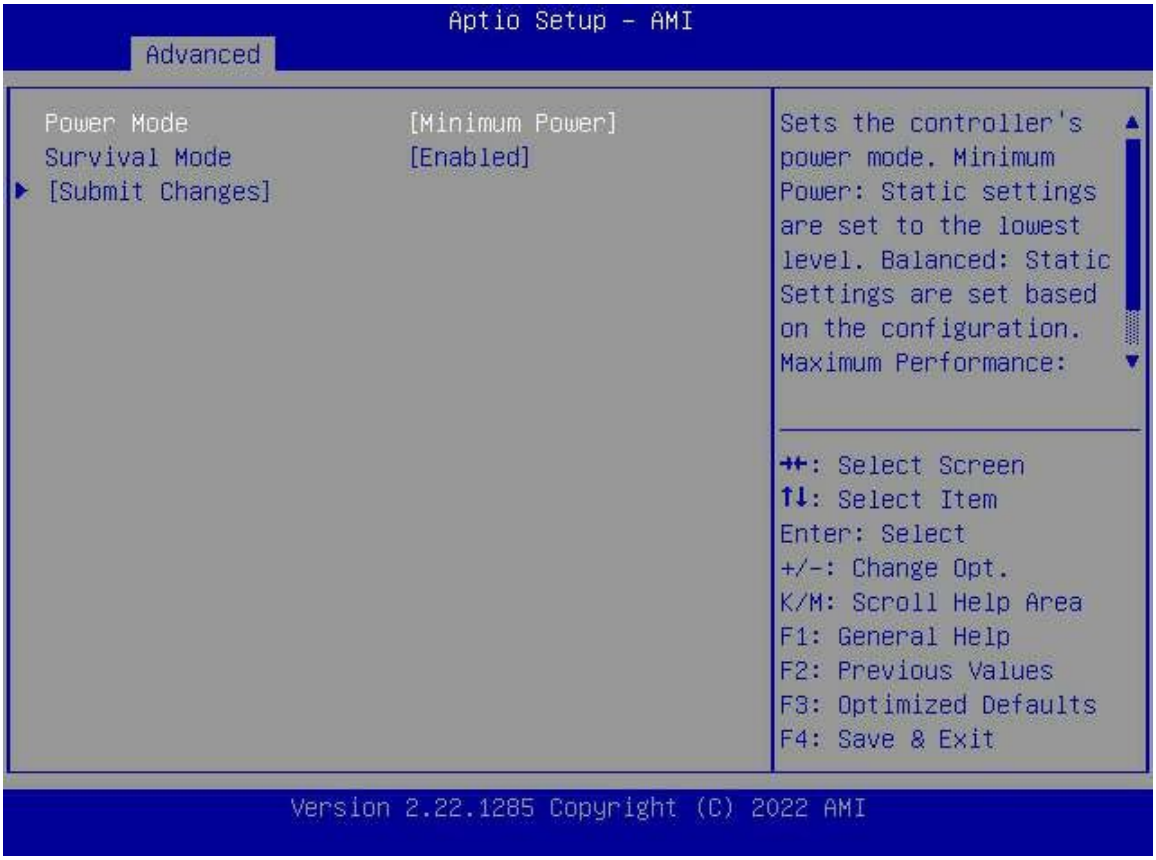

5. (Optional) Determine whether to enable **Survival Mode** as required.

Note

By default, **Survival Mode** is enabled, indicating that when the working temperature of the power supply exceeds the threshold, the RAID controller card is allowed to switch to the energy saving mode, but it may cause performance deterioration.

- Use the arrow keys to select **Survival Mode**, and then press **Enter**. The **Survival Mode** menu is displayed, see Figure 4-108.

Figure 4-108 Survival Mode Menu

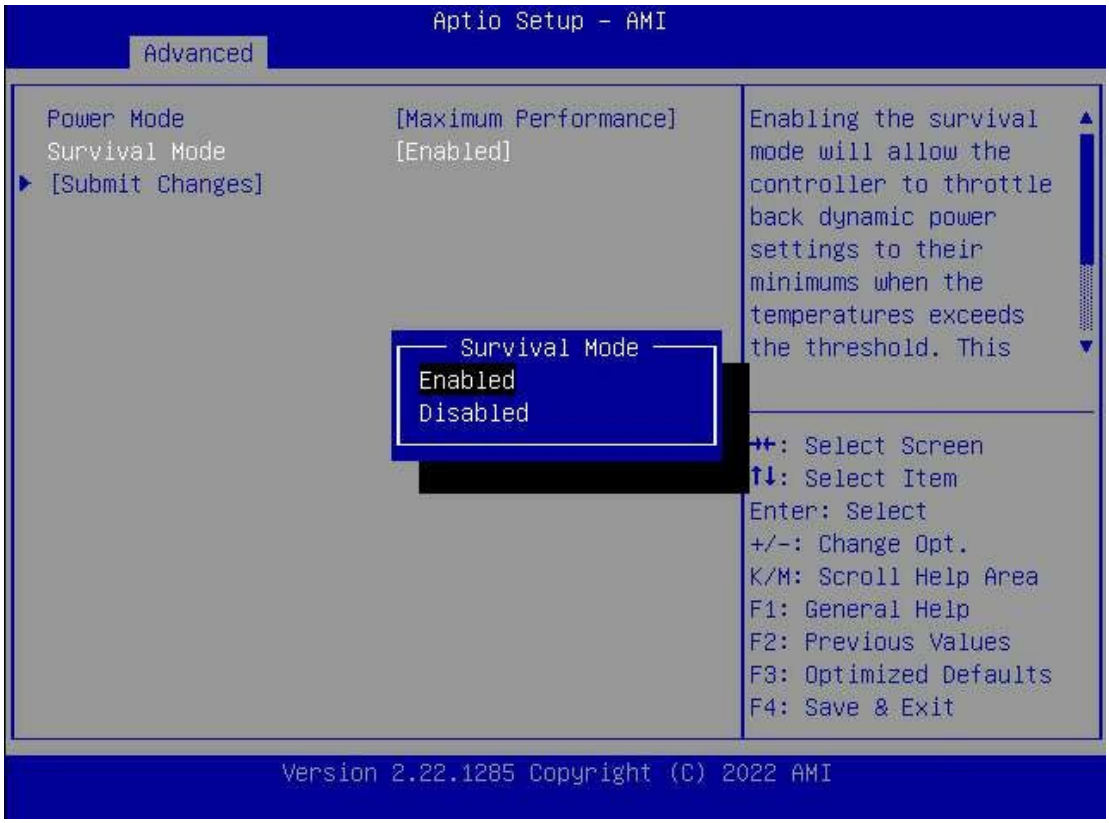

- Use the arrow keys to select **Disabled**, and then press **Enter**. The **Survival Mode** is disabled, see Figure 4-109.

Figure 4-109 Survival Mode Menu

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The power mode is set successfully, see Figure 4-110.

Figure 4-110 Power Mode Configured Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.7 Deleting a RAID Volume

Abstract

When a server no longer needs a RAID volume, you can delete the RAID volume to release the disk space.

Notice

- The data that is lost during deletion of the RAID volume cannot be restored. Therefore, you must make sure that you have backed up important data before deleting the volume.
- If the RAID logical volume to be deleted is the only logical volume on the current array, the array is also deleted after the RAID logical volume is deleted.

Prerequisite

The RAID volume is created successfully. For details, refer to "4.3.2 Creating a RAID Volume".

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-111.

Figure 4-111 Array Configuration Screen

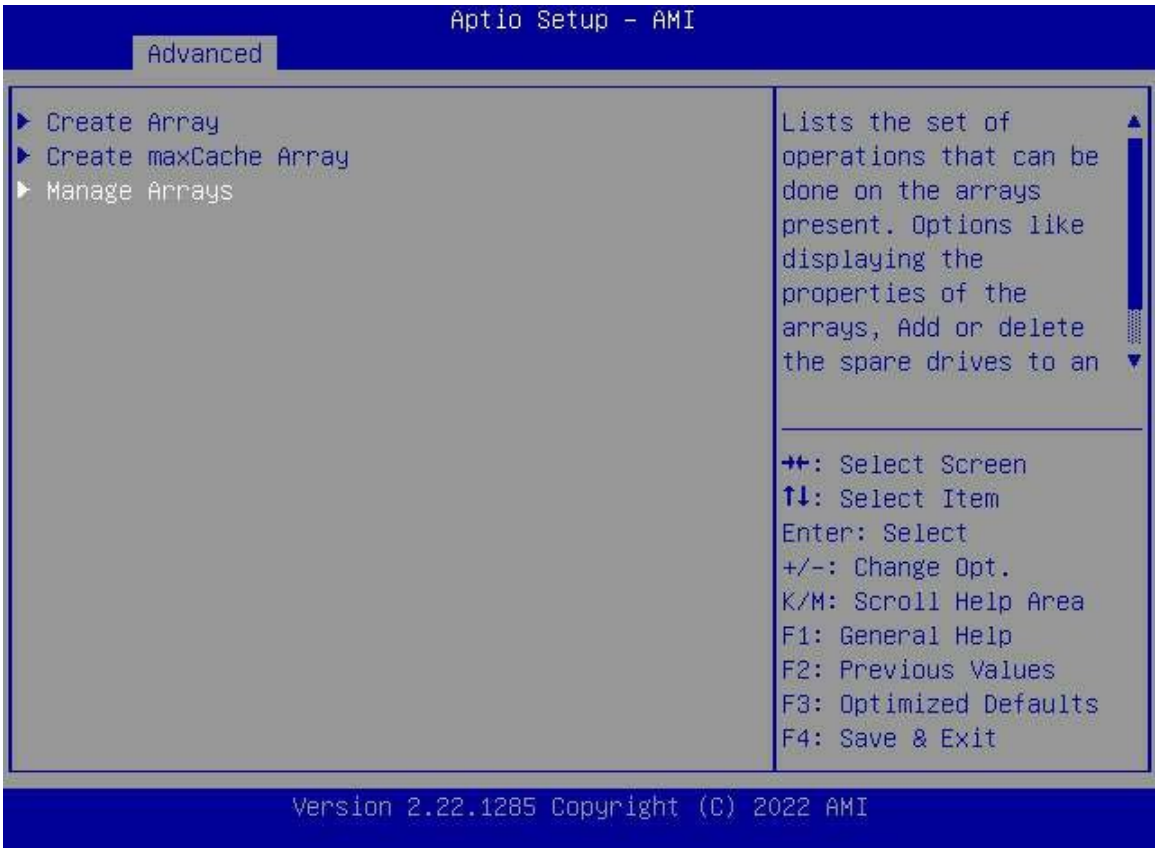

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-112.

Figure 4-112 Managing Arrays Screen

3. Use the arrow keys to select the array for which you need to delete a logical volume, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-113.

Figure 4-113 Managing the Selected Array

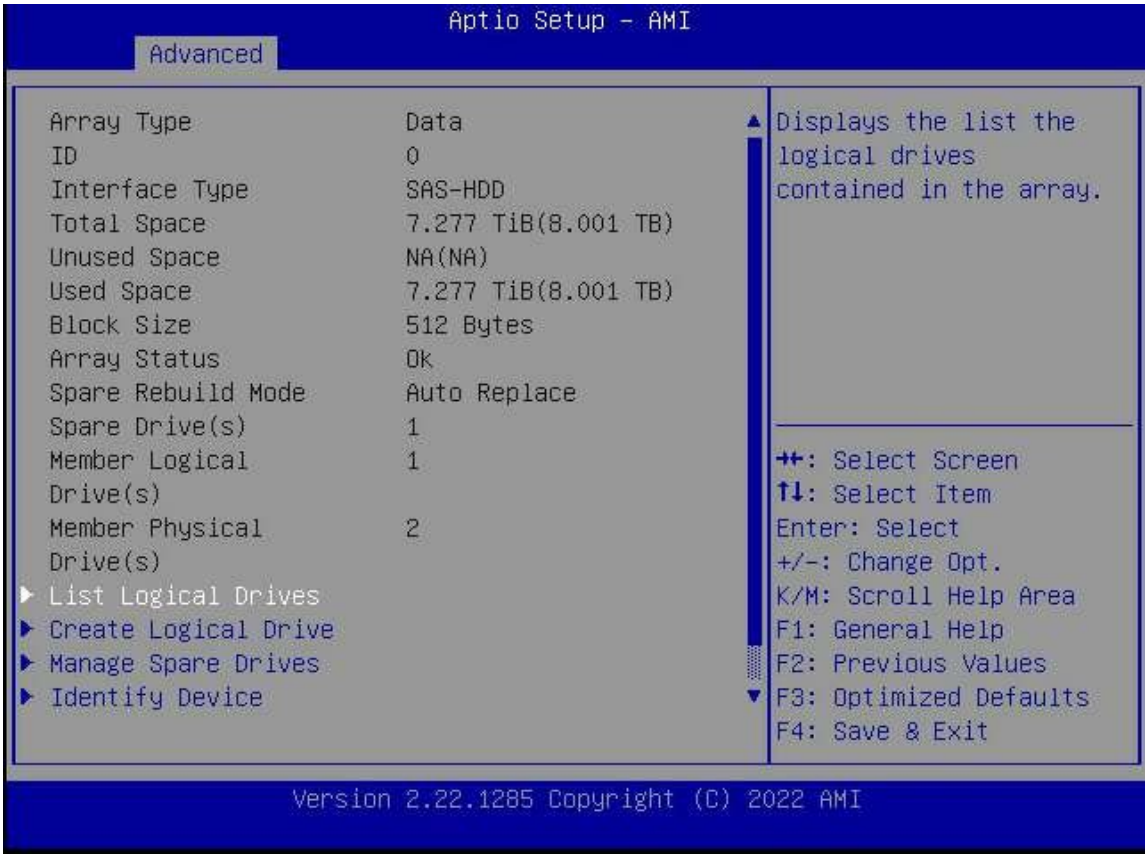

4. Use the arrow keys to select **List Logical Drives**, and then press **Enter**. The **List Logical Drives** screen is displayed, see Figure 4-114.

Figure 4-114 List Logical Drives Screen

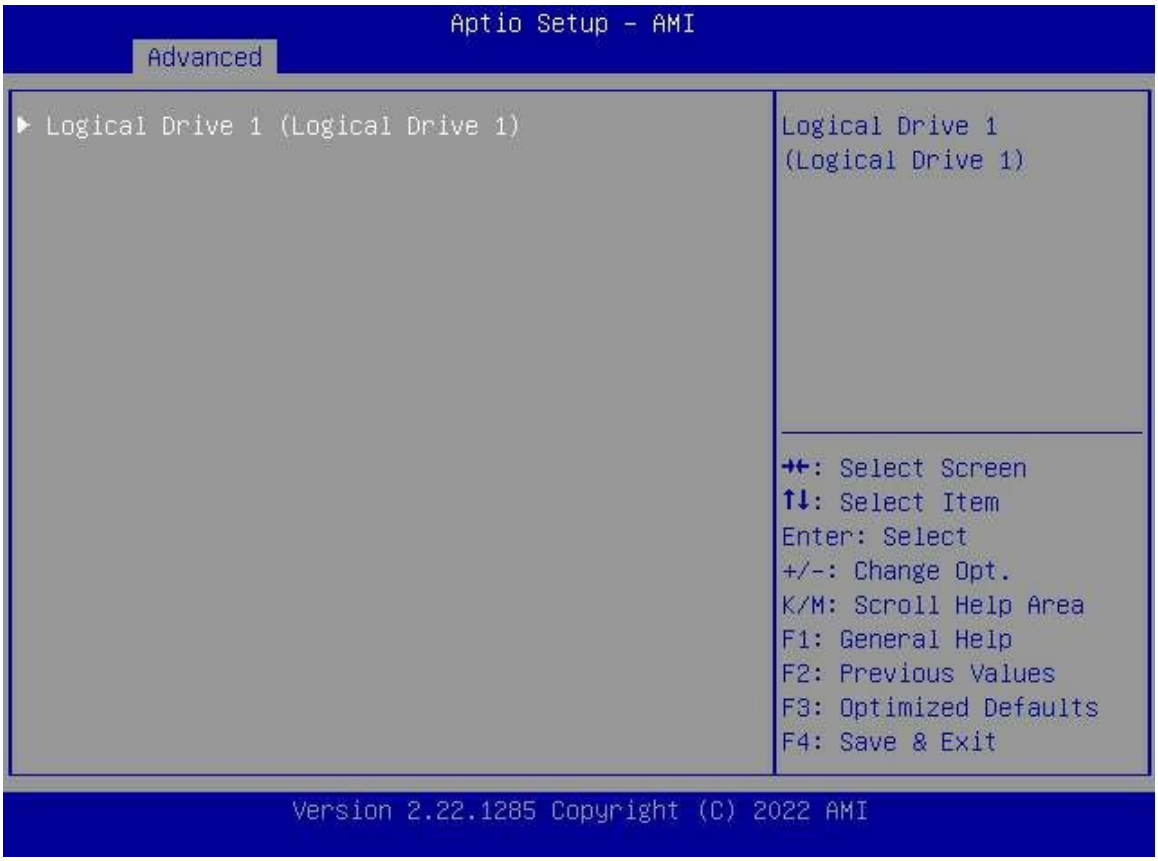

5. Use the arrow keys to select the logical volume to be deleted, and then press **Enter**. The screen for managing the selected logical volume is displayed, see Figure 4-115.

Figure 4-115 Managing the Selected Logical Volume

6. (Optional) To view the details of the logical volume, use the arrow keys to select **Logical Drive Details**, and then press **Enter**, see Figure 4-116.

Figure 4-116 Logical Volume Details

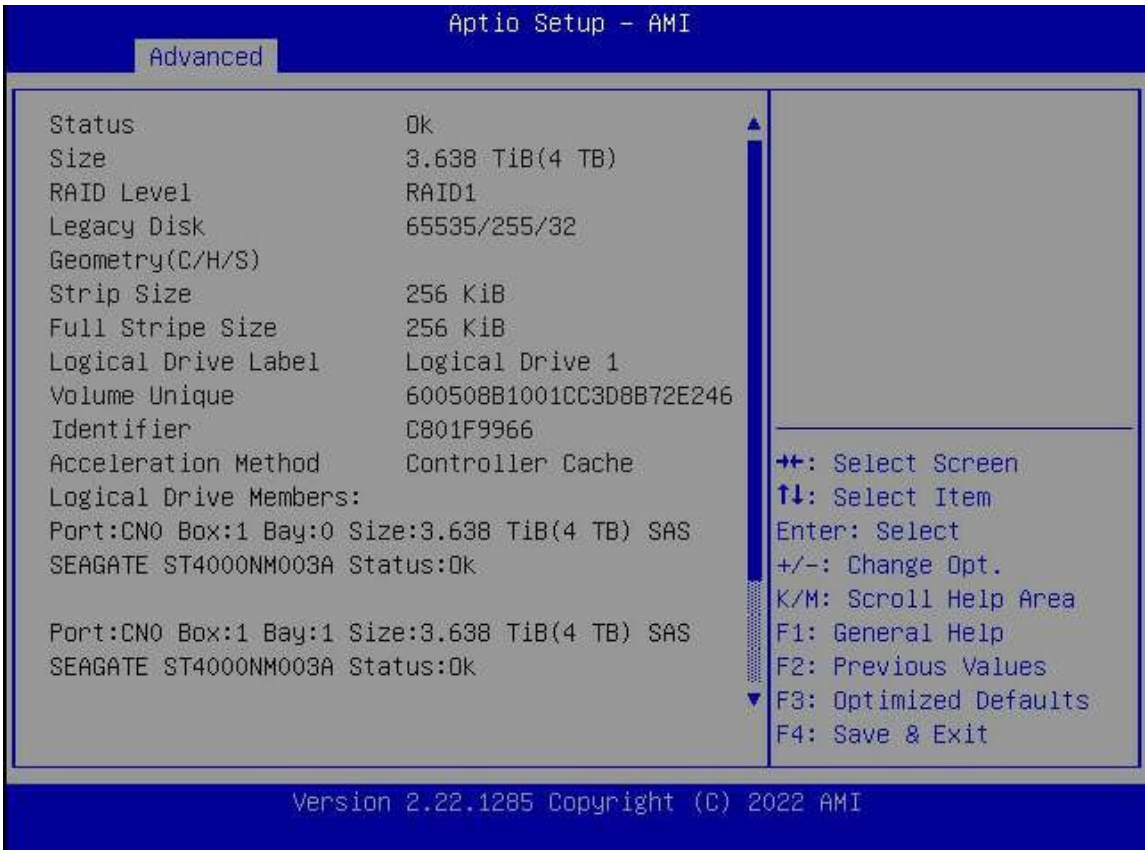

7. On the logical volume management screen, use the arrow keys to select **Delete Logical Drive**, and then press **Enter**. The confirmation screen for RAID volume deletion is displayed, see Figure 4-117.

Figure 4-117 Confirmation Screen for RAID Volume Deletion

8. Select **Submit Changes**, and then press **Enter**. The logical volume is deleted successfully, see Figure 4-118.

Figure 4-118 Logical Volume Deleted Successfully

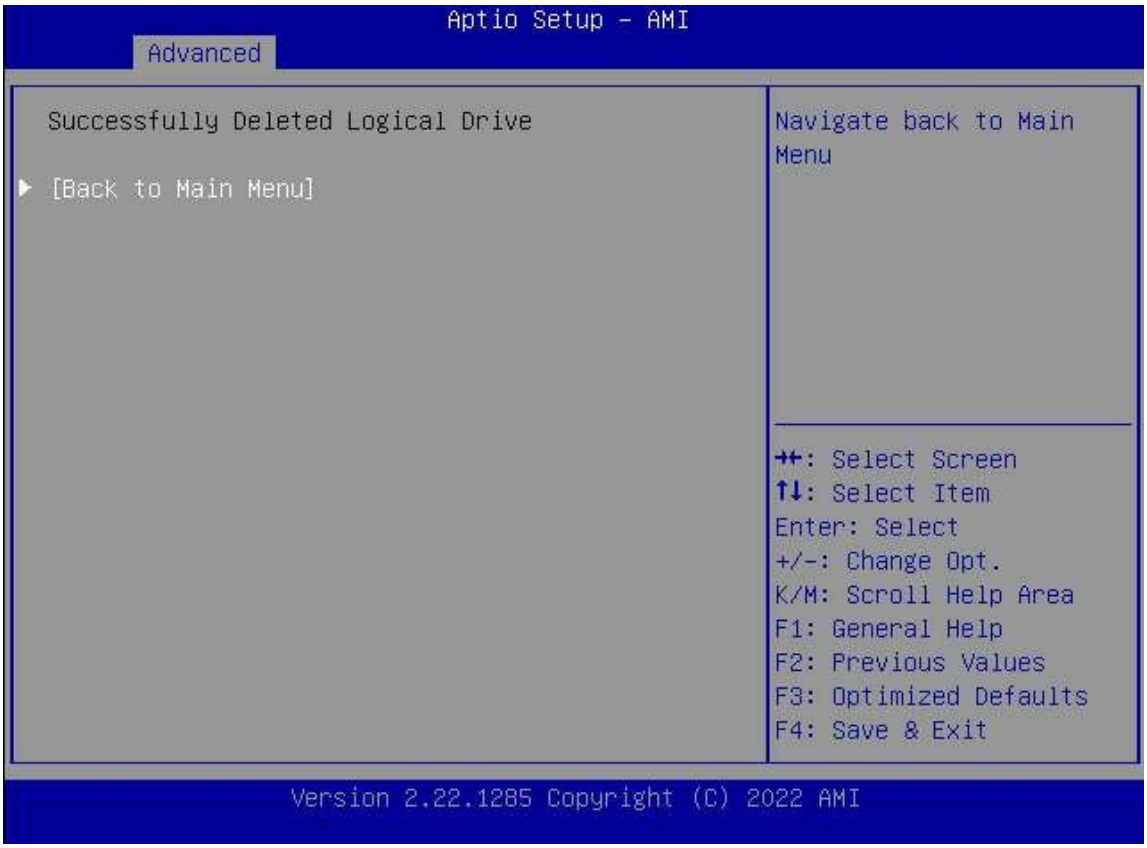

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.8 Deleting a Disk Group

Abstract

When a server no longer needs a disk group (array), you can delete it to release the disk space.

Notice

- When an array is deleted, the RAID logical volume built on it is also deleted, and the data lost during the deletion cannot be restored. Therefore, it is required to make sure that you have backed up important data before the deletion.
- If the array to be deleted is the only array of the RAID controller card, the configurations related to the RAID controller card are also cleared, and the default configurations are restored.

Prerequisite

A RAID volume is created successfully to form a disk group. For details, refer to “4.3.2 Creating a RAID Volume”.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The **Array Configuration** screen is displayed, see Figure 4-119.

Figure 4-119 Array Configuration Screen

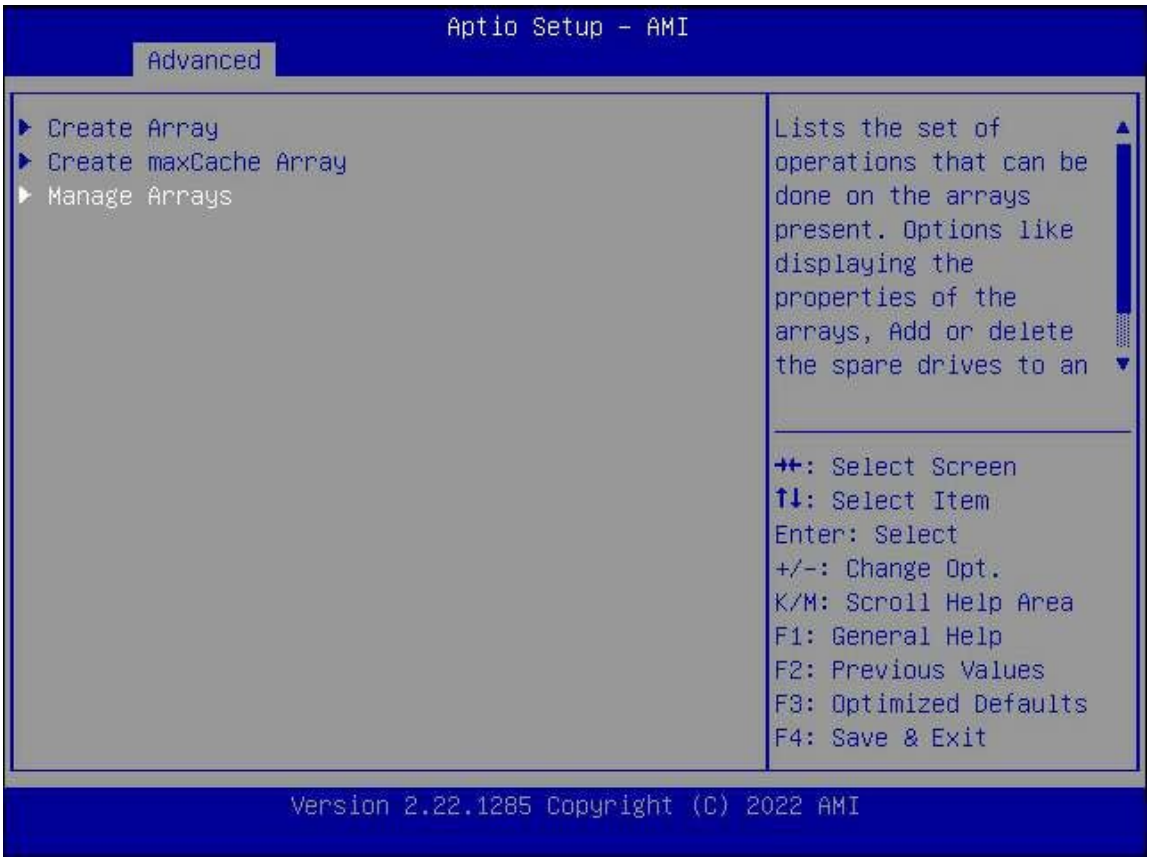

2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The **Manage Arrays** screen is displayed, see Figure 4-120.

Figure 4-120 Managing Arrays Screen

3. Use the arrow keys to select the array to be deleted, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-121.

Figure 4-121 Managing the Selected Array

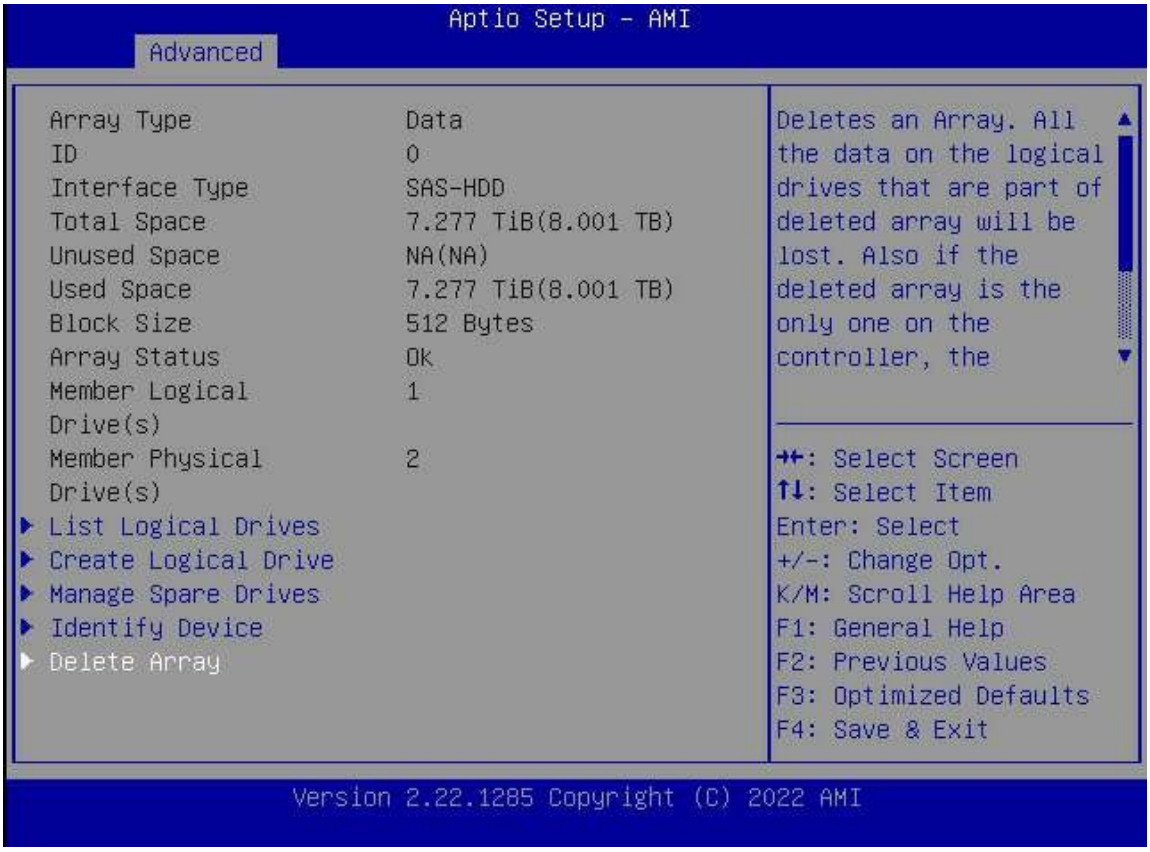

4. Use the arrow keys to select **Delete Array**, and then press **Enter**. The **Delete Array** screen is displayed, see Figure 4-122.

Figure 4-122 Delete Array Screen

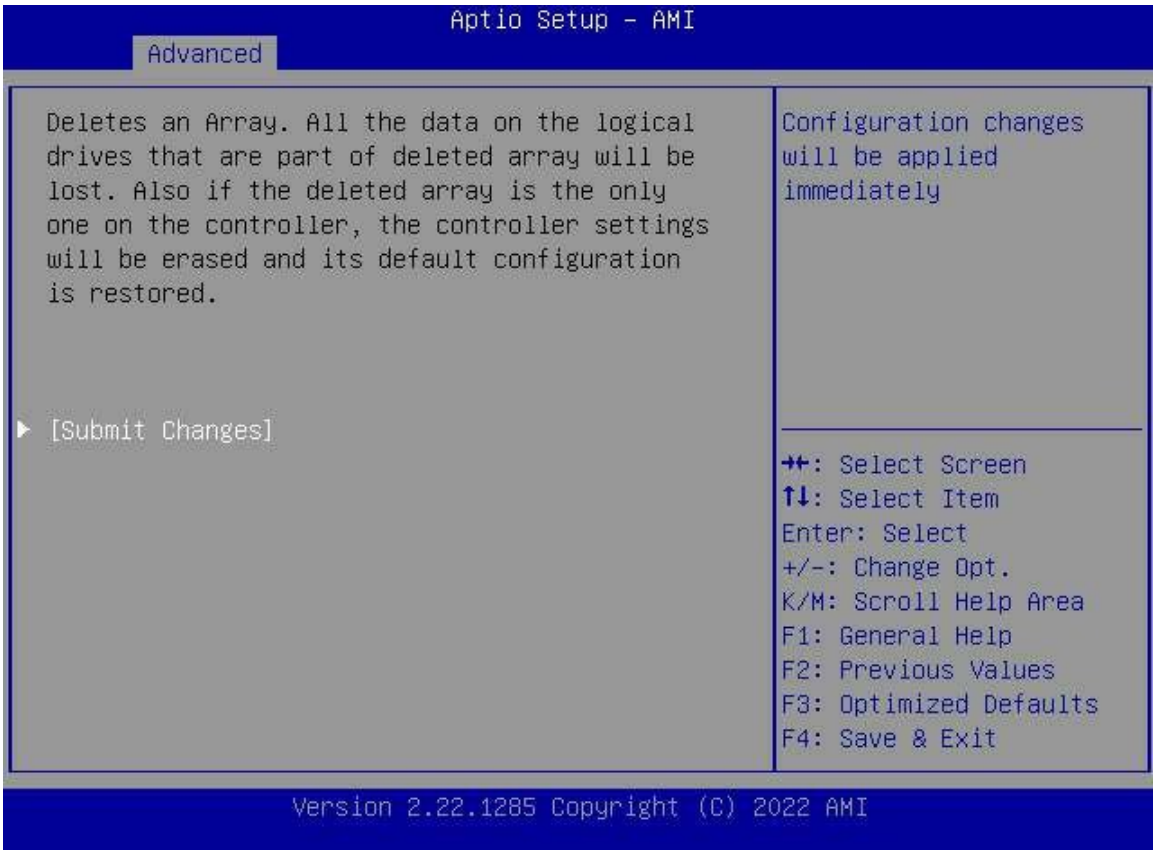

5. Use the arrow keys to select **Submit Changes**, and press **Enter**. The array is deleted successfully, see Figure 4-123.

Figure 4-123 Array Deleted Successfully

6. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.9 Clearing RAID Configuration Information

Abstract

This procedure describes how to clear all configuration information that is already created on a SmartROC 3100 RAID controller card.

Notice

The data that is lost during clearing of the configuration information on the RAID controller card cannot be restored. Therefore, it is required to make sure that you have backed up important data before the clearing operation.

Prerequisite

A RAID volume is created successfully and it has the corresponding RAID configuration information. For details, refer to “4.3.2 Creating a RAID Volume”.

Steps

1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-124.

Figure 4-124 Configure Controller Settings Screen

2. Use the arrow keys to select **Clear Configuration**, and then press **Enter**. The screen for clearing the configurations is displayed, see Figure 4-125.

Figure 4-125 Clearing the Configurations

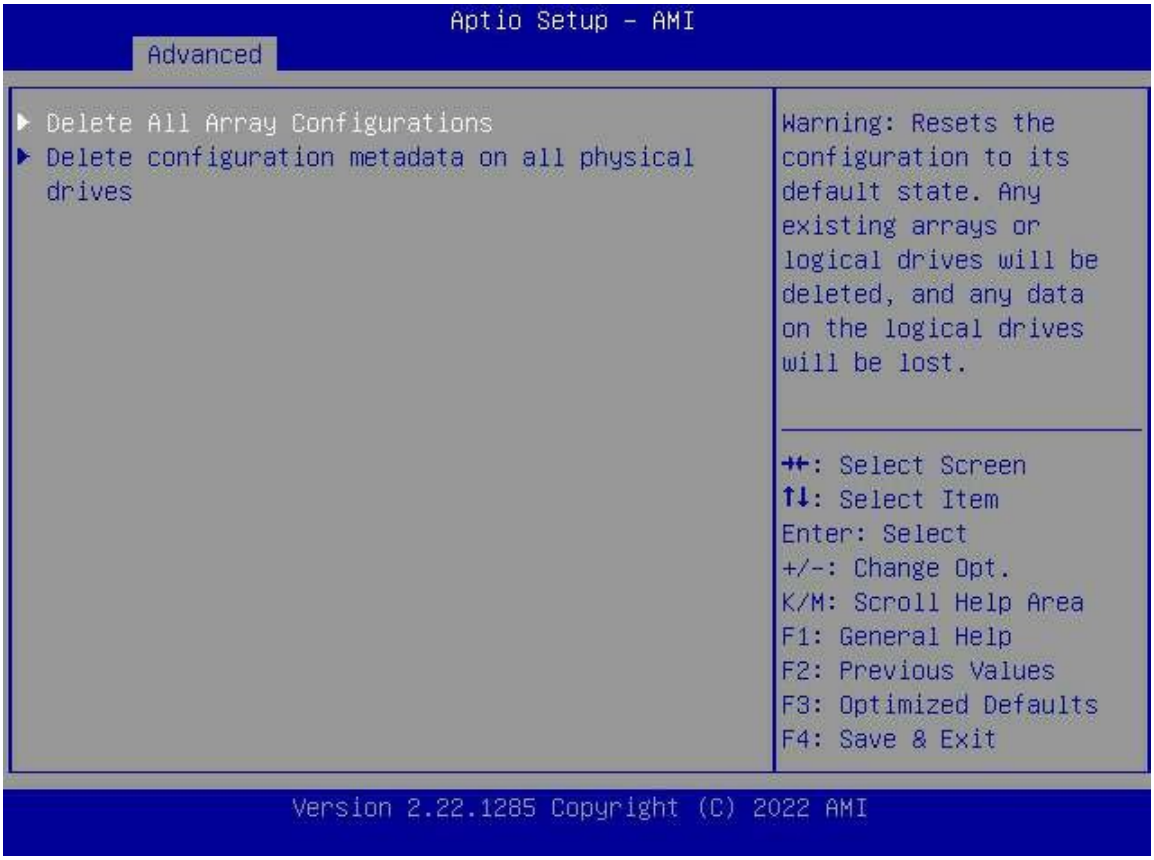

3. Use the arrow keys to select **Delete All Array Configuration**, and then press **Enter**. The screen for clearing the array configurations is displayed, see Figure 4-126.

Figure 4-126 Clearing the Array Configurations

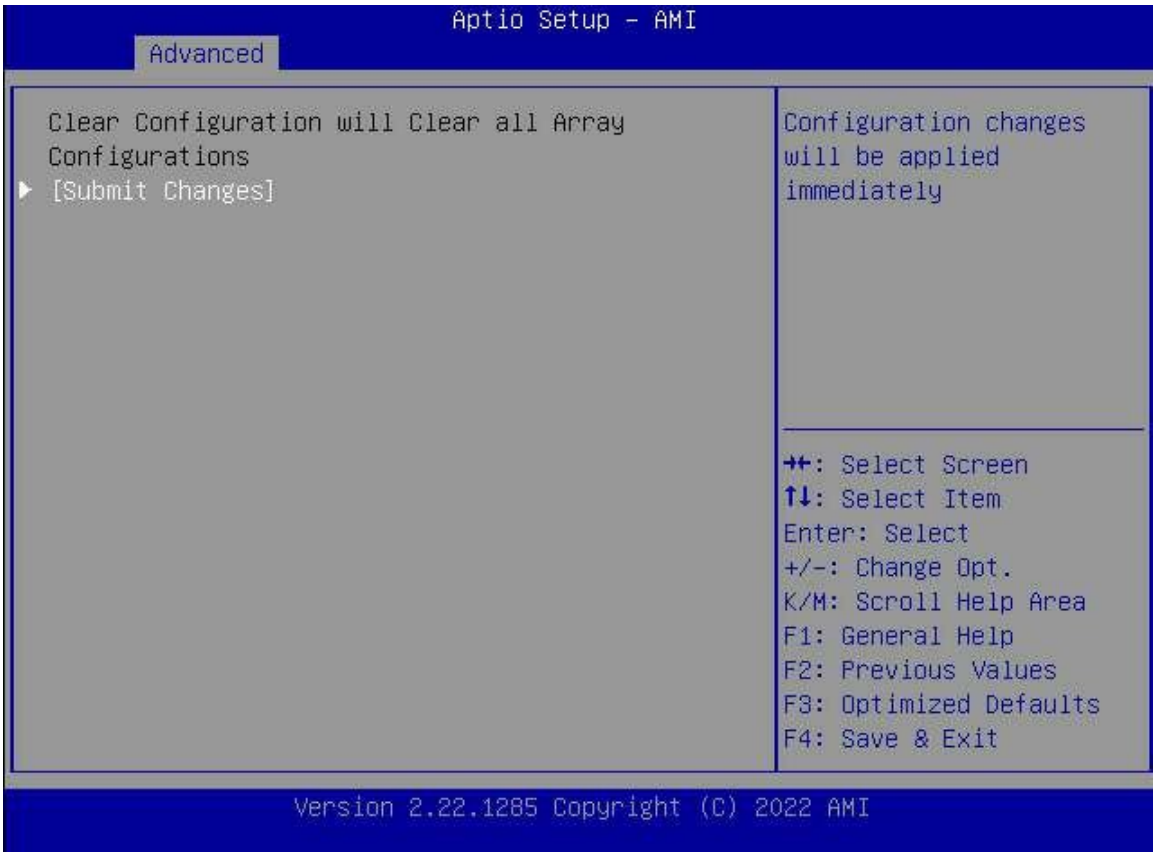

4. Use the arrow keys to select **Submit Changes**, and press **Enter**. The configuration information is cleared successfully, see Figure 4-127.

Figure 4-127 Configuration Cleared Successfully

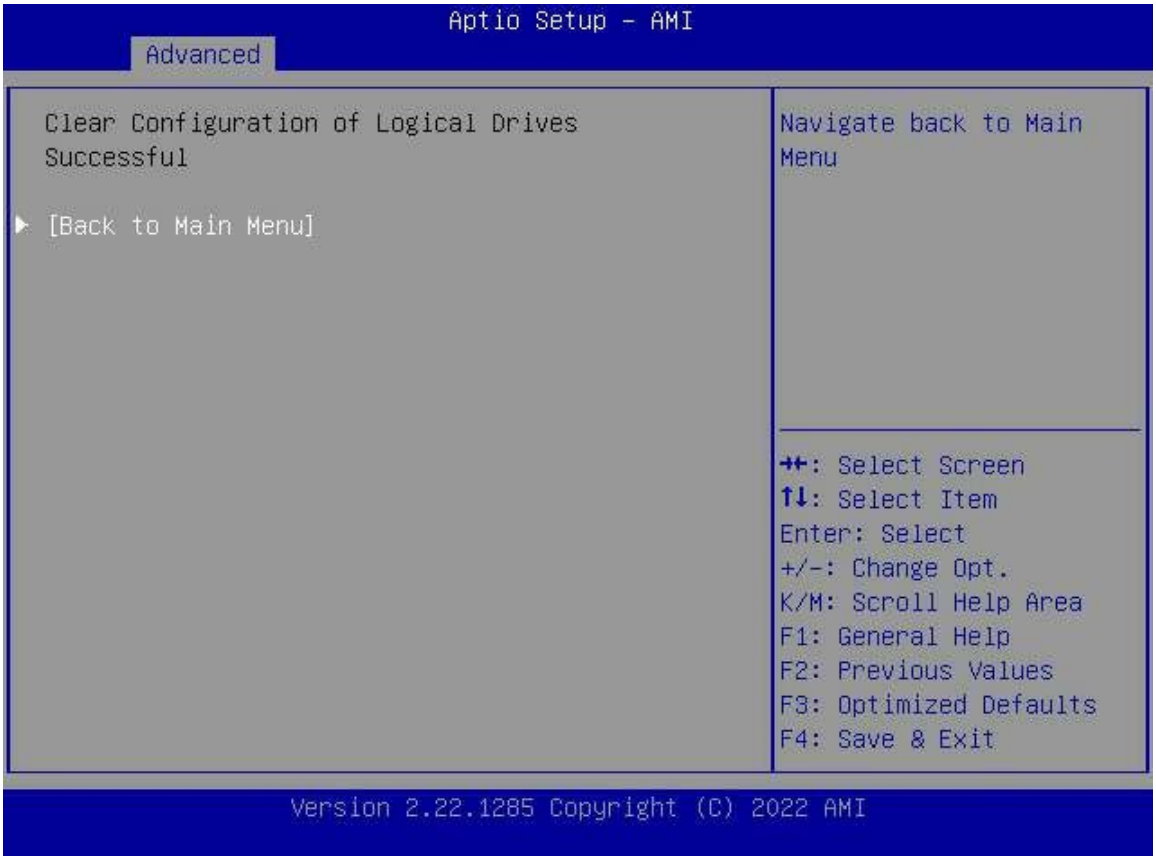

5. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.10 Configuring a Pass-Through Disk

Abstract

In accordance with the actual RAID configuration, you can configure a pass-through disk in either of the following ways:

- Setting the mode of all ports in batches:
  If the disks connected to all ports are not used to create a RAID volume, the mode of these ports can be set to HBA in a unified manner.
- Setting the mode of ports separately:
  If the disks connected to some ports have been used to create a RAID volume, the mode of these ports can be set to Mixed and that of other ports can be set to HBA.

Context

Port modes include RAID, HBA and Mixed, which are described as follows:

- In RAID mode, the connected disks can be used only after they form a RAID volume.

- In HBA mode, the connected disks are pass-through disks and cannot be used to create a RAID volume. Instead, they can only be used directly.
- In Mixed mode, the connected disks support both RAID and HBA mode.
  → The RAID mode is applicable to the disks that have been used to create a RAID volume.
  → The HBA mode (pass-through) is applicable to the disks that are not used to create a RAID volume.

Steps

- Setting the Mode of All Ports in Batches
  1. On the controller management screen, use the arrow keys to select **Controller Information**, and then press **Enter**. The **Controller Information** screen is displayed. The value of **Controller Mode** is **Mixed**, see Figure 4-128.

Figure 4-128 Controller Information Screen

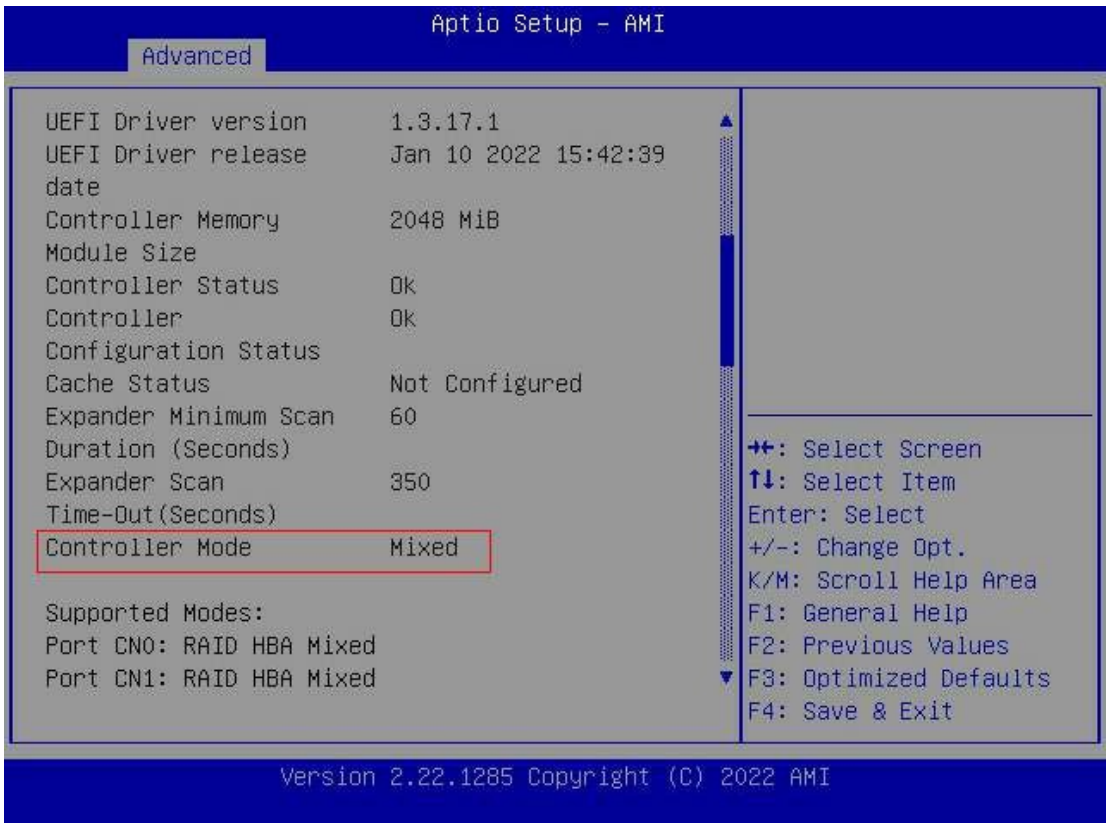

  2. Press **Esc** to return to the controller management screen. Use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-129.

Figure 4-129 Configure Controller Settings Screen

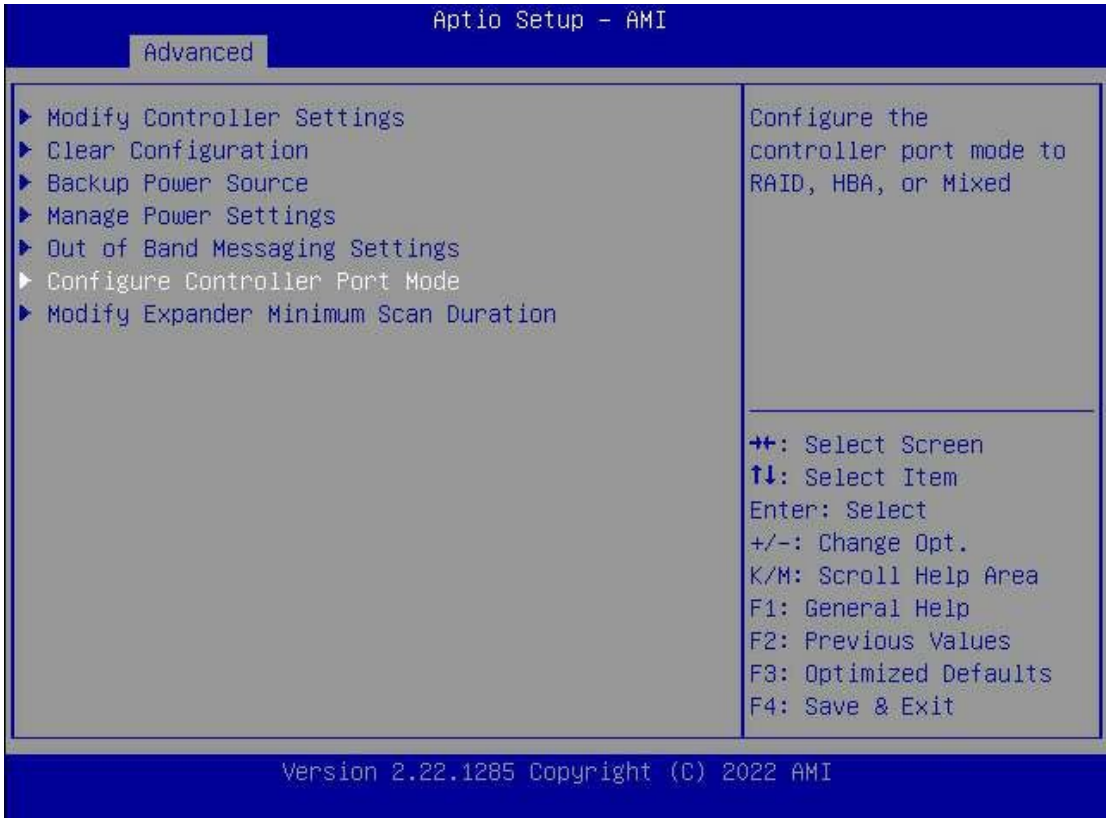

3. Use the arrow keys to select **Configure Controller Port Mode**, and then press **Enter**.
   The screen for setting the port mode is displayed, see Figure 4-130.

Figure 4-130 Setting the Port Mode

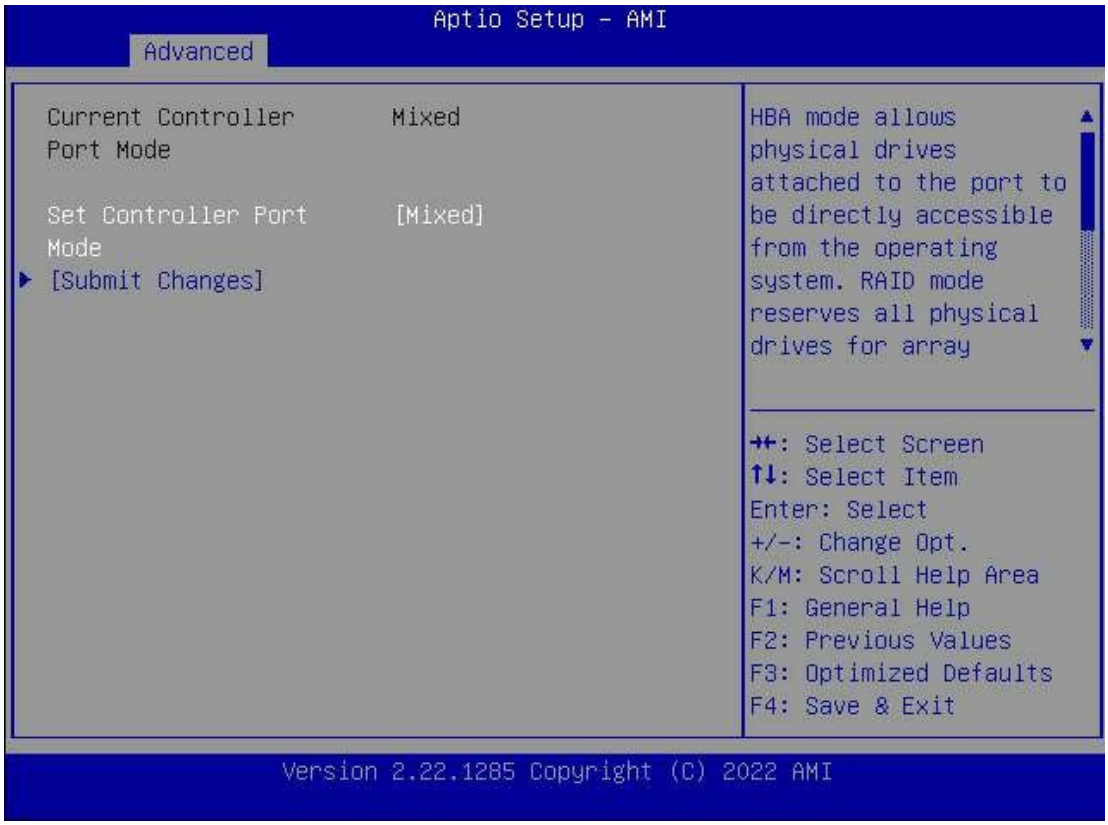

4. Use the arrow keys to select **Set Controller Port Mode**, and then press **Enter**. The **Set Controller Port Mode** dialog box is displayed, see Figure 4-131.

Figure 4-131 Set Controller Port Mode Dialog Box

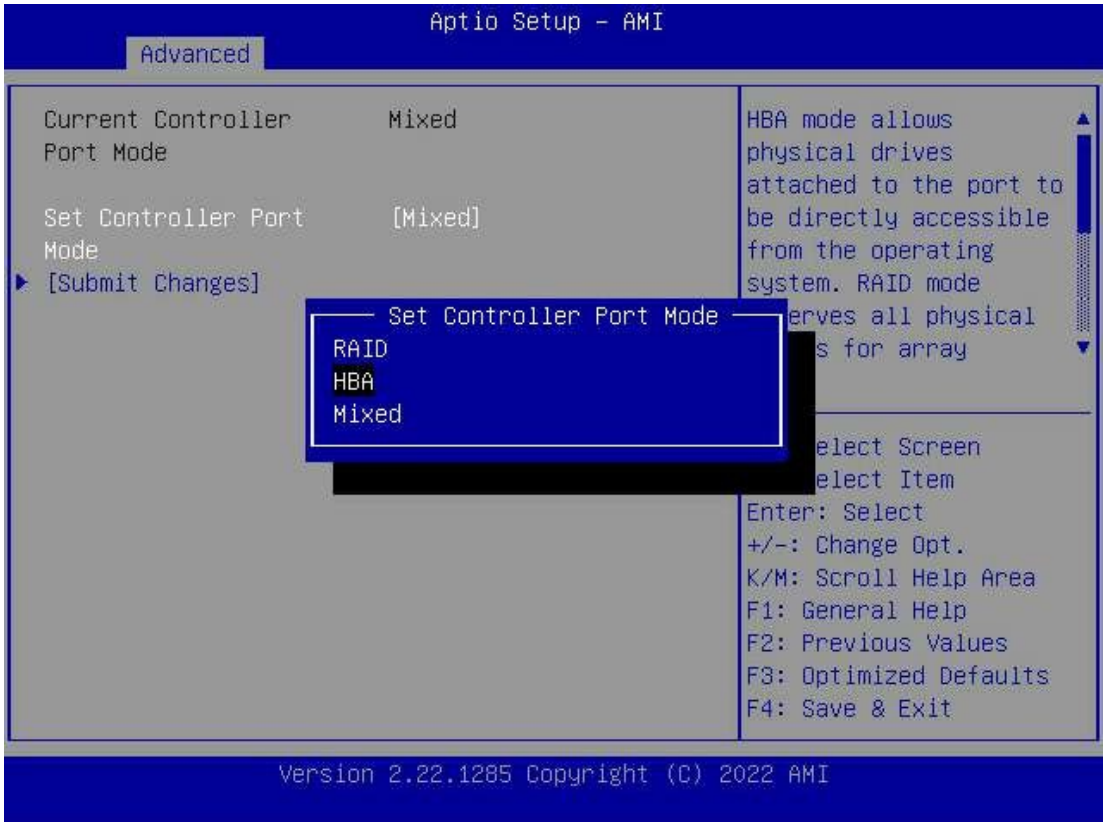

5. Use the arrow keys to select **HBA**, and then press **Enter**. The mode of all ports is set to **HBA**, see Figure 4-132.

Figure 4-132 Setting the Port Mode to HBA

6. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 4-133.

Figure 4-133 Port Mode Set Successfully

7. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen, see Figure 4-134.

Figure 4-134 Controller Management Page

8. Use the arrow keys to select **Controller Details**, and then press **Enter**. On the displayed **Controller Information** screen, view the port mode after the modification, see Figure 4-135. The value of **Controller Mode** is **HBA**.

Figure 4-135 Controller Information Screen

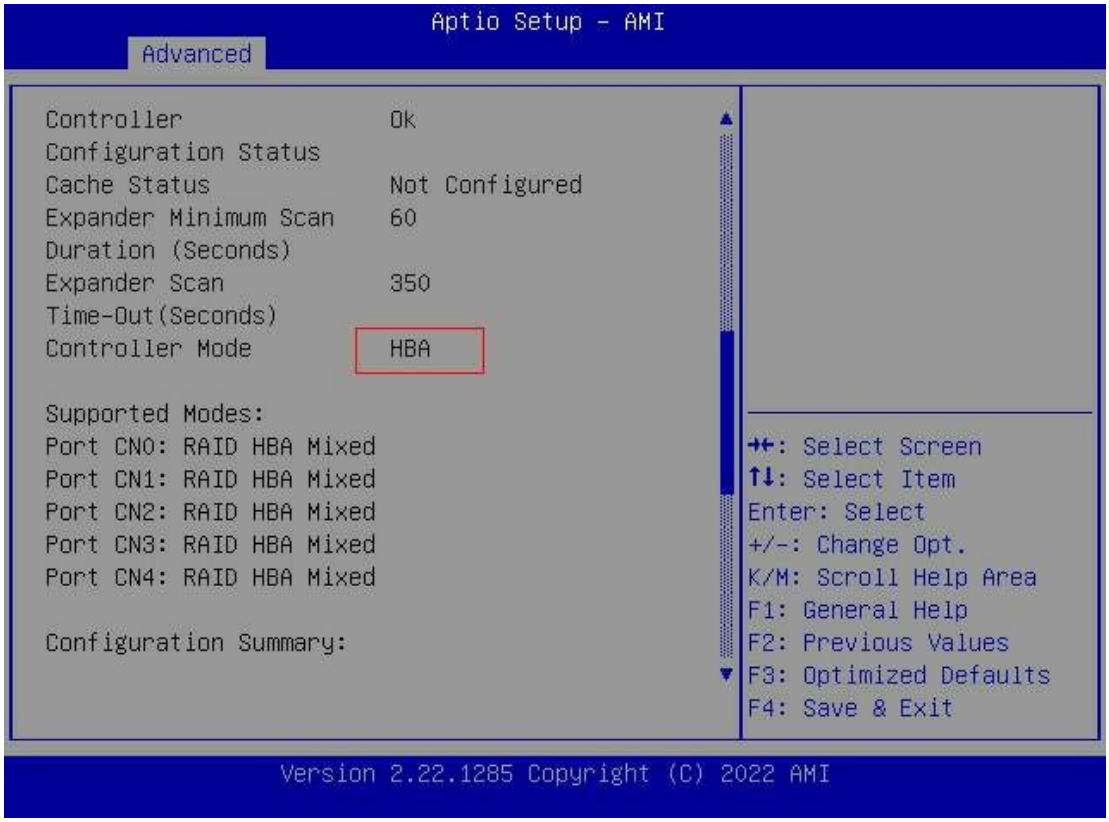

- Setting the Mode of a Port Separately
  1. On the controller management screen, use the arrow keys to select **Configure Controller Settings**, and then press **Enter**. The **Configure Controller Settings** screen is displayed, see Figure 4-136.

Figure 4-136 Configure Controller Settings Screen

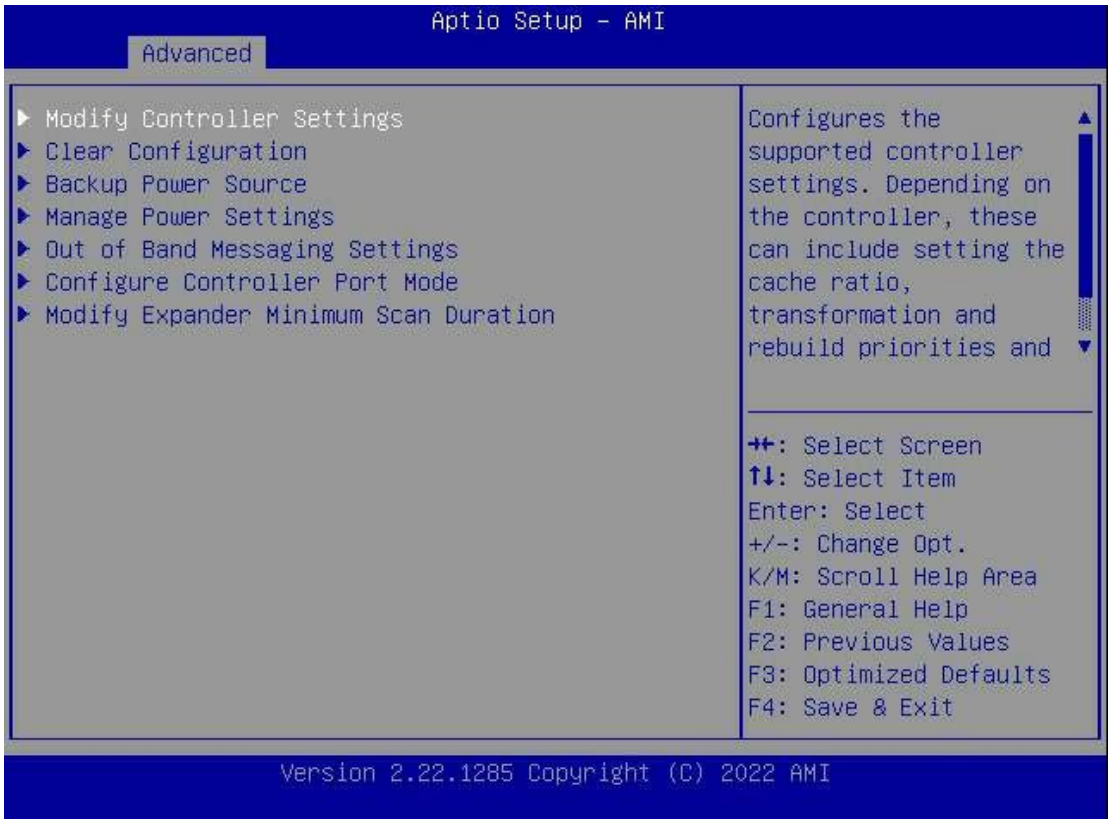

2. Use the arrow keys to select **Modify Controller Settings**, and then press **Enter**. The screen for setting the controller is displayed, see Figure 4-137.

Figure 4-137 Setting the Controller

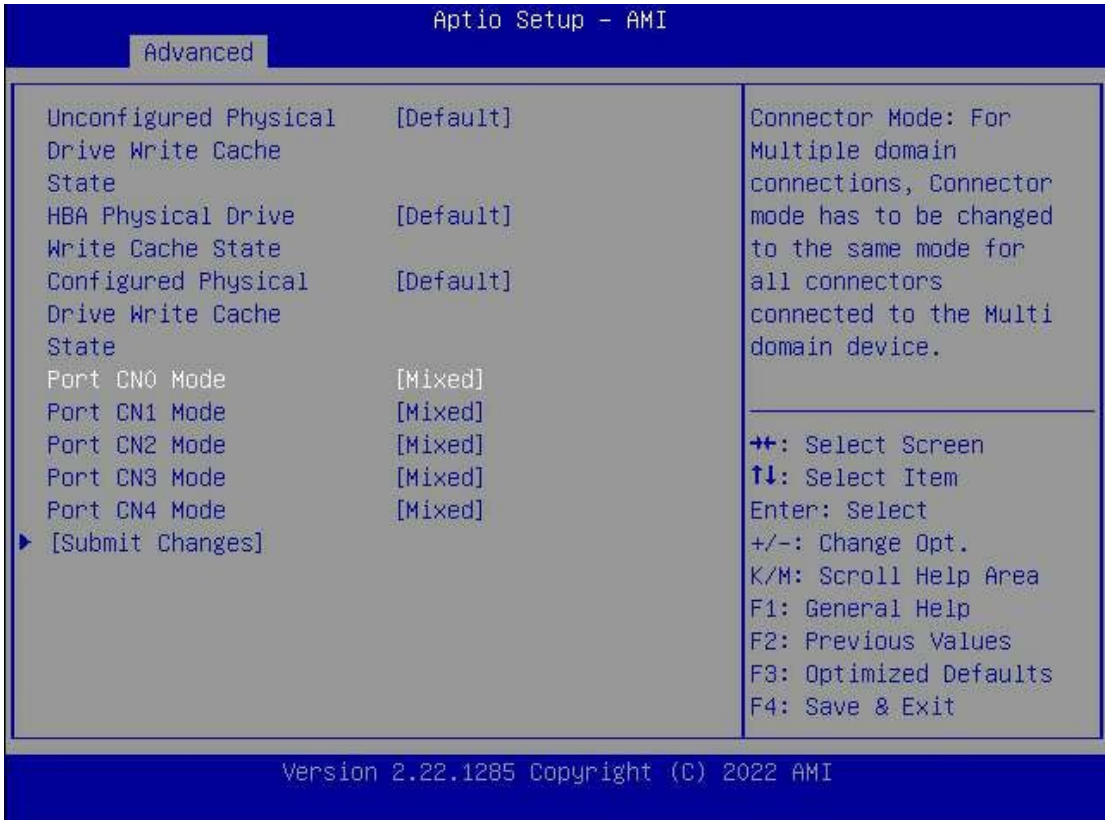

3. Use the arrow keys to select **Port CN0 Mode**, and then press **Enter**. The **Port CN0 Mode** dialog box is displayed, see Figure 4-138.

Figure 4-138 Port CNO Mode Dialog Box

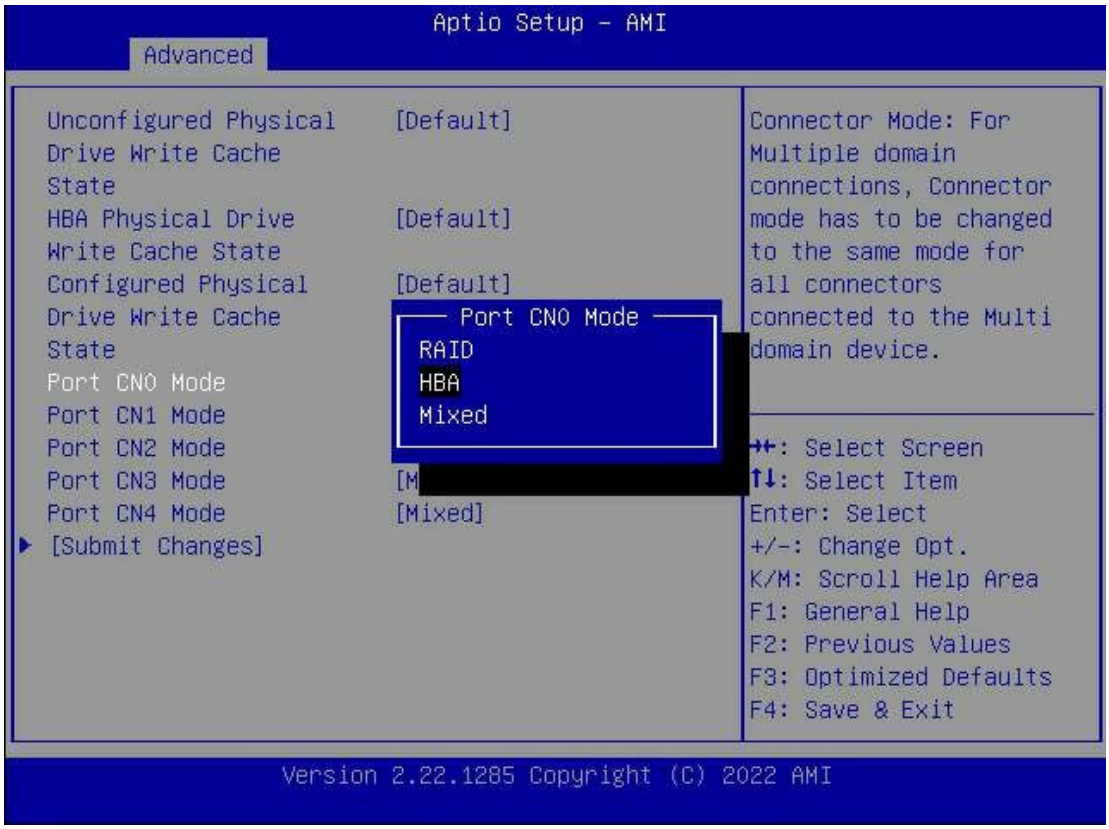

4. Use the arrow keys to select **HBA**, and then press **Enter**. The mode of the port whose connected disk is not used to create a RAID volume is set to **HBA**, see Figure 4-139.

Figure 4-139 Setting the Port Mode to HBA

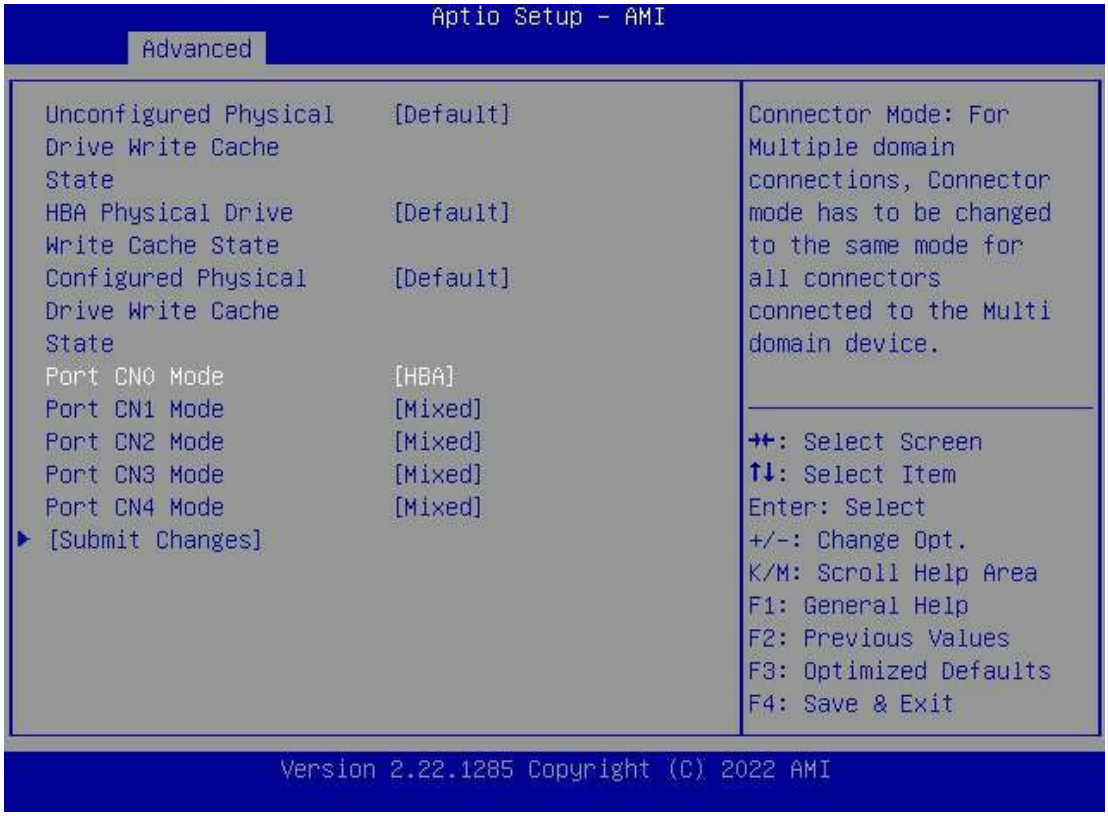

5. Repeat Step 3 and Step 4 to set the mode of another port whose connected disk is not used to create a RAID volume to **HBA**.
6. Use the arrow keys to select a port whose connected disk is already used to create a RAID volume (using **Port CN2 Mode** as an example), and then press **Enter**. The port mode setting menu is displayed, see Figure 4-140.

Figure 4-140 Port Mode Setting Menu

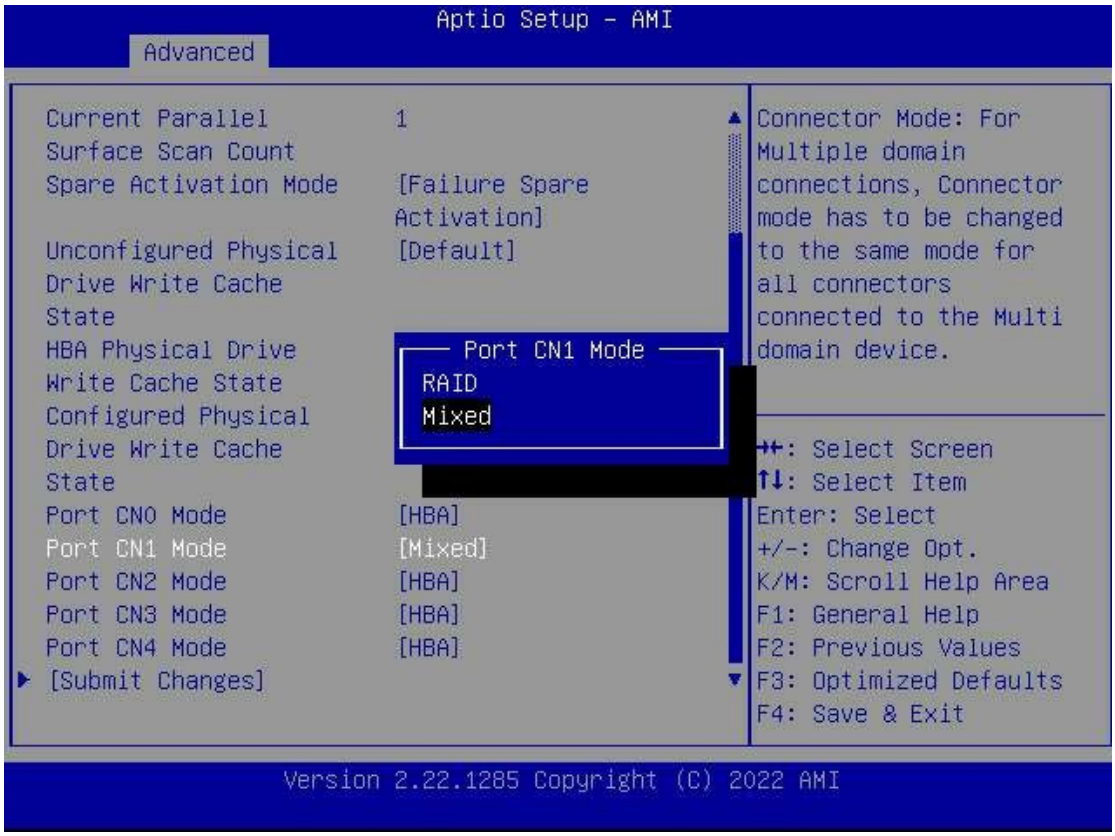

7. Use the arrow keys to select **Mixed**, and then press **Enter**. The mode of the port whose connected disk is used to create a RAID volume is set to **Mixed**, see Figure 4-141.

Figure 4-141 Setting the Port Mode to Mixed

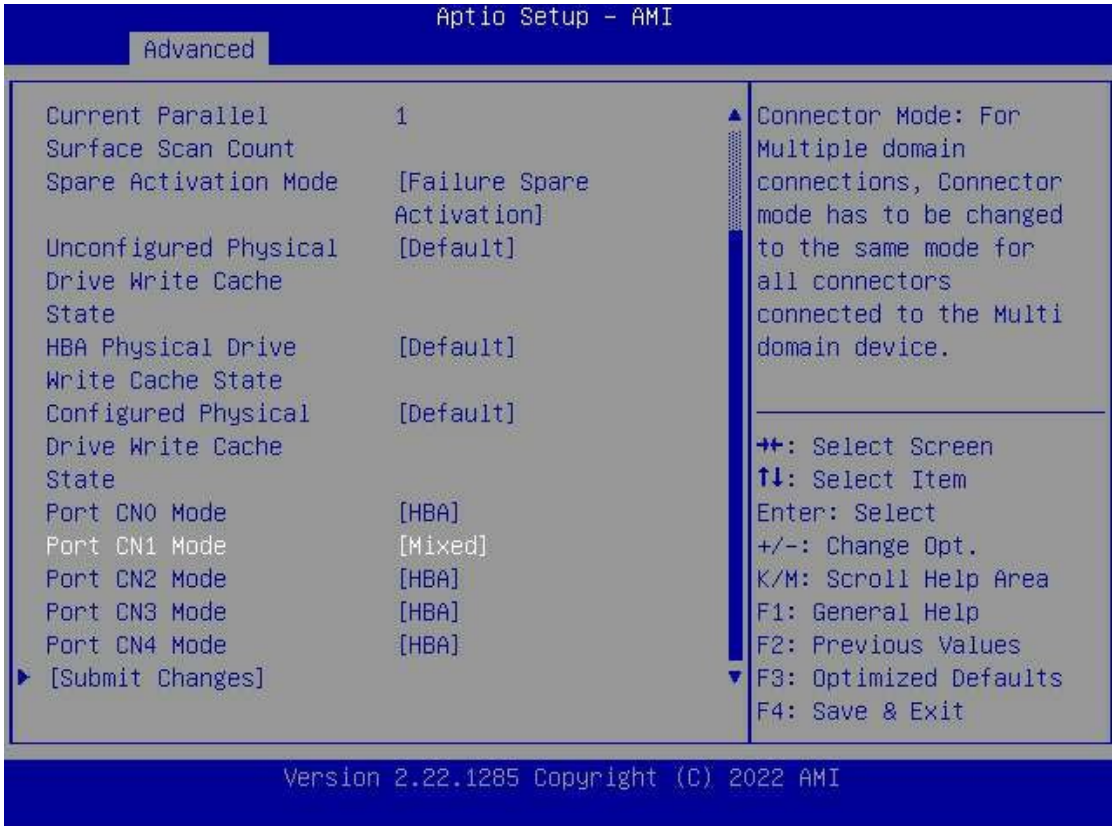

Note

The mode of the ports whose connected disks are already used to create a RAID volume cannot be set to HBA.

8. Use the arrow keys to select **Submit Changes**, and press **Enter**. The port mode is set successfully, see Figure 4-142.

Figure 4-142 Port Mode Set Successfully

9. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen.

4.5.11 Enabling the Caching Function

Abstract

A VT SmartROC 3100 RAID controller card supports the caching function. In UEFI mode, a SmartROC 3100 RAID supports the following two caching modes:

- Controller Cache: enables controller cache optimization. The read cache and write cache are used at the same time.
- None: disables the controller cache.

After you select the **Controller Cache** mode, the caching function is enabled, which improves the data read/write speed.

Note

You can enable the caching function only by referring to this procedure. The caching function cannot be enabled on the Web portal of the BMC.

Prerequisite

A RAID volume is created successfully. For details, refer to “3.3.2 Creating a RAID Volume”.

Context

Enabling the caching function improves the data read/write speed. The details are as follows:

- When a RAID controller card reads the data, if the data has been written into the Cache, the data can be directly read from the Cache to prevent the hard disk from searching for the data again, thus saving the response time and improving the data read speed.
- When a RAID controller card writes the data , the data is directly written into the Cache. The RAID controller card refreshes the data to the hard disk only when the written data is accumulated to a certain extent, achieving batch data write. In addition, the Cache is a fast read/write device, so the read/write speed of the Cache is higher than that of the hard disk, thus improving the data write speed.

Note

To ensure the data security in the Cache, you can configure a super capacitor for the RAID controller card. In case of unexpected power failure of the server, the super capacitor is used to supply power, and provides data security protection in the Cache.

Steps

1. On the controller management screen, perform the following operations as required.

If...	Then...
The super capacitor is configured	View the status of the backup power. The value of **Backup Power Status** is **charged**.
The super capacitor is not configured	a. Disable **Survival Mode**. i. Use the arrow keys to select **VT SmartROC3100 RS241-18i > Array Configuration > Manage Arrays**, and press **Enter**. The **Manage Arrays** screen is displayed. ii. Use the arrow keys to select the array for which you need to disable **Survival Mode**, and then press **Enter**. The screen for managing the selected array is displayed. iii. Use the arrow keys to select **I/O Bypass Settings**, and press **Enter**. In the displayed dialog box, set **I/O Bypass** to **Disabled**. b. Enable **No Battery Write Cache**. i. Use the arrow keys to select **VT SmartROC3100 RS241-18i > Configure Controller Settings > Modify Cache Settings**, and press **Enter**. The **Modify Cache Settings** screen is displayed. ii. Use the arrow keys to select **No Battery Write Cache**, and press **Enter**. In the displayed dialog box, set **No Battery Write Cache** to **Enabled**.

Note

When the super capacitor is not configured and the **No battery Write Cache** is disabled, setting **Acceleration Method** to **Controller Cache** will not take effect.

2. Press **Esc** multiple times until the controller management screen is displayed.
3. Use the arrow keys to select **Array Configuration**, and then press **Enter**. The screen for configuring an array is displayed, see Figure 4-143.

Figure 4-143 Configuring an Array

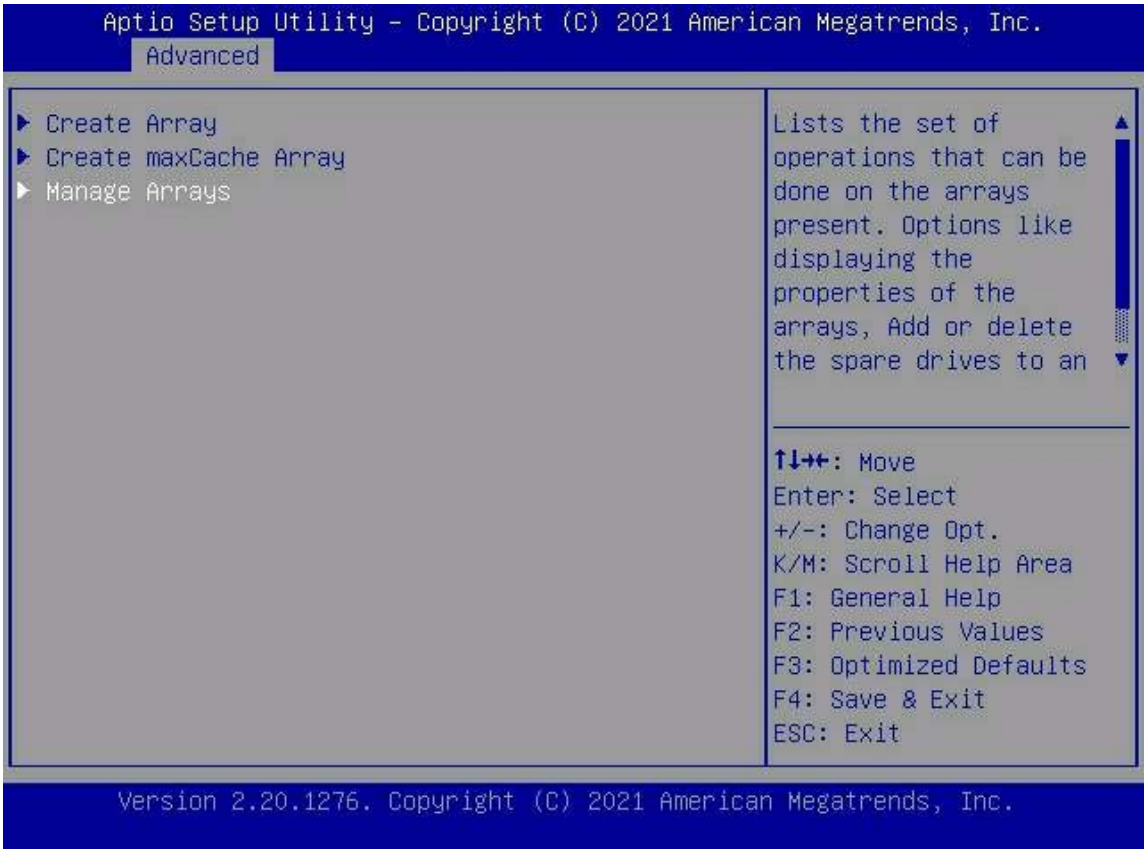

4. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The screen for managing arrays is displayed, see Figure 4-144.

Figure 4-144 Managing Arrays

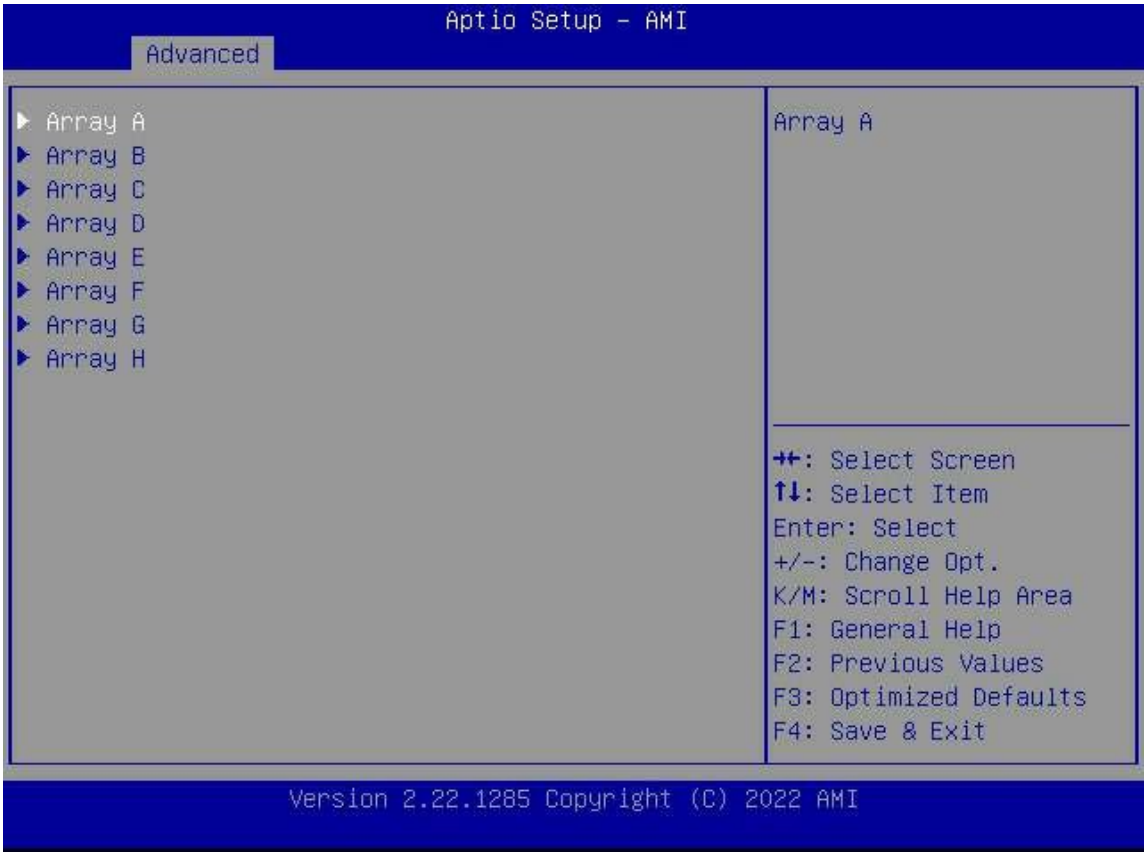

5. Use the arrow keys to select the array for which the caching function needs to be enabled, and then press **Enter**. The screen for managing the selected array is displayed, see Figure 4-145.

Figure 4-145 Managing the Selected Array

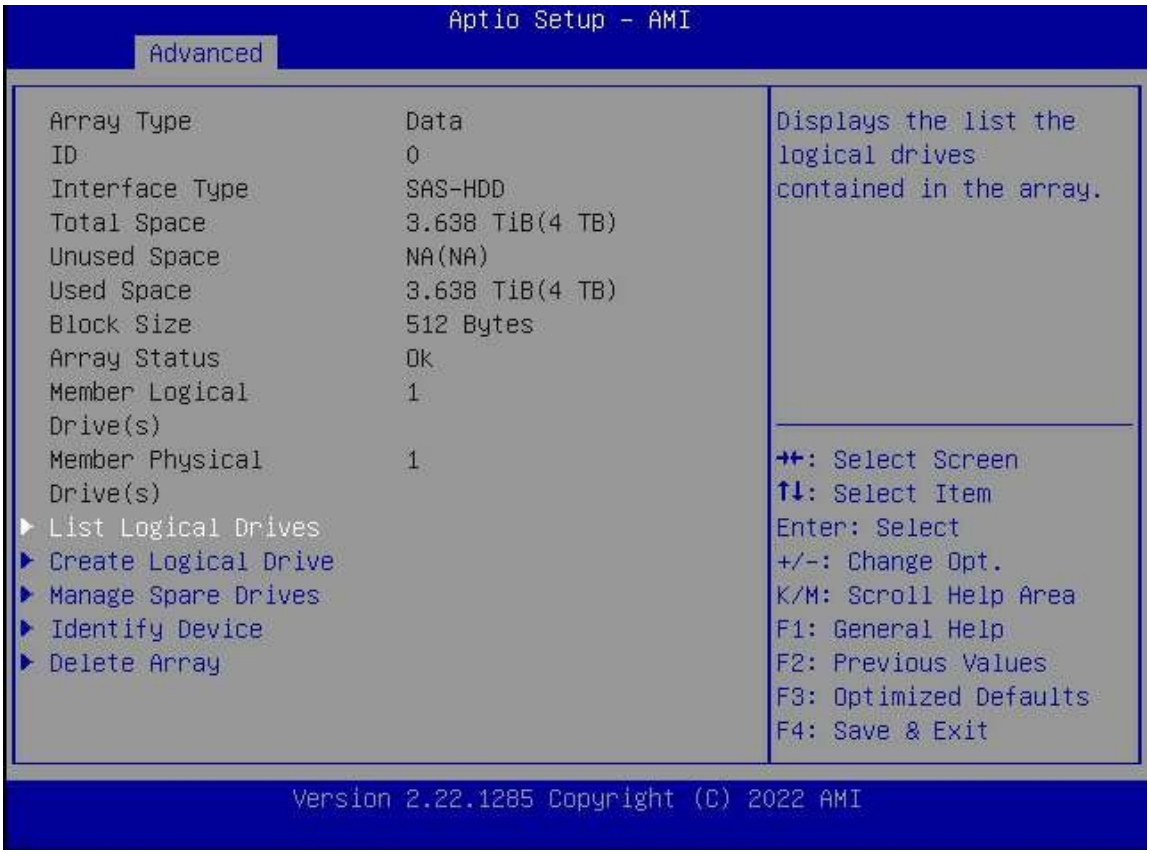

6. Use the arrow keys to select **List Logical Drives,** and then press **Enter**. The screen for managing RAID logical volumes is displayed, see Figure 4-146.

Figure 4-146 Managing Logical Volumes

7. Use the arrow keys to select the RAID logical volume for which the caching function needs to be enabled, and then press **Enter**. The screen for managing the selected RAID logical volume is displayed, see Figure 4-147.

Figure 4-147 Managing the Selected Logical Volume

8. Use the arrow keys to select **Logical Drive Details**, and press **Enter**. The RAID volume information is displayed, see Figure 4-148. The value of **Acceleration Method** is **None**, indicating that the caching function is disabled.

Figure 4-148 Viewing RAID Volume Information

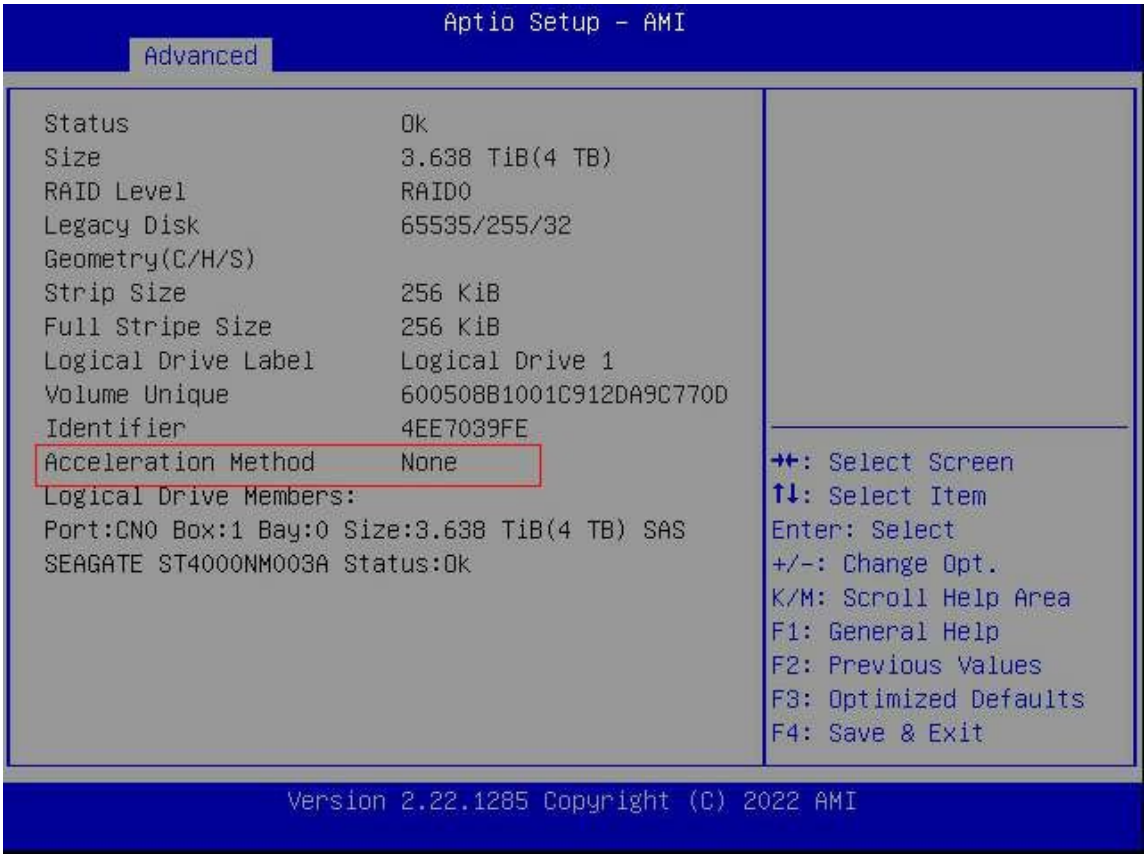

9. Press **Esc** to return to the RAID volume management screen. Use the arrow keys to select **Edit Logical Drive** and press **Enter**. The screen for setting RAID volume information is displayed, see Figure 4-149.

Figure 4-149 Setting RAID Volume Information

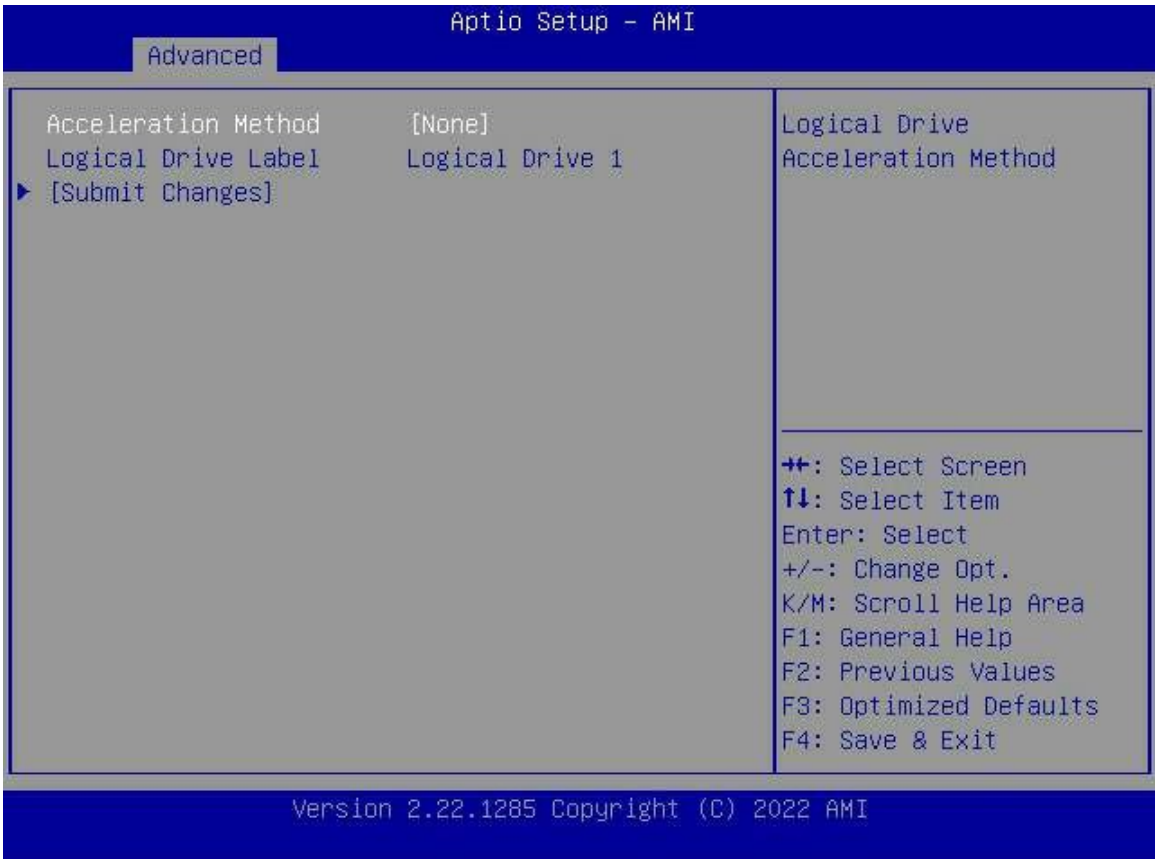

10. Use the arrow keys to select **Acceleration Method**, and then press **Enter**. The **Acceleration Method** dialog box is displayed, see Figure 4-150.

Figure 4-150 Acceleration Method Dialog Box

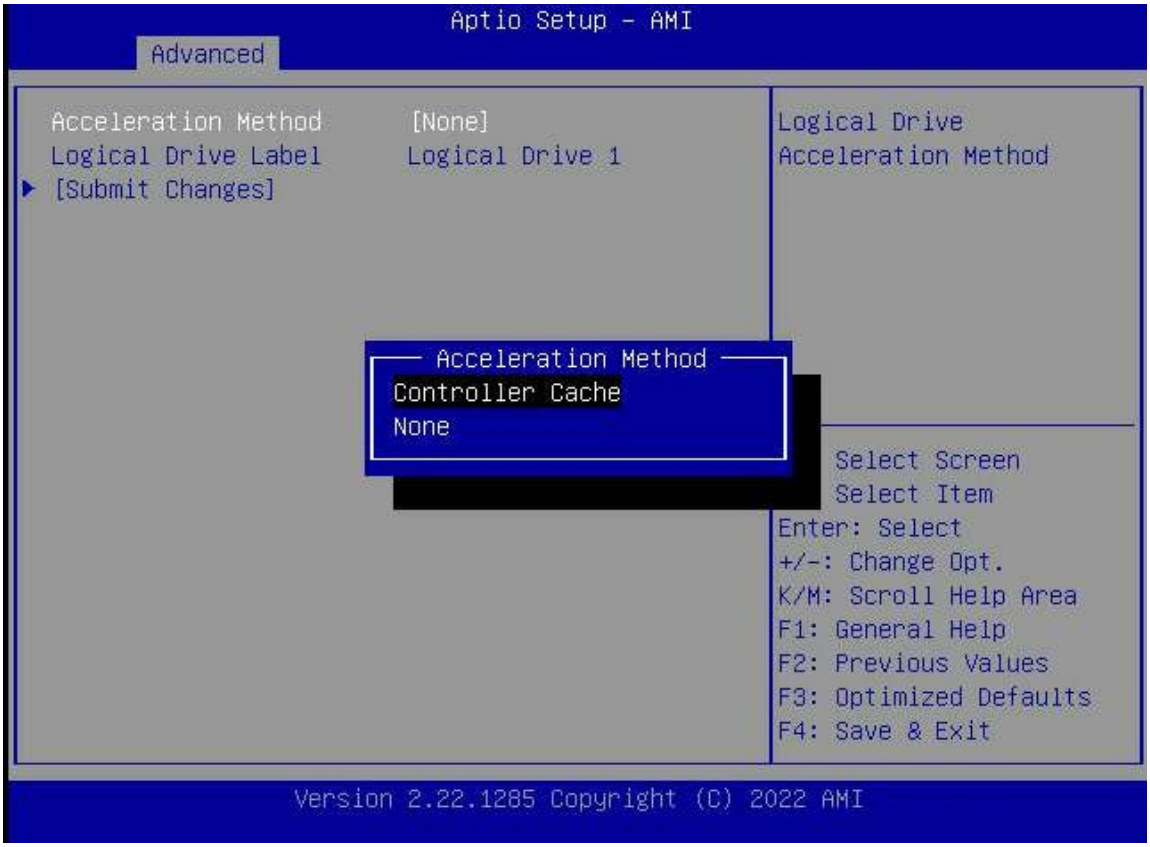

Note

There is no **I/O Bypass** option in **Acceleration Method** when HDDs are used. The **Acceleration Method** dialog box is displayed based on the actual hard disk configuration.

11. Use the arrow keys to select **Controller Cache**, and then press **Enter**. The caching mode is set to **Controller Cache**, see Figure 4-151.

Figure 4-151 Setting the Caching Mode

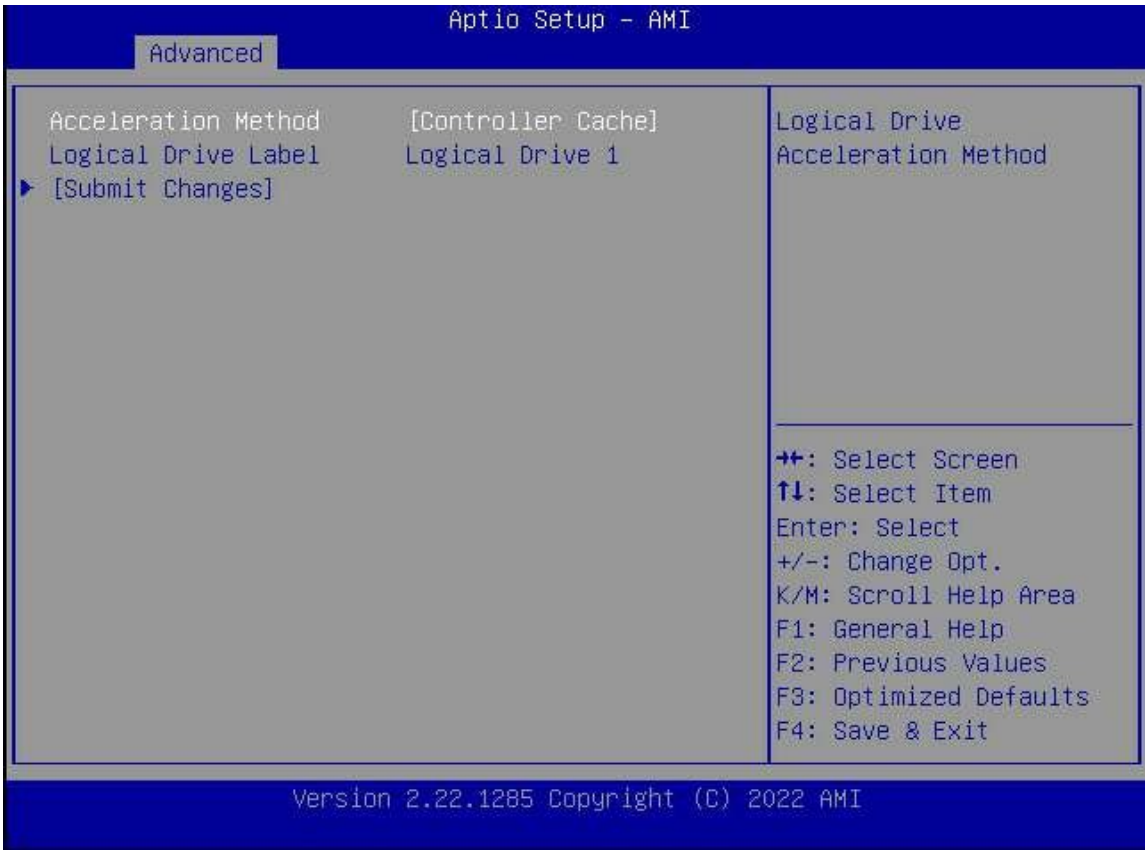

12. Use the arrow keys to select **Submit Changes**, and press **Enter**. The caching mode is set successfully, see Figure 4-152.

Figure 4-152 Caching Mode Set Successfully

13. Select **Back to Main Menu**, and then press **Enter** to return to the controller management screen, see Figure 4-153.

Figure 4-153 Managing the Controller

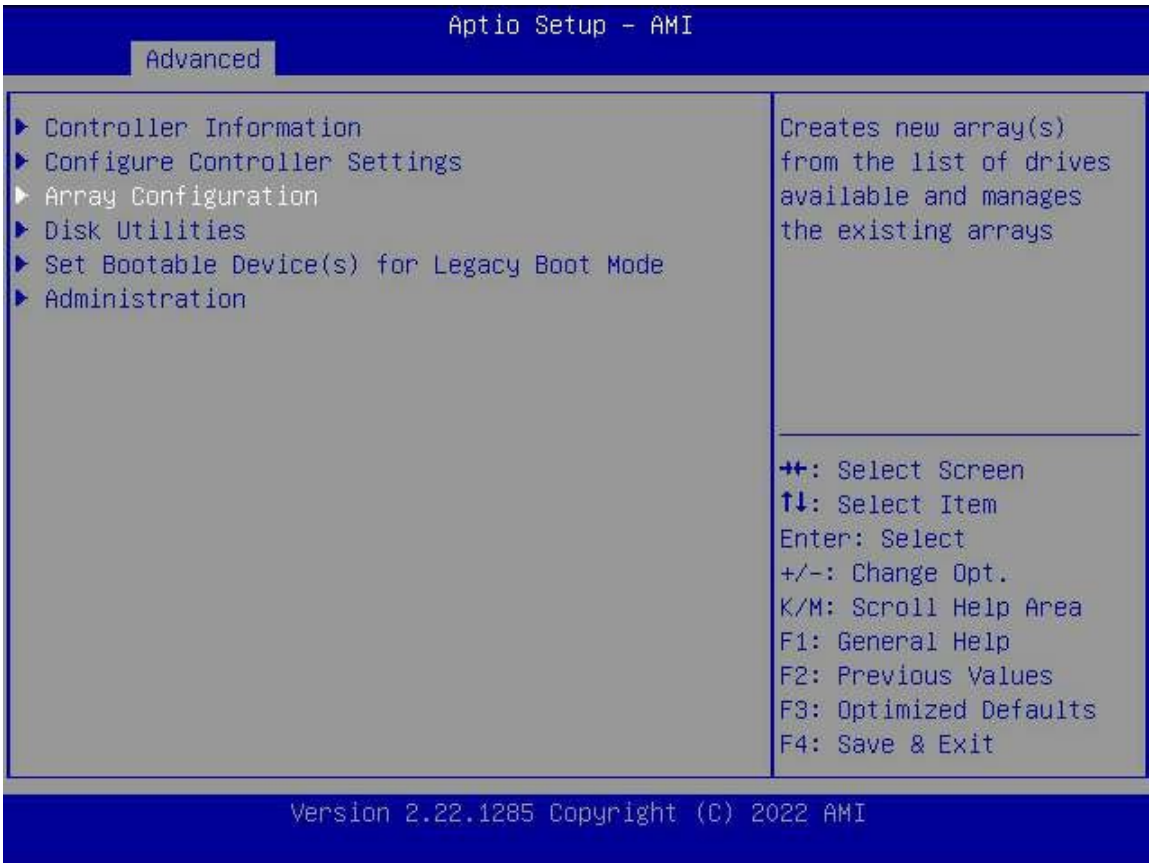

14. Repeat Step 3 through Step 8. The RAID volume information is displayed, see Figure 4-154. The value of **Acceleration Method** is **Controller Cache**, indicating that the caching function is enabled.

Figure 4-154 Viewing RAID Volume Information

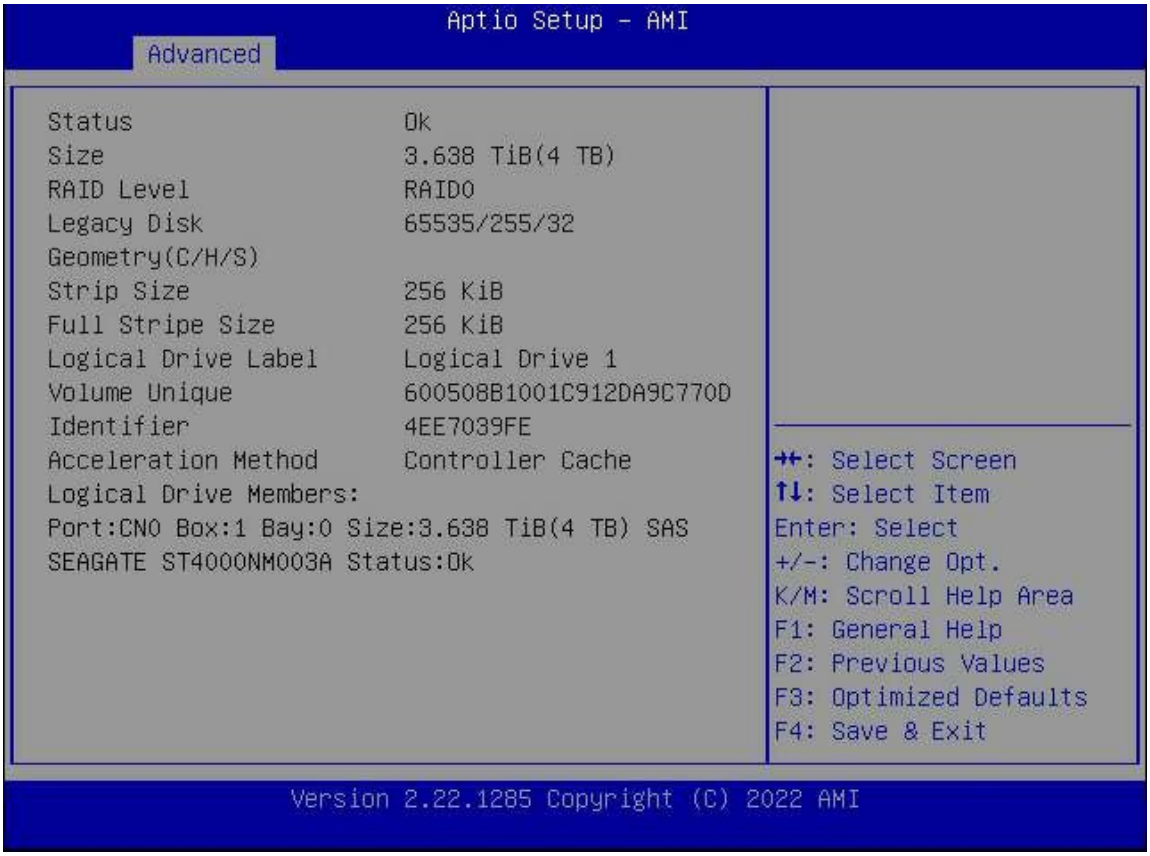

4.6 Typical Scenarios for Replacing a Disk (Legacy Mode)

For a description of the common scenarios for replacing a disk in a RAID volume on a SmartROC 3100 RAID card in legacy mode, refer to Table 4-12.

Table 4-12 Common Scenarios for Replacing a Disk in a RAID Volume on a SmartROC 3100 RAID Controller Card

Scenario	Description
Scenario 1	A newly inserted disk is converted into a RAID member disk. For details, refer to "4.6.1 Converting a Newly Inserted Disk Into a RAID Member Disk".
Scenario 2	After a faulty SmartROC 3100 RAID card is replaced, all the member disks in the RAID 1 array managed by the original faulty RAID card are moved to a new RAID card. For details, refer to “4.6.2 Moving All Member Disks of a RAID 1 Volume”.

4.6.1 Converting a Newly Inserted Disk Into a RAID Member Disk

Abstract

To convert a newly inserted disk into a RAID member disk, perform the following operations:

- The newly inserted disk is a foreign disk.

- The disk in a slot is removed and inserted back.

Inserting a Foreign Disk as a New Disk

When a disk in a RAID volume created on a server is faulty and needs replacement, remove the faulty disk from the disk slot on the server, and insert the prepared disk into the disk slot of the faulty disk.
After the disk is replaced, the configuration utility of the RAID controller card automatically synchronizes data on the newly inserted disk in the RAID volume.

Note

Data is automatically synchronized to the newly inserted disk no matter whether it carries RAID information or not.

Installing a Disk in the Original Slot After Removing It from the Slot

After a disk on a server is used to create a RAID volume, if the disk is removed from its slot and then inserted back, the RAID controller card configuration utility automatically rebuild the disk.

Note

The RAID 0 volume does not support the above functions.

4.6.2 Moving All Member Disks of a RAID 1 Volume

Abstract

If a SmartROC 3100 RAID controller card on a server fails and needs to be replaced, all the member disks in the RAID 1 on the faulty RAID controller card need to be moved to a new SmartROC 3100 RAID controller card.

Notice

It is risky to move the member disks of the RAID volume, and therefore it is recommended that you contact VANTAGEO technical support for help.

Steps

1. Shut down the server, and replace the faulty SmartROC 3100 RAID controller card with a new one.
2. Connect all member disks of the RAID 1 be moved to the new SmartROC 3100 RAID controller card.

3. Power on the server again and start the server system.
4. Start the BIOS configuration utility. For details, refer to “4.2.1 Starting the Configuration Utility”.
5. Query RAID volume information. For details, refer to “4.4.1 Querying RAID Volume Information”.
6. Contact VANTAGEO technical support to move member disks.

4.7 Typical Scenarios for Replacing a Disk (UEFI Mode)

For a description of the common scenarios for replacing a disk in a RAID volume on a SmartROC 3100 RAID card in UEFI mode, refer to Table 4-13.

Table 4-13 Common Scenarios for Replacing Disks in a RAID Volume on a SmartROC 3100 RAID Controller Card

Scenario	Description
Scenario 1	When a RAID 0 member disk is faulty, the RAID controller card is reconfigured. For details, refer to “4.7.1 A RAID 0 Member Disk Fails”.
Scenario 2	When a member disk of a logical volume with no hot spare disk configured is faulty, the faulty disk is replaced. For details, refer to "4.7.2 A Member Disk of a RAID Redundant Logical Volume (Without a Configured Hot Spare Disk) Fails".
Scenario 3	When a member disk of a logical volume with a hot spare disk configured is faulty, the faulty disk is replaced. For details, refer to "4.7.3 A Member Disk of a RAID Redundant Logical Volume (with a Configured Hot Spare Disk) Fails".

4.7.1 A RAID 0 Member Disk Fails

RAID 0 does not support data redundancy or backup. As a result, data cannot be restored after a fault occurs in the RAID 0 logical volume. It is necessary to install a new disk and reconfigure RAID.

4.7.2 A Member Disk of a RAID Redundant Logical Volume (Without a Configured Hot Spare Disk) Fails

Abstract

If a fault occurs in a member disk of a redundant logical volume (with no hot spare disk configured) on a SmartROC 3100 RAID controller card, the SmartROC 3100 RAID controller card can automatically restore the data after the faulty disk is replaced with a new disk. During the process, the member disk status may be **OK**, but the logical disk status may be **Failed**. In this case, you need to restore the logical disk status.

Steps

1. On the controller management screen, use the arrow keys to select **Array Configuration**, and then press **Enter**. The screen for configuring an array is displayed.
2. Use the arrow keys to select **Manage Arrays**, and then press **Enter**. The screen for managing arrays is displayed.
3. Use the arrow keys to select the array for which you need to manage the logical volume, and then press **Enter**. The screen for managing the selected array is displayed.
4. Use the arrow keys to select **List Logical Drives**, and then press **Enter**. The screen for managing logical volumes is displayed.
5. Use the arrow keys to select the logical volume to be corrected, and then press **Enter**. The screen for managing the selected logical volume is displayed.
6. Use the arrow keys to select **Re-Enable Logical Drive**, and then press **Enter**. The screen for restoring logical volume status is displayed.
7. Press **Enter**. The status of the logical disk is restored.

Verification

On the logical volume management screen, use arrow keys to select **Logical Drive Details**, and then press **Enter**. The logical volume details are displayed. Verify that the logical volume status is **Ok**.

4.7.3 A Member Disk of a RAID Redundant Logical Volume (with a Configured Hot Spare Disk) Fails

When a fault occurs in a member disk of a redundant logical volume (with a hot spare disk configured) on a SmartROC 3100 RAID controller card, the RAID controller card automatically replaces the faulty disk with the hot spare disk and restores the data.

- When the hot spare disk is of the **Dedicated** type, the RAID controller card temporarily replaces the faulty disk with the hot spare disk and automatically restores the data. After the faulty disk is replaced with a new disk, the hot spare disk is restored to Hot Spare status.
- When the hot spare disk is of the **Auto Replace** type, the RAID controller card immediately replaces the faulty disk with the hot spare disk and automatically restores the data. After the new disk is inserted, the new disk becomes a hot spare disk.

Chapter 5
Appendixes

Table of Contents

5.1 ARCCONF CLI Tool

During the proper operation of a server, the ARCCONF CLI tool allows you to configure a RAID controller card without restarting the server.

5.1.1 Downloading and Installing the ARCCONF Tool

Abstract

Before using the ARCCONF tool to operate a RAID controller card, you must download and install the ARCCONF tool.

Steps

Downloading the ARCCONF Tool

1. On the address bar of the browser on the maintenance PC, enter https://vantageo.com/tools-detail.html?id=9, and then press **Enter**. The page for downloading the tool is displayed.
2. Click the latest version. The page for downloading it is displayed.

Note

For example, the filename of the downloaded tool is "*LogTool_V01.20.01.02.zip*," where "V01.20.01.02" is the tool version number and vary with the update of the tool version.
This manual uses the tool of the *LogTool_V01.20.01.02.zip* version as an example to describe how to download and install the ARCCONF tool.

3. Click . The latest version of the tool is downloaded to the maintenance PC.

Installing the ARCCONF Tool

4. Decompress the *LogTool_V01.20.01.02.zip* tool package to the *LogTool_V01.20.01.02* folder, and then enter this folder.
5. Depending on the server operating system, perform the following operations:

If…	Then…
The OS is Linux	a. Decompress the *LogCollect_linux.tar.gz* package in the *LogTool_V01.20.01.02* folder to the *LogCollect_linux* folder. Access the *LogCollect_linux\components\raid\tools* directory. *Arcconf-linux* in the *tools* folder is the ARCCONF tool. b. Connect the maintenance PC to the server directly through a network cable, and upload the *arcconf-linux* file to any directory of the server OS through the SSH tool on the maintenance PC. c. Rename the *arcconf-linux* file as *arcconf*. d. Grant the executable permission to the *arcconf* files.
The OS is Windows	a. Decompress the *LogCollect_windows.7z* package in the *LogTool_V01.20.01.02* folder to the *LogCollect_windows* folder. Access the *LogCollect_windows\components\raid\tools\x86_64* directory. *Arcconf64.exe* in the *x86_64* folder is the ARCCONF tool. b. On the maintenance PC, upload the *arcconf64.exe* file to the *C:\Windows\System32* directory of the server OS through the SSH tool. c. Rename the *arcconf64.exe* file as *arcconf.exe*.
The OS is VMware	a. Use a file transfer tool (for example, Putty) to upload the tool package applicable to VMware to any directory of the server OS. The following uses */tmp* as an example. b. Run the following command to install the ARCCONF tool: **esxcli software vib install -v=***/tmp/vmware-esx-provider-arcconf.vib --no-sig-check* In the above command, */tmp/vmware-esx-provider-arcconf.vib* must be the complete path of the ARCCONF tool file. c. Enter the */opt/pmc* (installation path of the ARCCONF tool) directory and execute the commands related to the RAID controller card.

Note

If the OS is Linux, the ARCCONF tool can be used only after the executable permission is granted to it. The following command is used to grant the highest permission. You must exercise caution when executing the following command:

 # chmod 777 *arcconf*

5.1.2 Common Commands

When you configure a RAID controller card by using the ARCCONF CLI tool, it is recommended that you comply with the initial configuration flow shown in Figure 5-1.

Figure 5-1 Initial Configuration Flow

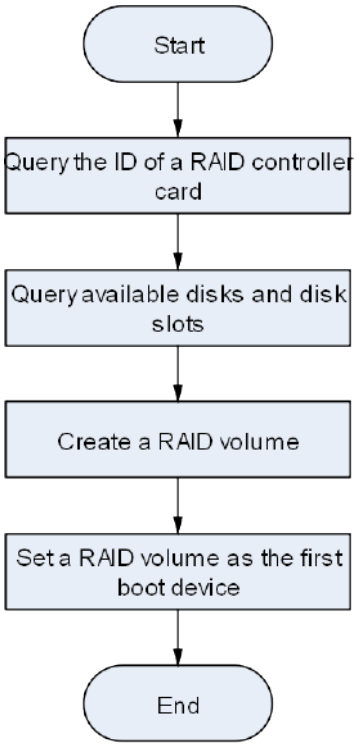

Note

- The ID of a queried RAID controller card is used as a parameter for executing other common commands. For the commands, refer to "5.1.2.1 Querying the Basic Information About a RAID Controller Card".
- The queried available disks and disk slots are used as the parameters of the RAID creation command. For the command, refer to "5.1.2.6 Creating and Deleting a RAID Array".
- A RAID volume is set as the first boot device. For the command, refer to "5.1.2.7 Setting a Boot Device".

5.1.2.1 Querying the Basic Information About a RAID Controller Card

Command Function

This command displays all the RAID controller cards identified by the ARCCONF tool or the basic information about the specified RAID controller card.

Syntax

- arcconf list
- arcconf list [Controller#]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-

Usage Guidelines

None

Example

- The following example shows how to query the IDs of all RAID controller cards identified by the ARCCONF tool:

  # arcconf list

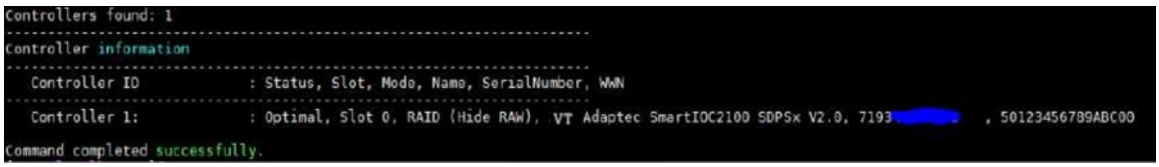

For a description of the fields in the command output, refer to Table 5-1.

Table 5-1 Descriptions of the Fields in the Command Output

Field	Description
Controller ID	ID of a RAID controller card
Status	Status of a RAID controller card
Slot	Slot number of a RAID controller card
Mode	Work mode of a RAID controller card

- The following example shows how to query information about RAID controller card 1:

  # arcconf list 1

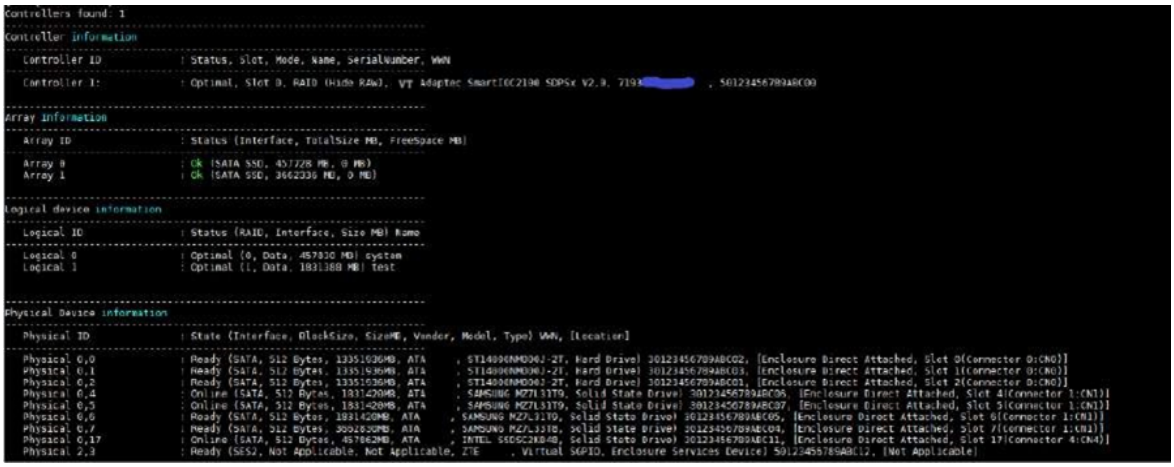

For a description of the fields in the command output, refer to Table 5-2.

Table 5-2 Descriptions of the Fields in the Command Output

Field	Description
Array ID	ID of the disk array
Logical ID	ID of the virtual disk
Physical ID	ID of the physical disk

5.1.2.2 Querying the Detailed Information About a RAID Controller Card

Command Function

This command displays detailed information about a RAID controller card, disk array, virtual disk, and physical disk.

Syntax

arcconf getconfig <Controller#> [AD | LD [LD#] | AR [AR#] | PD [Channel# ID# Channel# ID#...]| MC | CN | [AL]] [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
AD	Specifies that RAID controller card attributes are queried	-
LD	Specifies that virtual disk attributes are queried	-
[LD#]	ID of a virtual disk	-
[AR#]	ID of a disk array	-
PD	Specifies that physical disk attributes are queried	-
<Channel# ID#>	Channel ID (slot number) of a disk	-
MC	Specifies that the maximum cache information is queried	-
CN	Specifies that the RAID controller card information is queried	-
[AL]	Specifies that all device information is queried	-

Usage Guidelines

None

Example

- The following example shows how to query the information about a RAID controller card:

  # arcconf getconfig 1 ad

```
Controllers found: 1
----------------------------------------------------------------------
Controller information
----------------------------------------------------------------------
   Controller Status                        : Optimal
   Controller Mode                          : RAID (Hide RAW)
   Channel description                      : SCSI
   Controller Model                         : VT  Adaptec SmartIOC2100 SDPSx V2.0
   Controller Serial Number                 : 71934[redacted]
   Controller World Wide Name               : 50123456789ABC00
   Physical Slot                            : 0
   Temperature                              : 46 C/ 114 F (Normal)
   Host bus type                            : PCIe 3.0
   Host bus speed                           : 7880 MBps
   Host bus link width                      : 8 bit(s)/link(s)
   PCI Address (Bus:Device:Function)        : 0:3c:0:0
   Number of Ports                          : 5
   Internal Port Count                      : 5
   External Port Count                      : 0
   Defunct disk drive count                 : 0
   NCQ status                               : Enabled
   Queue Depth                              : Automatic
   Monitor and Performance Delay            : 60 minutes
   Elevator Sort                            : Enabled
   Degraded Mode Performance Optimization   : Disabled
   Latency                                  : Disabled
   Statistics data collection mode          : Disabled
   Post Prompt Timeout                      : 15 seconds
   Boot Controller                          : False
   Primary Boot Volume                      : None
   Secondary Boot Volume                    : None
   Driver Name                              : smartpqi
   Driver Supports SSD I/O Bypass           : Yes
   Manufacturing Part Number                : Not Applicable
   Manufacturing Spare Part Number          : Not Applicable
   Manufacturing Wellness Log               : Not Applicable
   NVRAM Checksum Status                    : Passed
   Sanitize Lock Setting                    : None
   -------------------------------------------------------
   Power Settings
   -------------------------------------------------------
   Power Consumption                        : Not Available
   Current Power Mode                       : Maximum Performance
   Pending Power Mode                       : Not Applicable
   Survival Mode                            : Enabled
   -------------------------------------------------------
   I2C Settings
   -------------------------------------------------------
   I2C Address                              : 0x0
   -------------------------------------------------------
   Cache Properties
   -------------------------------------------------------
```

```
-------------------------------------------------------
Physical Drive Write Cache Policy Information
-------------------------------------------------------
Configured Drives                          : Default
Unconfigured Drives                        : Default
HBA Drives                                 : Default
-------------------------------------------------------
RAID Properties
-------------------------------------------------------
Logical devices/Failed/Degraded            : 2/0/0
Automatic Failover                         : Disabled
Background consistency check               : High
Consistency Check Delay                    : 0 seconds
Parallel Consistency Check Supported       : Enabled
Parallel Consistency Check Count           : 1
Inconsistency Repair Policy                : Disabled
Consistency Check Inconsistency Notify     : Disabled
Rebuild Priority                           : Medium
Expand Priority                            : Medium
-------------------------------------------------------
Controller Version Information
-------------------------------------------------------
Firmware                                   : 2.66[0]
Driver                                     : Linux 1.1.4-115
Hardware Revision                          : A
-------------------------------------------------------
Temperature Sensors Information
-------------------------------------------------------
Sensor ID                                  : 0
Current Value                              : 37 deg C
Max Value Since Powered On                 : 46 deg C
Location                                   : Inlet Ambient

Sensor ID                                  : 1
Current Value                              : 46 deg C
Max Value Since Powered On                 : 56 deg C
Location                                   : ASIC

-------------------------------------------------------
```

For a description of the fields in the command output, refer to Table 5-3.

Table 5-3 Descriptions of the Fields in the Command Output

Field	Description
Controller Status	Status of the RAID controller card
Controller Mode	Work mode of the RAID controller card
Controller Model	Model of the RAID controller card

Field	Description
Firmware	Firmware version of the RAID controller card
Driver	Driver version of the RAID controller card

- The following example shows how to query the information about array 1:

  # arcconf getconfig 1 ar 1

```
Controllers found: 1
----------------------------------------------------------------------
Array Information
----------------------------------------------------------------------
Array Number 1
   Name                                    : B
   Status                                  : Ok
   Interface                               : SATA SSD
   Total Size                              : 3662336 MB
   Unused Size                             : 0 MB
   Block Size                              : 512 Bytes
   Array Utilization                       : 100.00% Used, 0.00% Unused
   Type                                    : Data
   Transformation Status                   : Not Applicable
   Spare Rebuild Mode                      : Dedicated
   SSD I/O Bypass                          : Enabled
   ----------------------------------------------------------
   Array Logical Device Information
   ----------------------------------------------------------
   Logical 1                               : Optimal (1, Data, 1831388 MB) test
   ----------------------------------------------------------
   Array Physical Device Information
   ----------------------------------------------------------
```

- The following example shows how to query the information about virtual disk 1:

  # arcconf getconfig 1 ld

```
Controllers found: 1
----------------------------------------------------------------------
Logical device information
----------------------------------------------------------------------
Logical Device number 0
   Logical Device name                     : system
   Disk Name                               : /dev/sdd
   Block Size of member drives             : 512 Bytes
   Array                                   : 0
   RAID level                              : 1
   Status of Logical Device                : Optimal
   Size                                    : 457830 MB
   Stripe-unit size                        : 256 KB
   Full Stripe Size                        : 256 KB
   Interface Type                          : SATA SSD
   Device Type                             : Data
   Boot Type                               : None
   Heads                                   : 255
   Sectors Per Track                       : 32
   Cylinders                               : 65535
   Caching                                 : Enabled
   Mount Points                            : /boot/efi 200 MB  Partition Number 1 /boot 1024 MB  Partition Number 2
   LD Acceleration Method                  : Controller Cache
   Volume Unique Identifier                : 600508B1001C1FF69AB000FC028B878C
   ----------------------------------------------------------
   Array Physical Device Information
   ----------------------------------------------------------
   Device 16                               : Present (457862MB, SATA, SSD, Channel:0, Device:16) PHYS7295[redacted]
   Device 17                               : Present (457862MB, SATA, SSD, Channel:0, Device:17) BTHC450[redacted]
```

For a description of the fields in the output displayed after the command for querying virtual disk 1 is executed, refer to Table 5-4.

Table 5-4 Description of the Fields in the Command Output

Parameter	Description
Logical Device name	Name of the virtual disk
Disk Name	Drive letter of the virtual disk
Array	ID of the virtual disk
RAID level	RAID level
Status of Logical Device	RAID status
Boot Type	Boot type of the virtual disk
Caching	Whether cache is enabled

- The following example shows how to query the information about the disk in slot 5:

  # arcconf getconfig 1 pd 0 5

```
Controllers found: 1
----------------------------------------------------------------------
Physical Device information
----------------------------------------------------------------------
      Device #0
         Device is a Hard drive
         State                              : Online
         Drive has stale RIS data           : False
         Block Size                         : 512 Bytes
         Physical Block Size                : 4K Bytes
         Transfer Speed                     : SATA 6.0 Gb/s
         Reported Channel,Device(T:L)       : 0,5(5:0)
         Reported Location                  : Enclosure Direct Attached, Slot 5(Connector 1:CN1)
         Array                              : 1
         Vendor                             : ATA
```

5.1.2.3 Querying Virtual Disk Information

Command Function

This command displays virtual disk information.

Syntax

arcconf getconfig <Controller#> LD [LD#]

Syntax Description

Parameter	Description	Range
<Controller#>	ID of the RAID controller card that the disks are connected to	-
LD	Specifies that virtual disk attributes are queried	-
[LD#]	ID of a virtual disk	-

Usage Guidelines

None

Example

The following example shows how to display the information about all virtual disks in RAID controller card 1:

arcconf getconfig 1 ld

```
Controllers found: 1
----------------------------------------------------------------------
Logical device information
----------------------------------------------------------------------
Logical Device number 0
   Logical Device name                     : system
   Disk Name                               : /dev/sdd
   Block Size of member drives             : 512 Bytes
   Array                                   : 0
   RAID level                              : 1
   Status of Logical Device                : Optimal
   Size                                    : 457830 MB
   Stripe-unit size                        : 256 KB
   Full Stripe Size                        : 256 KB
   Interface Type                          : SATA SSD
   Device Type                             : Data
   Boot Type                               : None
   Heads                                   : 255
   Sectors Per Track                       : 32
   Cylinders                               : 65535
   Caching                                 : Enabled
   Mount Points                            : /boot/efi 200 MB  Partition Number 1 /boot 1024 MB  Partition Number 2
   LD Acceleration Method                  : Controller Cache
   Volume Unique Identifier                : 600508B1001C1FF69AB000FC028B878C
   ----------------------------------------------------
   Array Physical Device Information
   ----------------------------------------------------
   Device 16                               : Present (457862MB, SATA, SSD, Channel:0, Device:16) PHYS7295[redacted]
   Device 17                               : Present (457862MB, SATA, SSD, Channel:0, Device:17) BTHC450[redacted]
```

For a description of the fields in the command output, refer to Table 5-5.

Table 5-5 Descriptions of the Fields in the Command Output

Field	Description
Logical Device name	Name of the virtual disk
Disk Name	Drive letter of the virtual disk
Array	ID of the virtual disk
RAID level	RAID level of the virtual disk
Status of Logical Device	RAID status
Boot Type	Boot type of the virtual disk
Caching	Whether the caching function is enabled

5.1.2.4 Querying Physical Disk Information

Command Function

This command displays physical disk information.

Syntax

arcconf getconfig <Controller#> PD [Channel# ID# Channel# ID#...]

Syntax Description

Parameter	Description	Range
<Controller#>	ID of the RAID controller card that the disks are connected to	-
PD	Specifies that physical disk attributes are queried	-
<Channel# ID#>	Channel ID (slot number) of a physical disk	-

Usage Guidelines

None

Example

The following example shows how to display the information about the physical disk in slot 5:

arcconf getconfig 1 pd 0 5

5.1.2.5 Querying the Firmware/Driver Version of a RAID Controller Card

Command Function

This command displays the firmware/driver version of a RAID controller card.

Syntax

arcconf getversion [Controller#]

Syntax Description

Parameter	Description	Setting
[Controller#]	ID of a RAID controller card	-

Usage Guidelines

None

Example

The following example shows how to query the firmware/driver version of the RAID controller card whose ID is 1:

arcconf getversion 1

```
Controllers found: 1
Controller #1
==============
Firmware                                : 2.66[0] (0)
Driver                                  : Linux 1.1.4-115 (0)


Command completed successfully.
```

For a description of the fields in the command output, refer to Table 5-6.

Table 5-6 Descriptions of the Fields in the Command Output

Field	Description
Firmware	Firmware version of the RAID controller card
Driver	Driver version of the RAID controller card

5.1.2.6 Creating and Deleting a RAID Array

Command Function

The following commands create and delete a RAID array.

Syntax

- arcconf create <Controller#> logicaldrive [Options] <Size> <RAID#> <Channel# ID#> [Channel# ID#] ... [noprompt] [nologs]
- arcconf delete <Controller#> logicaldrive <Id#> [noprompt] [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
[Options]	Options for creating a RAID array	• stripesize: optional, specifying the stripe size in KB, range: 16, 32, 64, 128*, 256, 512, 1024. • legs: optional parameter, specifying the number of sub-logical devices.

Parameter	Description	Setting
		● name: optional, used to specify the name of the logical disk. ● method: initialization method of the logical device. Valid options include BUILD and DEFAULT. ● wcache: write cache setting of the logical device. Valid options include WT and WB. ● ldcache: sets the cache of the logical device. Valid options include LON and LOFF. ● ssdoverprovisioningoptimization <enable | disable>: specifies the initialization to support the fast parity check function. ● cachelinesize: specifies the size of the maximum cache. The default value is 64.
<Size>	Logical device size (MB). You can set it to all available space by using the MAX keyword.	-
<RAID#>	RAID level	RAID 0, RAID 1, RAID 1ADM, RAID 10, RAID 10ADM, RAID 5, RAID 6, RAID 50, RAID 60
<Channel# ID#>	Channel ID (slot number) of a disk	-
<ld#>	ID of a virtual disk	-

Usage Guidelines

By default, the disks without data are used to create a RAID array.

Example

- The following example shows how to query available disks. In the command output, the disk whose **State** is **Ready** can be used to create a RAID array, and the disk whose **State** is **Online** has been used to create a RAID array.

  # arcconf getconfig 1 pd

```
Controllers found: 1
----------------------------------------------------------------------
Physical Device information
----------------------------------------------------------------------
   Channel #0:
      Device #0
         Device is a Hard drive
         State                              : Online
         Drive has stale RIS data           : False
         Block Size                         : 512 Bytes
         Physical Block Size                : 4K Bytes
         Transfer Speed                     : SATA 6.0 Gb/s
         Reported Channel,Device(T:L)       : 0,0(0:0)
         Reported Location                  : Direct Attached, Slot 0(Connector 0:CN0)
         Array                              : 0
         Vendor                             : ATA
         Model                              : Micron_5300_MTFD
         Firmware                           : D3MU001
         Serial number                      : 21022CAFD844
         World-wide name                    : 30015EBE03850E41
         Reserved Size                      : 32768 KB
         Used Size                          : 457830 MB
         Unused Size                        : 0 MB
         Total Size                         : 457862 MB
         Write Cache                        : Disabled (write-through)
         S.M.A.R.T.                         : No
         S.M.A.R.T. warnings                : 0
         SSD                                : Yes
         Device On Boot Connector           : No
         NCQ supported                      : Supported
         NCQ status                         : Enabled
         Boot Type                          : None
         Current Temperature                : 28 deg C
         Maximum Temperature                : 47 deg C
         Threshold Temperature              : 70 deg C
         PHY Count                          : 1
         Drive Configuration Type           : Data
         Drive Exposed to OS                : False
         Sanitize Erase Support             : True
         Sanitize Lock Freeze Support       : True
         Sanitize Lock Anti-Freeze Support  : True
         Sanitize Lock Setting              : None
         Usage Remaining                    : 100 percent
         Estimated Life Remaining           : Not Applicable
         SSD Smart Trip Wearout             : False
         56 Day Warning Present             : False
         Drive Unique ID                    : 6C0D214B7B889236EEB49804E39D9427
         Drive SKU Number                   : Not Applicable
```

- The following example shows how to create RAID 1:

  # arcconf create 1 logicaldrive stripesize 64 name test max 1 0 4 0 5

```
Controllers found: 1

Do you want to add a logical device to the configuration?
Press y, then ENTER to continue or press ENTER to abort: y

Creating logical device: test

Command completed successfully.
```

- The following example shows how to delete the virtual disk whose ID is 1:

  # arcconf delete 1 logicaldrive 1

```
Controllers found: 1

WARNING: Deleting this logical device will automatically delete array 1 because it is the only logical device present on that array.
All data in logical device 1 will be lost.
Delete the logical device?
Press y, then ENTER to continue or press ENTER to abort: y

Deleting: logical device 1 (*test*)

Command completed successfully.
```

5.1.2.7 Setting a Boot Device

Command Function

The following two commands set a RAID volume or physical disk as a boot device.

- The configuration takes effect when the work mode of a RAID volume on a RAID controller card is mixed or RAID.
- The configuration takes effect when the work mode of a physical disk on a RAID controller card is HBA or mixed.

Syntax

- arcconf setboot <Controller#> logicaldrive <LogicalDrive#> type <Boot Type>
- arcconf setboot <Controller#> device <Channel# ID#> type <Boot Type>

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<LogicalDrive#>	ID of a virtual disk	-
<Boot Type>	Boot order	• Primary: sets the virtual disk or physical disk as the primary boot device. • Secondary: sets the virtual disk or physical disk as the secondary boot device.
<Channel# ID#>	ID of a physical disk	-

Usage Guidelines

None

Example

- The following example shows how to set virtual disk 0 as the primary boot device:
  # arcconf setboot 1 logicaldrive 0 type Primary

```
Controllers found: 1
Do you wish to change the boot type of the selected device.
Press y, then ENTER to continue or press ENTER to abort: y


Command completed successfully.
```

- The following example shows how to set the disk in slot 6 as the secondary boot device:
  # arcconf setboot 1 device 0 6 type Secondary

```
Controllers found: 1
Do you wish to change the boot type of the selected device.
Press y, then ENTER to continue or press ENTER to abort: y


Command completed successfully.
```

5.1.2.8 Setting a Hot Spare Disk

Command Function

This command sets a global or dedicated hot spare disk.

Syntax

arcconf setstate <Controller#> device <Channel# ID#> <State> [ARRAY <AR#> [AR#] ...] [SPARETYPE <TYPE>] [noprompt] [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<Channel# ID#>	Channel ID (slot number) of a disk	-
<State>	Adds or deletes a hot spare disk	● HSP: adds a hot spare disk for one or more arrays ● RDY: deletes a hot spare disk
[ARRAY <AR#> [AR#] ...]	ID of a virtual disk	-
[SPARETYPE <TYPE>]	Sets the replacement mode of the hot spare disk	● 1: dedicated spare, which can be shared between arrays ● 2: automatically replaced but cannot be shared between arrays

Usage Guidelines

None

Example

- The following example shows how to set the disk in slot 6 as a dedicated hot spare disk:
  # arcconf setstate 1 device 0 6 hsp array 1 sparetype 2

```
Controllers found: 1

Any existing hot-spare drives of a different sparetype will be removed.

Command completed successfully.
```

- The following example shows how to cancel the disk in slot 6 as a dedicated hot spare disk:
  # arcconf setstate 1 device 0 6 rdy

```
Controllers found: 1

Command completed successfully.
```

5.1.2.9 Setting the Work Mode of a RAID Controller Card

Command Function

This command sets the work mode of a RAID controller card.

Syntax

arcconf setcontrollermode <Controller#> <Controller Mode>

Syntax Descriptions

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<Controller Mode>	Work mode of the RAID controller card	• 2: indicates HBA mode. This mode does not allow RAID array creation. All hard disks are reported to the OS as raw disks. • 3: indicates RAID mode. The RAID controller card only reports the disks that form a RAID array to the OS. • 5: indicates Mixed mode. The RAID controller card reports that the RAID disks that form a RAID array. If a disk has no RAID configuration information, the RAID controller card reports to the OS as a raw disk. The OS can directly operate the disk.

Usage Guidelines

None

Example

- The following example shows how to set the work mode of a RAID controller card to RAID:
  # arcconf setcontrollermode 1 3

```
Controllers found: 1

Command completed successfully.
```

- The following example shows how to query the work mode of a RAID controller card:
  # arcconf getconfig 1

```
Controllers found: 1
----------------------------------------------------------------------
Controller information
----------------------------------------------------------------------
   Controller Status                       : Optimal
   Controller Mode                         : RAID (Hide RAW)
   Channel description                     : SCSI
   Controller Model                        : VT  Adaptec SmartIOC2100 SDPSx V2.0
```

5.1.2.10 Setting the Read/Write Cache Ratio of a RAID Controller Card

Command Function

This command sets the respective percentage of the read cache and write cache of a RAID controller card.

Syntax

arcconf setcache <Controller#> cacheratio read write

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
read	Percentage of the read cache	0–100
write	Percentage of the write cache	0–100

Usage Guidelines

The sum of the percentage of the read cache and that of the write cache is 100%.

Example

The following example shows how to set the percentage of the read cache to 50% and that of the write cache to 50%:
arcconf setcache 1 cacheratio 50 50

```
Controllers found: 1


Command completed successfully.
```

5.1.2.11 Setting the Caching Policy

Command Function

This command enables or disables the caching function for a virtual disk.

Syntax

- arcconf setcache <Controller#> logicaldrive <LogicalDrive#> con
- arcconf setcache <Controller#> logicaldrive <LogicalDrive#> coff

Syntax Description

Parameter	Description	Range
<Controller#>	ID of a RAID controller card	-
<LogicalDrive#>	ID of a virtual disk	-

Usage Guidelines

None

Example

- The following example shows how to enable the caching function for virtual disk 0:

  # arcconf setcache 1 logicaldrive 0 con

```
Controllers found: 1
Cache mode is already set to Enabled.


Command aborted.
```

- The following example shows how to disable the caching function for virtual disk 0:

  # arcconf setcache 1 logicaldrive 0 coff

```
Controllers found: 1


Command completed successfully.
```

5.1.2.12 Using Cache Without a Capacitor

Command Function

This command enables the use of the cache of a RAID controller card without a capacitor.

Notice

If the cache of the RAID controller card is used when there is no capacitor, data may be lost if the RAID controller card is powered off. Therefore, you must use this command with caution.

Syntax

arcconf setcache <Controller#> nobatterywritecache enable

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-

Usage Guidelines

None

Example

The following example shows how to enable the use of the cache of a RAID controller card without a capacitor:

arcconf setcache 1 nobatterywritecache enable

```
Controllers found: 1
WARNING : Enabling write caching on a cache module without a fully charged battery or
capacitor could cause data loss in the event of a power failure.

Are you sure you want to continue?
Press y, then ENTER to continue or press ENTER to abort: y

Command completed successfully.
```

5.1.2.13 Setting a Disk as a Pass-Through Disk

Command Function

After a disk is set as a pass-through disk, the OS can directly manage it.

Syntax

arcconf uninit <Controller#> [Channel# ID#] ... [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
[Channel# ID#]	Channel ID (slot number) of a disk	-

Usage Guidelines

None

Example

The following example shows how to set the disk in slot 5 as a pass-through disk:

arcconf uninit 1 0 5

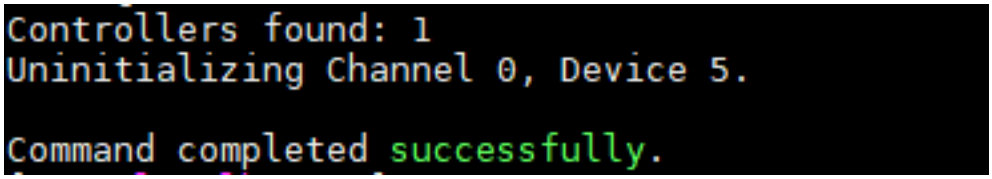

5.1.2.14 Querying and Setting the Write Cache Policy for Disks

Command Function

The following commands query and set the write cache policy for disks.

Syntax

- arcconf getconfig <Controller#> logicaldrive ad
- arcconf setcache <Controller#> drivewritecachepolicy <drivetype> <cachepolicy> <drivetype> <cachepolicy>...

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<drivetype>	Type of disk for which you want to set the write cache policy	• Configured: sets the write cache policy for RAID member disks in RAID or mixed mode. • Unconfigured: sets the write cache policy for non-RAID member disks in RAID or mixed mode. • HBA: set the write cache policy for disks in HBA mode.
<cachepolicy>	Write cache policy	• 0: Default (keeps the default status) • 1: Enabled (enables write caching) • 2: Disabled (disables write caching)

Usage Guidelines

None

Example

- The following example shows how to set the write cache policy for disks in **Configured** status to **Enabled** and set the write cache policy for disks in **Unconfigured** status to **Disabled**:

arcconf setcache 1 drivewritecachepolicy Configured 1 Unconfigured 2

```
Controllers found: 1
Enabling controller drive write cache can increase write performance but risks losing
the data in the cache on sudden loss.
Command completed successfully.
```

- The following example shows how to query the write cache policy for disks:

  # arcconf getconfig 1 ad

```
Controllers found: 1
----------------------------------------------------------------------
Controller information
----------------------------------------------------------------------
   Controller Statue              : Optimal
   Controller Mode                :Mixed
   Channel description            :SCSI

...
----------------------------------------------------------------------
Physical Drive Write Cache Policy Information
----------------------------------------------------------------------
Configured Drives                 :Enabled
Unconfigured Drives               :Disabled
HBA Drives                        :Default
----------------------------------------------------------------------
...
```

5.1.2.15 Setting the Status of a Disk Locating Indicator

Command Function

This command turns on or off a disk locating indicator.

Syntax

- arcconf identify <Controller#> all [TIME <BlinkTime>] [STOP] [nologs]
- arcconf identify <Controller#> logicaldrive <logicaldrive#> [TIME <BlinkTime>] [nologs]
- arcconf identify <Controller#> device <Channel# ID#> [TIME <BlinkTime>] [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
[TIME <BlinkTime>]	Time period (in seconds) that the locating indicator is lit for	-
[STOP]	Specifies that the locating indicator is off	-
<LogicalDrive#>	ID of a virtual disk	-
<Channel# ID#>	Channel ID (slot number) of a physical disk	-

Usage Guidelines

None

Example

- The following example shows how to light the locating indicators of disks in RAID group 1 for 5 seconds:

  # arcconf identify 1 logicaldrive 1 time 5

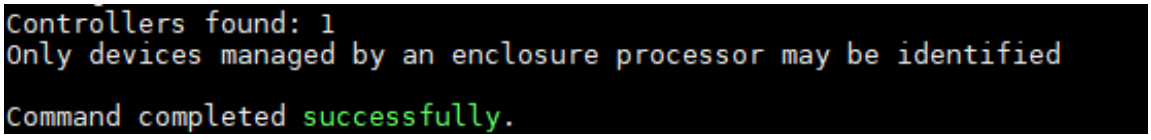

- The following example shows how to light the locating indicators of disks in slot 7 for 30 seconds:

  # arcconf identify 1 device 0 7 time 30

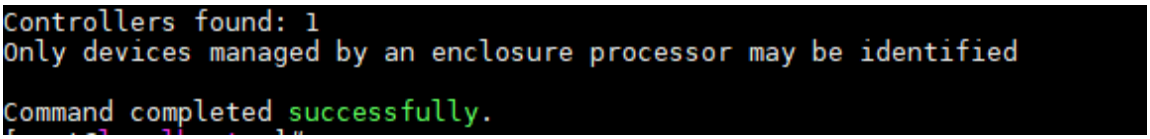

- The following example shows how to turn off the locating indicators of disks in all slots:

  # arcconf identify 1 all stop

5.1.2.16 Modifying the RAID Stripe Size

Command Function

This command modifies the RAID stripe size.

Syntax

arcconf modify <Controller#> from <LogicalDrive#> to stripesize <STRIPE> <Size> <RAID#> <Channel# ID#> [Channel# ID#] ... [noprompt] [nologs]

Syntax Description

Parameter	Description	Range
<Controller#>	ID of a RAID controller card.	-
<LogicalDrive#>	ID of the virtual disk to be set.	-
<STRIPE>	Stripe size to be set.	16, 32, 64, 128, 256, 512, and 1024. Unit: KB.
<Size>	Size of the virtual disk. The **max** parameter can be used to specify all the available space of the virtual disk.	Unit: MB.
<RAID#>	RAID level of the virtual disk.	-
<Channel# ID#> [Channel# ID#]	Channel ID (slot number) of the virtual disk.	-

Usage Guidelines

- The **noprompt** parameter is used in the command to indicate forcible execution.
- If the **nologs** parameter is specified, this command is not executed when data is stored in the virtual disk.
- The stripe size and RAID capacity cannot be modified at the same time.
- When the RAID capacity is set to the maximum value and the stripe size needs to be modified, the **max** parameter cannot be used in the command to represent a number.

Note

For example, the **Changing the size and the stripe size simultaneously is not allowed** message is displayed if the following command is executed to modify the stripe size:
arcconf modify 1 from 1 to stripesize 256 max 0 0 3
The following command can be executed to modify the stripe size:
arcconf modify 1 from 1 to stripesize 256 915683 0 0 3

Example

The following example shows how to modify the stripe size to 128 KB without adding any hard disk:

arcconf modify 1 from1 to stripesize 128 1831388 1 0 4 0 5

```
Controllers found: 1
Reconfiguration of a logical device is a long process.  Are you sure you want to continue?
Press y, then ENTER to continue or press ENTER to abort: y

Reconfiguring logical device: test

Command completed successfully.
```

5.1.2.17 Setting the Initialization Function for a Physical Disk

Command Function

The following two commands set the initialization function for a physical disk and query the task status of the physical disk.

Syntax

- arcconf task start <Controller#> device <Channel# ID#> <task> [PATTERN <erasePattern>][noprompt]
- arcconf task stop <Controller#> device <Channel# ID#>

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<Channel# ID#>	Channel ID (slot number) of a disk	-
<task>	Task type	secureerase: securely erases disk data
[PATTERN <erasePattern>]	Initialization type	● 1: zero initialization, which initializes all blocks as zero ● 2: random zero initialization, which randomly initializes blocks as zero ● 2: random zero initialization, which randomly initializes blocks as zero ● 4: enables the physical devices (HDDs and SSDs) to change the encryption key to prevent correct decryption of previously stored information. ● 5: block erasing and clearing method, applicable to SSDs only

Parameter	Description	Setting
		● 6: overwrites user data with specific pattern data, applicable to HDDs only

Usage Guidelines

Because the command uses the **noprompt** parameter to indicate forcible execution, which omits the reconfirmation operation during the setting process. Therefore, the **noprompt** parameter is not recommended.

Example

- The following example shows how to initialize the disk in slot 5:
  # arcconf task start 1 device 0 5 secureerase

```
Controllers found: 1
Secure erase of a Hard drive is a long process.
Warning: Rebooting the machine would stop the Secure erase task.

Are you sure you want to continue?
Press y, then ENTER to continue or press ENTER to abort: y

Secure Erasing Channel 0, Device 5.

Command completed successfully.
[root@localhost home]#
```

- The following example shows how to query task status:
  # arcconf getstatus 1

```
Controllers found: 1
Physical Device Task:
   Channel,Device                  : 0,5
   Task ID                         : 100
   Current operation               : Secure Erase(Zero)
   Status                          : In Progress
   Priority                        : High
   Percentage complete             : 0


Command completed successfully.
[root@localhost home]#
```

5.1.2.18 Setting the Priority of a Background Task

Command Function

This command sets and applies the priority of a background task.

Syntax

- arcconf setpriority <Controller#> priority [current]
- arcconf setpriority <Controller#> [Task ID#] <New Priority> [current]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
[Task ID#]	Task type	• rebuild: rebuild • expand: expansion
<New Priority>	Task priority	• high: high priority • medium: medium priority • low: low priority

Usage Guidelines

None

Example

The following example shows how to set and apply the priority of a background rebuild task to medium:

arcconf setpriority 1 rebuild medium

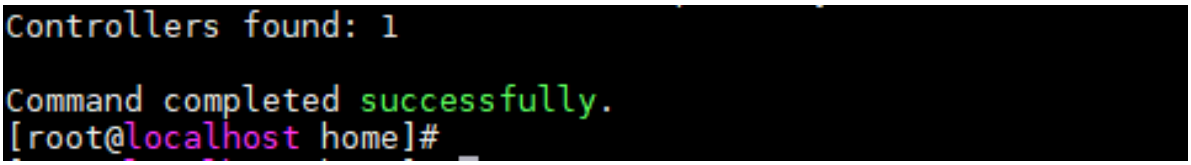

5.1.2.19 Querying Disk Array Rebuild Status

Command Function

This command queries the rebuild status of disk arrays on a RAID controller card.

Syntax

arcconf getstatus [Controller#]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-

Usage Guidelines

None

Example

The following example shows how to query the rebuild status of disk arrays on RAID controller card 1:

arcconf getstatus 1

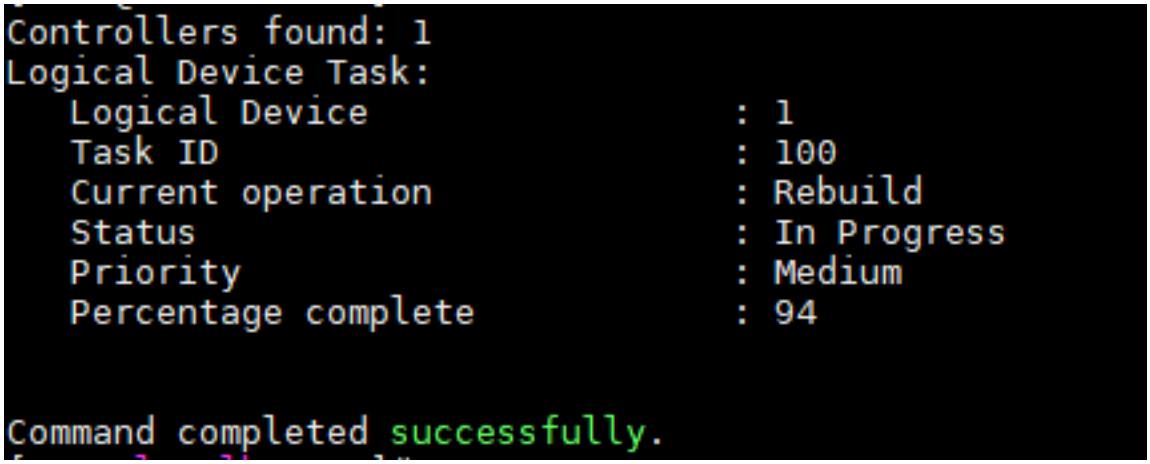

For a description of the fields in the command output, refer to Table 5-7.

Table 5-7 Description of the Fields in the Command Output

Field	Description
Logical Device	ID of the virtual disk
Current operation	Current operation
Status	Status
Priority	Priority
Percentage complete	Completion percentage

5.1.2.20 Setting Consistency Check Status

Command Function

This command enables the consistency check function.

Syntax

arcconf consistencycheck <Controller#> <On [Delay]|Off> [noprompt] [nologs]

Syntax Description

Parameter	Description	Setting
<Controller#>	ID of a RAID controller card	-
<On [Delay]|Off>	Consistency check status	• on: enabled • off: disabled

Usage Guidelines

None

Example

The following example shows how to enable the automatic consistency check function:

arcconf consistencycheck 1 on

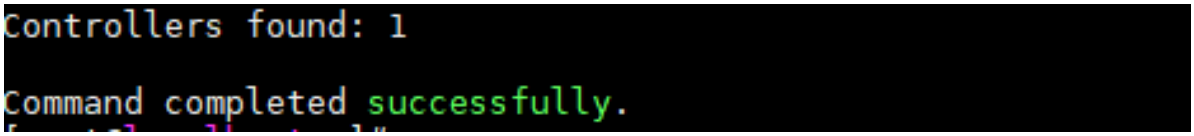

5.1.2.21 Upgrading Firmware

Command Function

This command upgrades the firmware of a RAID controller card.

Syntax

arcconf romupdate controller_id fwfile

Syntax Description

Parameter	Description	Setting
controller_id	ID of a RAID controller card	-
fwfile	Name of the .bin file required for firmware upgrade	-

Usage Guidelines

None

Example

The following example shows how to upgrade the firmware of a RAID controller card:

arcconf romupdate 1 SmartFWx100.bin

```
Controllers found: 1
```

```
Are you sure you want to continue?
Press y, then ENTER to continue or press ENTER to abort: y

Updating controller 1 firmware...
Succeeded
You must restart the system for firmware updates to take effect.

Command completed successfully.
```

5.1.2.22 Collecting Firmware Logs

Command Function

This command collects firmware logs of a RAID controller card.

Syntax

arcconf savesupportarchive

Syntax Description

None

Usage Guidelines

None

Example

The following example shows how to collect firmware logs of a RAID controller card:

arcconf savesupportarchive

```
Controllers found: 1

Usage: SAVESUPPORTARCHIVE [Path] [Firmware|GUI|Arcconf|Storlib|Basecode|Redfish]
Example: SAVESUPPORTARCHIVE C:\Adaptec\maxView Firmware
Example: SAVESUPPORTARCHIVE /var/log/maxView Storlib
Example: SAVESUPPORTARCHIVE
Example: SAVESUPPORTARCHIVE .
================================================================================

Saves  all the logs
    Path          : directory other than default
```

```
    Firmware      : saves Firmware logs
    GUI           : saves GUI logs
    Arcconf       : saves Arcconf logs
    StorLib       : saves StorLib logs
    Basecode      : saves Basecode logs
    Redfish       : saves Redfish logs

The path is not specified, Using the default.
The logs are saved in relevant folders in /var/log/Support

Command completed successfully.
```

5.2 Configuring RAID Through the VROC

Intel®VROC is a hybrid RAID storage solution for NVMe SSDs directly connected to CPUs. With the support of NVMe VROC Key, the VROC can create virtual RAID arrays on a CPU to implement the RAID function.

The advantage of the VROC is that it combines the RAID function and the strong performance of NVMe SSDs. That is, it can directly connect the NVMe SSDs to the PCIe channel of a CPU without the need to use the RAID HBA, avoiding the complexity and power consumption of traditional hardware.

5.2.1 Preparations

Before you configure RAID through the VROC, you need to make the following preparations:

- Make sure the cables, NVMe SSDs and NVMe VROC Key are installed correctly.
- Make sure that the server CPU supports the VROC function.
- Make sure that the model of the NVMe VROC Key matches the brand model of the NVMe SSDs. For the specifications of the NVMe VROC Key, refer to Table 5-8.

Table 5-8 NVMe VROC Key Specifications

Model	Brand	Supported RAID Level	Supported NVMe SSD
VROCPREMMOD 951606	Intel	RAID 0/1/5/10	Third-party NVMe SSDs except for those of the Intel brand
VROCSTANMOD 951605	Intel	RAID 0/1/10	Third-party NVMe SSDs except for those of the Intel brand
VROCISSDMOD 956822	Intel	RAID 0/1/5/10	NVMe SSDs of the Intel brand

Note

NVMe SSDs include those of the Intel brand and third-party ones. The NVMe SSDs are directly connected to CPUs through PCIe channels to implement high-rate read and write operations.

5.2.2 Enabling the VMD Function

Abstracts

VMD is an Intel NVMe SSD deployment solution. It supports online upgrade and replacement of NVMe SSDs from the PCIe bus on the server CPU. Therefore, if the VMD function is enabled, the operation and maintenance of the NVMe SSDs can be completed online without interrupting the kernel services.

Steps

1. Start the server system, and press the **F2**/**DEL** key as prompted. The **Aptio Setup** screen is displayed, see Figure 5-2.

Figure 5-2 Aptio Setup Screen

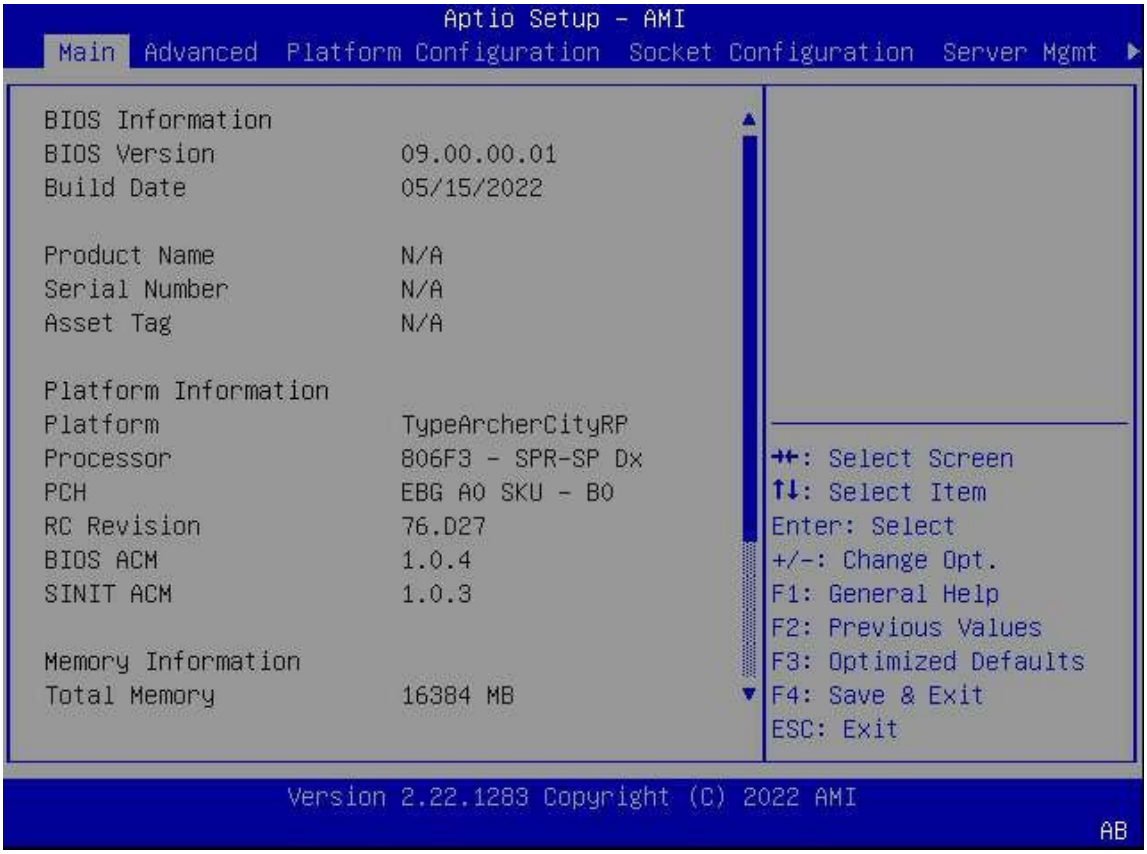

2. Use the arrow keys to select **Socket Configuration**, and press **Enter**. The **Socket Configuration** screen is displayed, see Figure 5-3.

Figure 5-3 Socket Configuration Screen

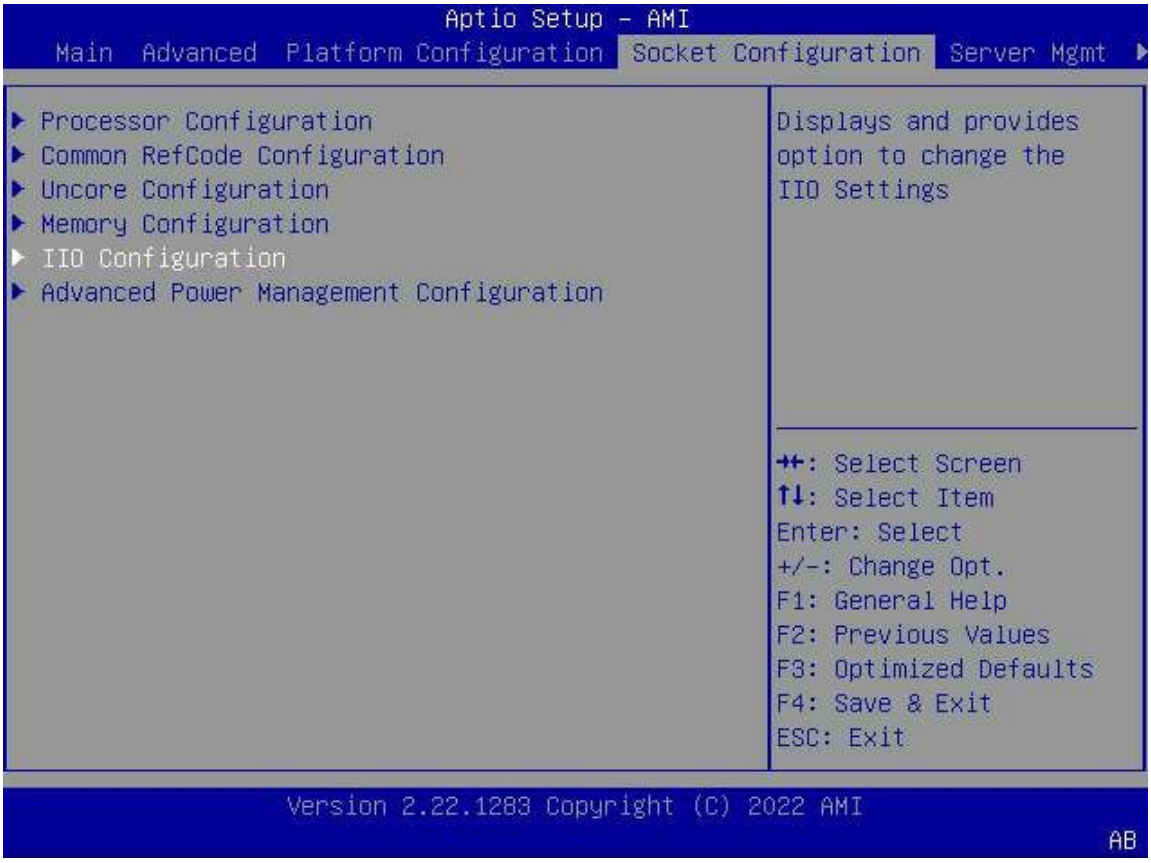

3. Use the arrow keys to select **IIO Configuration**, and then press **Enter**. The **IIO Configuration** screen is displayed, see Figure 5-4.

Figure 5-4 IIO Configuration Screen

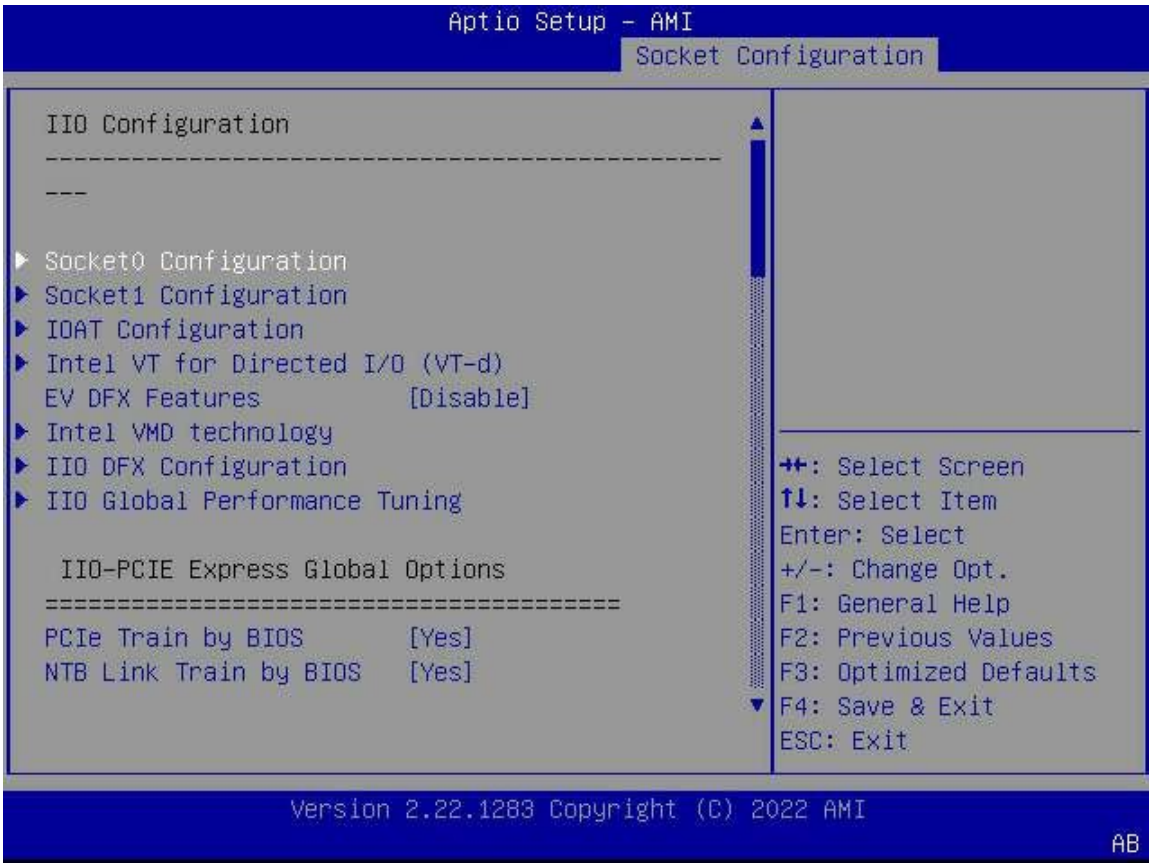

4. Use the arrow keys to select **Intel® VMD technology**, and then press **Enter**. The **Intel® VMD technology** screen is displayed, see Figure 5-5.

Figure 5-5 Intel® VMD Technology Screen

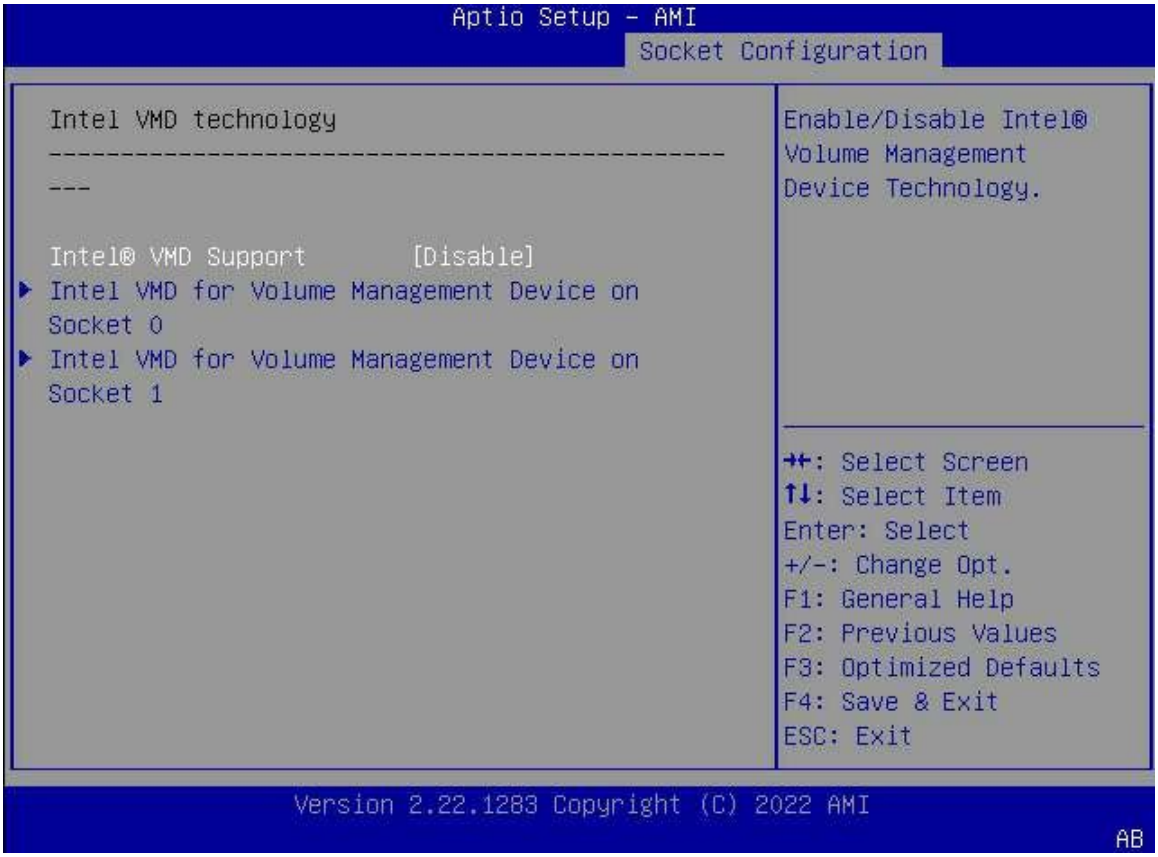

5. Use the arrow keys to select **Intel® VMD Support**, and press **Enter**. In the displayed dialog box, switch **Disabled** to **Enabled**, and press **Enter** to save the configuration, see Figure 5-6 .

Figure 5-6 Confirming the Configuration

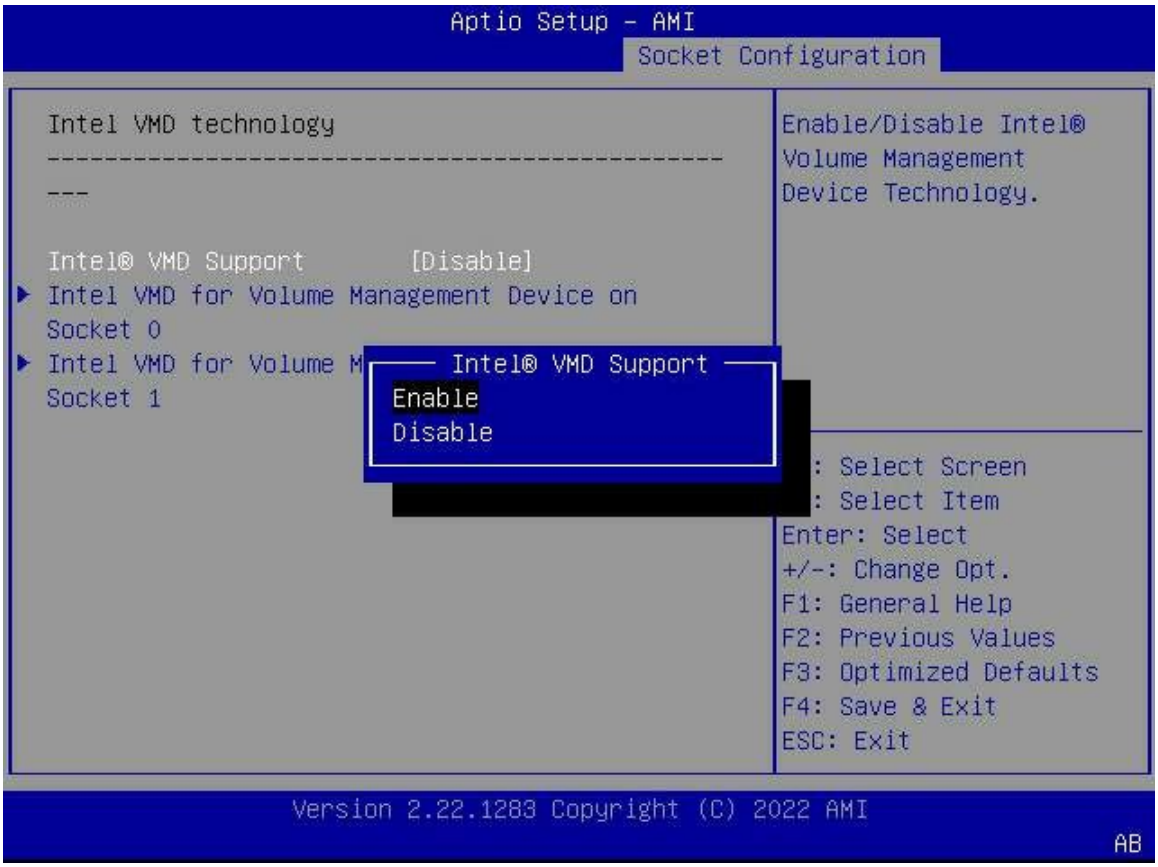

6. Press **Esc** multiple times until the **Socket Configuration** screen is displayed, see Figure 5-7.

Figure 5-7 Socket Configuration Screen

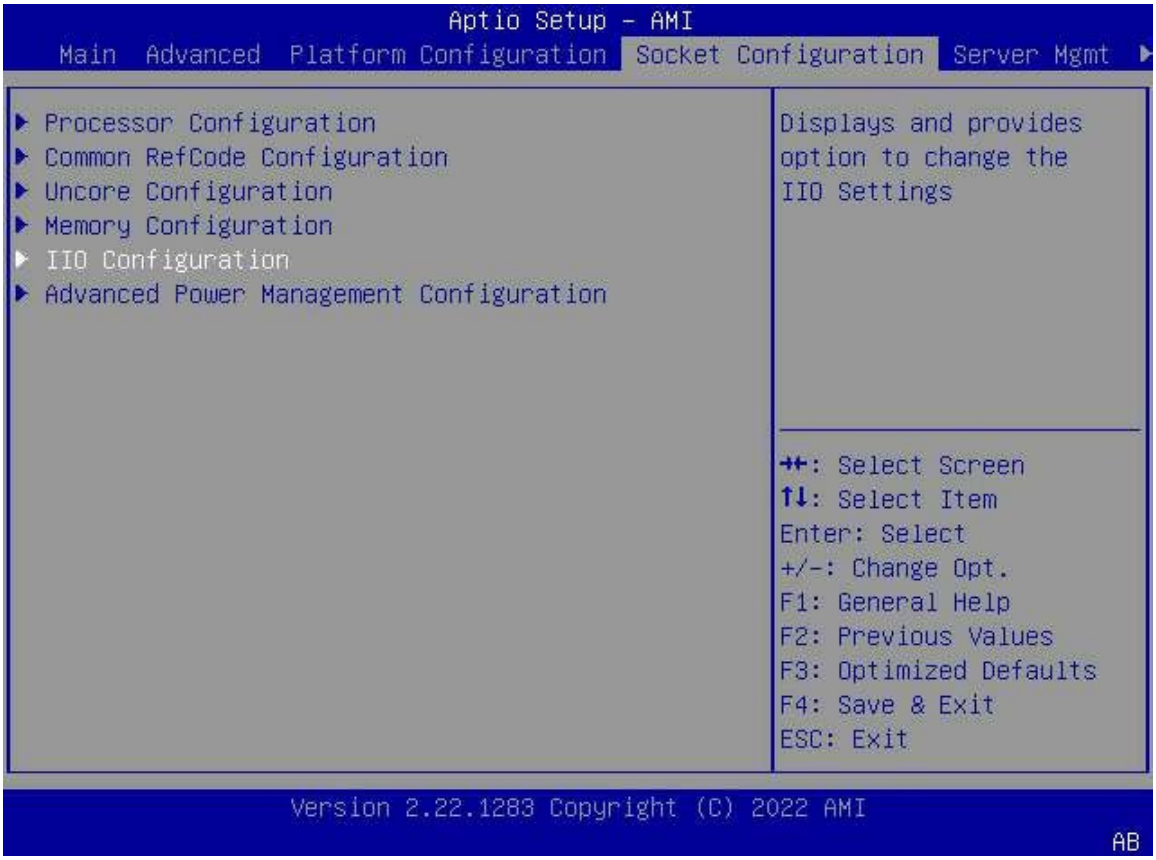

7. Press **F4**. In the displayed dialog box, select **Yes**, and press **Enter** to exit. The server system is restarted automatically
8. After power-on, press **F2**//**DEL**. The **Aptio Setup** screen is displayed, see Figure 5-8.

Figure 5-8 Aptio Setup Screen

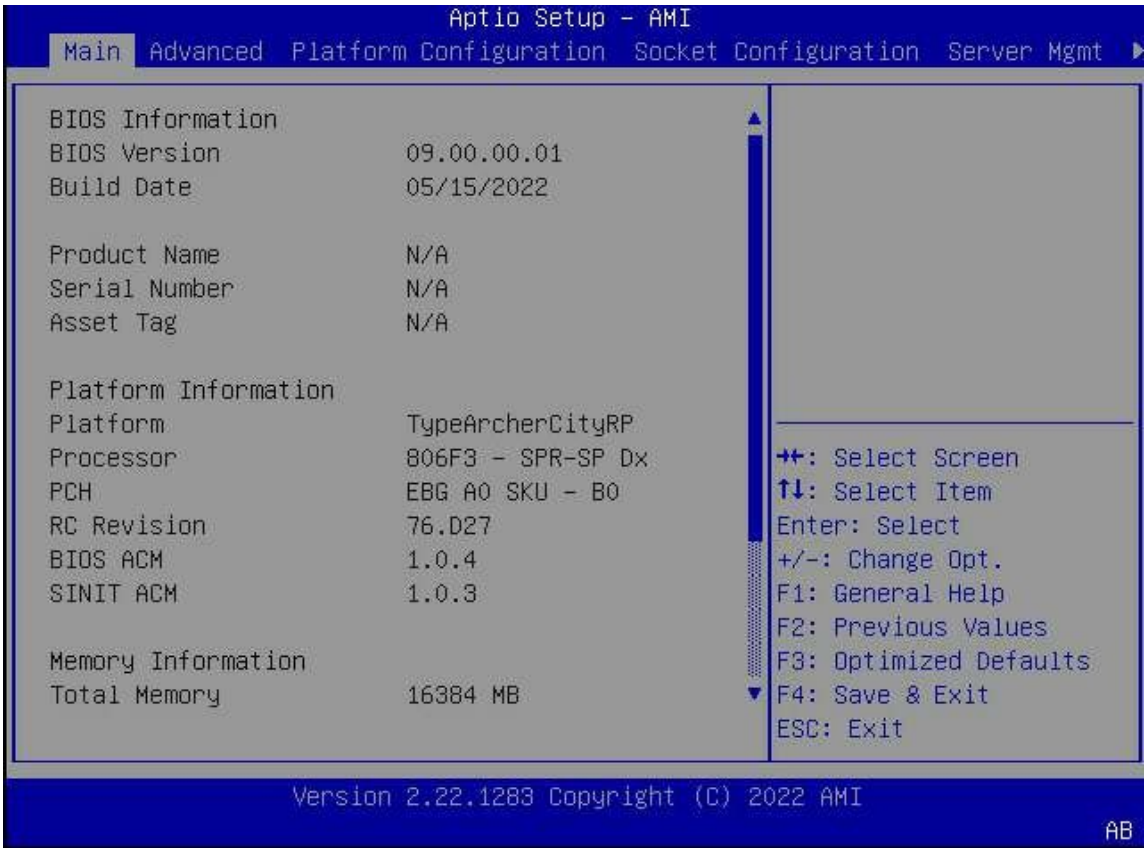

9. Use the arrow keys to select **Advanced**. On the displayed **Advanced** screen (see Figure 5-9), the **Intel(R) Virtual RAID on CPU** option is displayed, indicating that the VMD function is enabled for the BIOS.

Figure 5-9 Advanced Screen

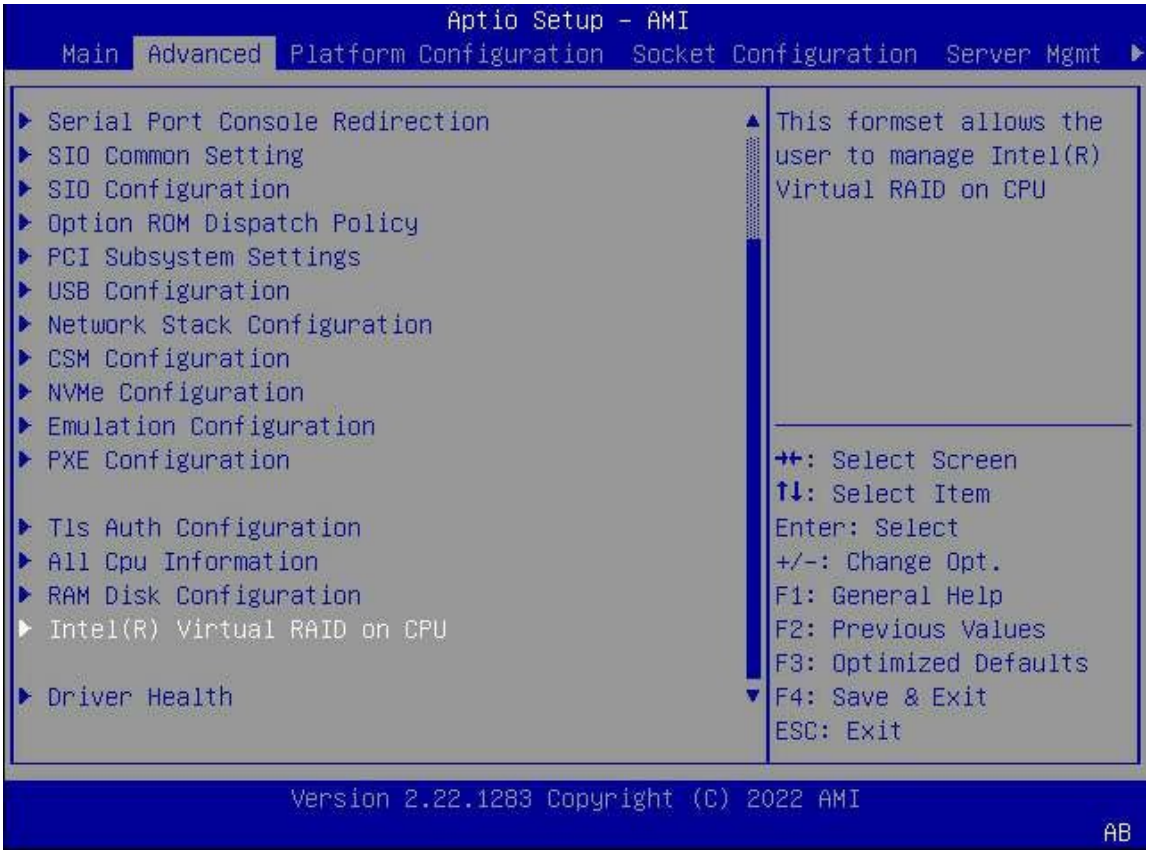

5.2.3 Creating a RAID Volume

Abstract

You can create RAID volumes at different levels as required.

The operations for creating RAID volumes at different levels are similar. This procedure uses creating a RAID 0 volume as an example.

Prerequisite

- Sufficient NVMe SSDs are installed on the server.
- The VMD function is enabled. For details, refer to “5.2.2 Enabling the VMD Function”.

Context

For a description of the number of disks required to create a RAID volume, refer to Table 5-9.

Table 5-9 Number of Disks Required for Creating a RAID Volume

RAID Level	Description
RAID 0	RAID 0 requires at least one disk.
RAID 1	RAID 1 requires at least two disks.

RAID Level	Description
	Disks with different capacities can be used in a RAID 1 volume, but the logical capacity of each member disk depends on the space of the disk with the smallest capacity.
RAID 5	RAID 5 requires at least three disks.
RAID 10	RAID 10 requires at least four disks. A RAID 10 volume consists of multiple RAID 1 volumes (at least two).

Steps

1. On the **Aptio Setup** screen, use the arrow keys to select **Advanced > Intel(R) Virtual RAID on CPU**, and press **Enter**. The **Intel(R) Virtual RAID on CPU** screen is displayed, see Figure 5-10.

Figure 5-10 Intel(R) Virtual RAID on CPU Screen

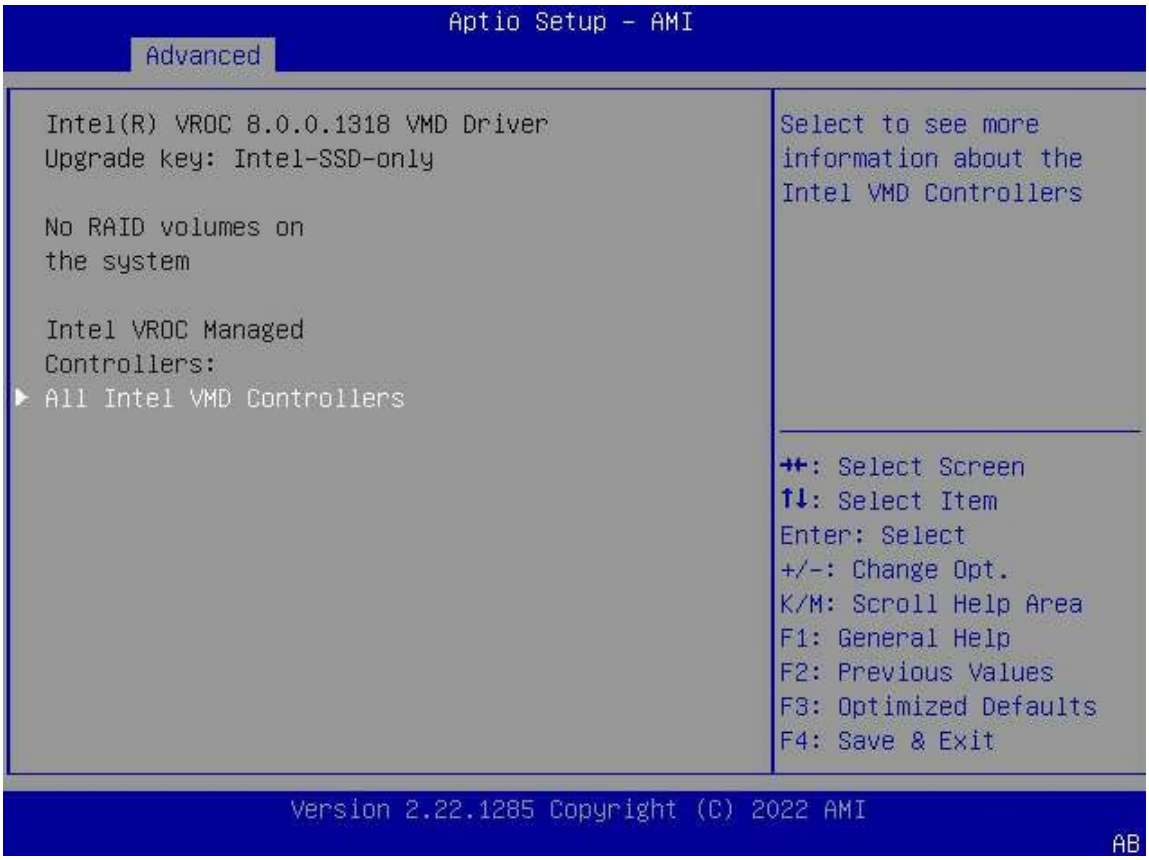

2. Use the arrow keys to select **All Intel VMD Controllers**, and press **Enter**. On the displayed **All Intel VMD Controllers** screen (see Figure 5-11), all the NVMe SSDs mounted to the CPU are displayed.

Figure 5-11 All Intel VMD Controllers Screen

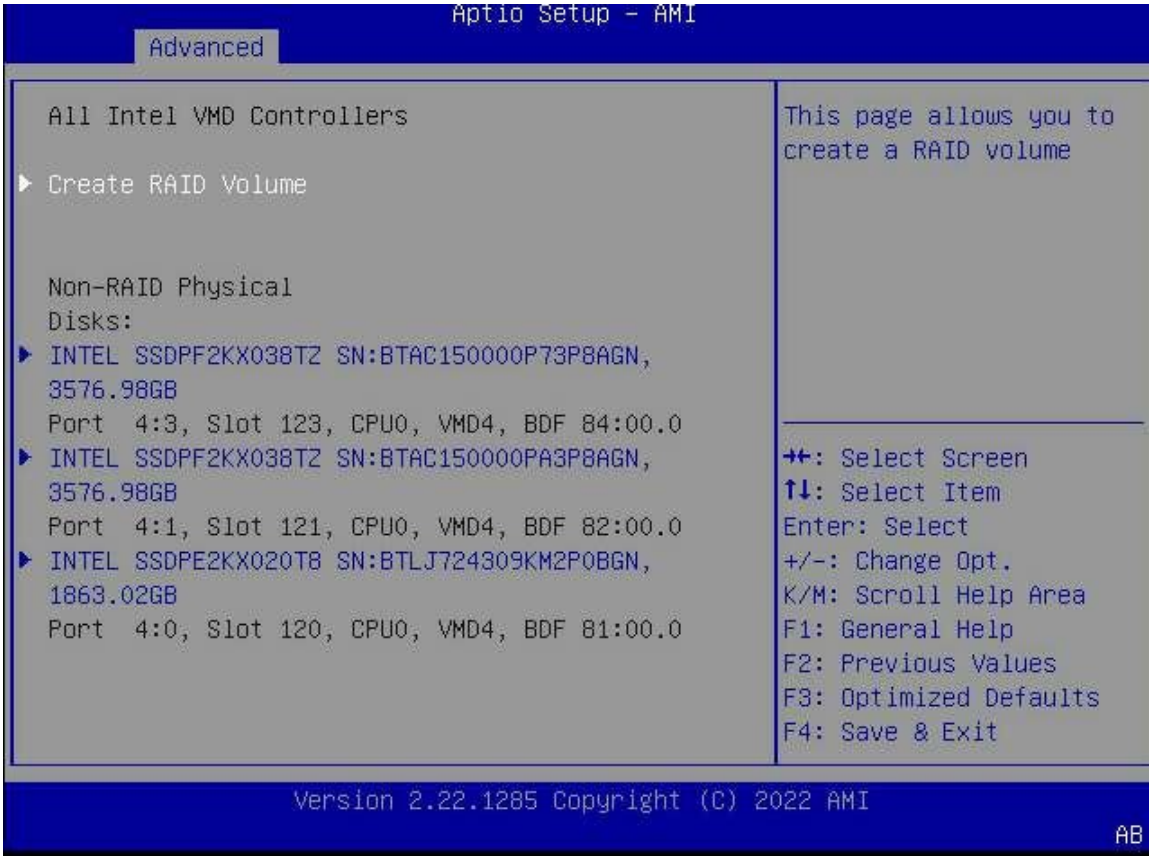

3. Use the arrow keys to select **Create RAID Volume**. The screen for creating a RAID volume is displayed, see Figure 5-12.

Figure 5-12 Creating a RAID Volume

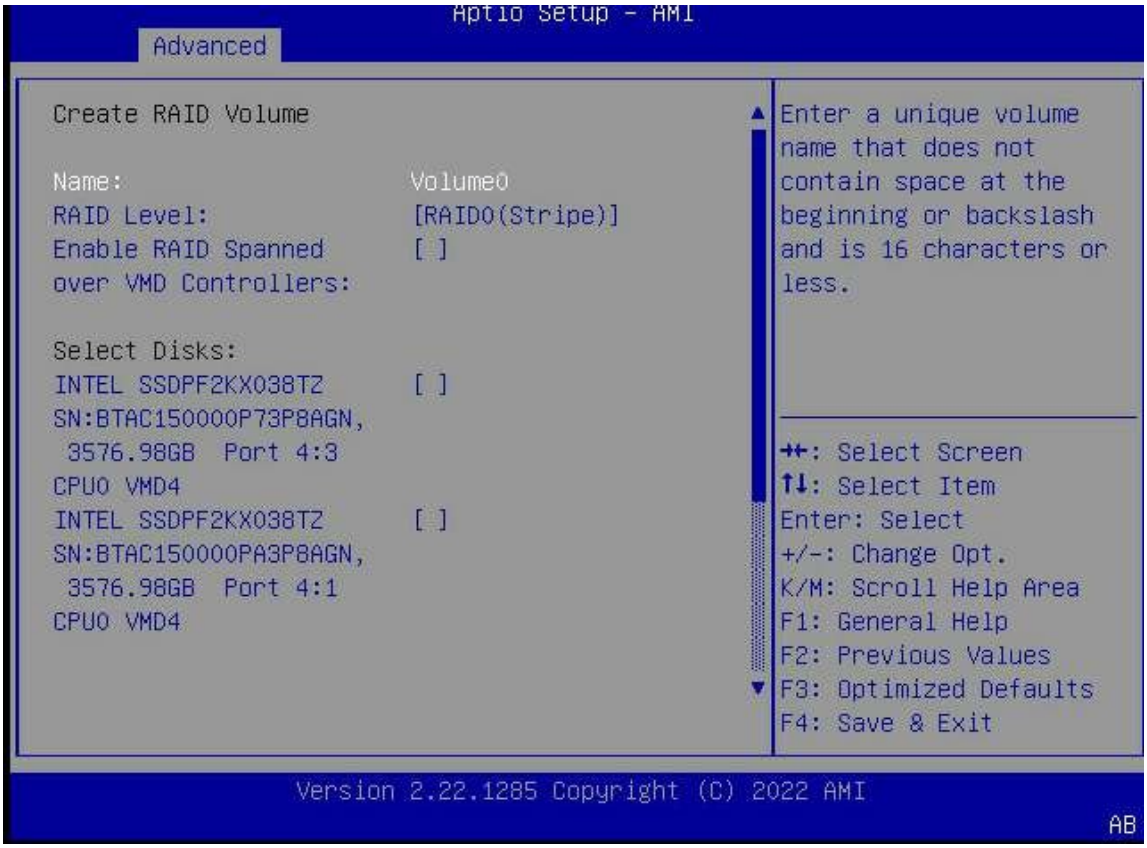

4. Use the arrow keys to select **Name**, and then press **Enter**. On the displayed screen, set the name of the RAID volume, see Figure 5-13.

Figure 5-13 Setting the RAID Volume Name

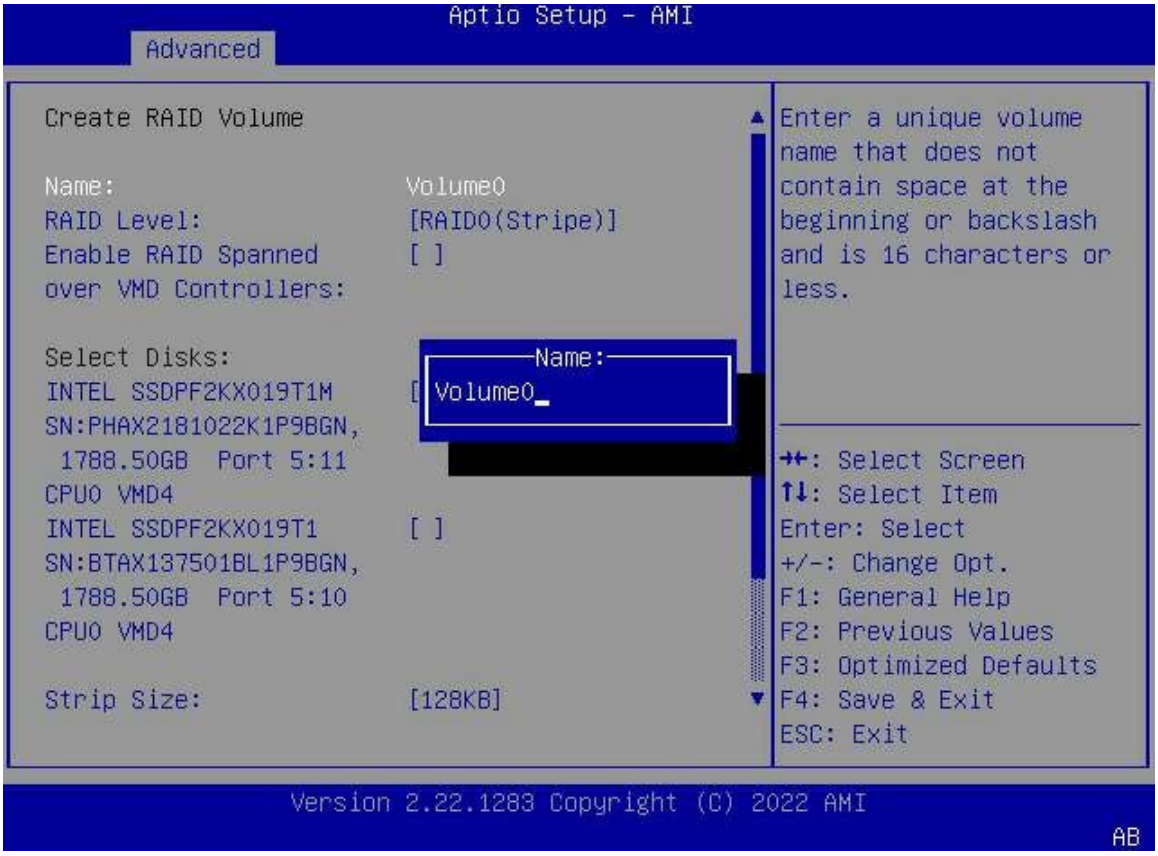

5. Use the arrow keys to select **RAID Level**, and then press **Enter**. From the displayed short-cut menu, select the desired RAID level, see Figure 5-14.

Figure 5-14 Selecting a RAID Level

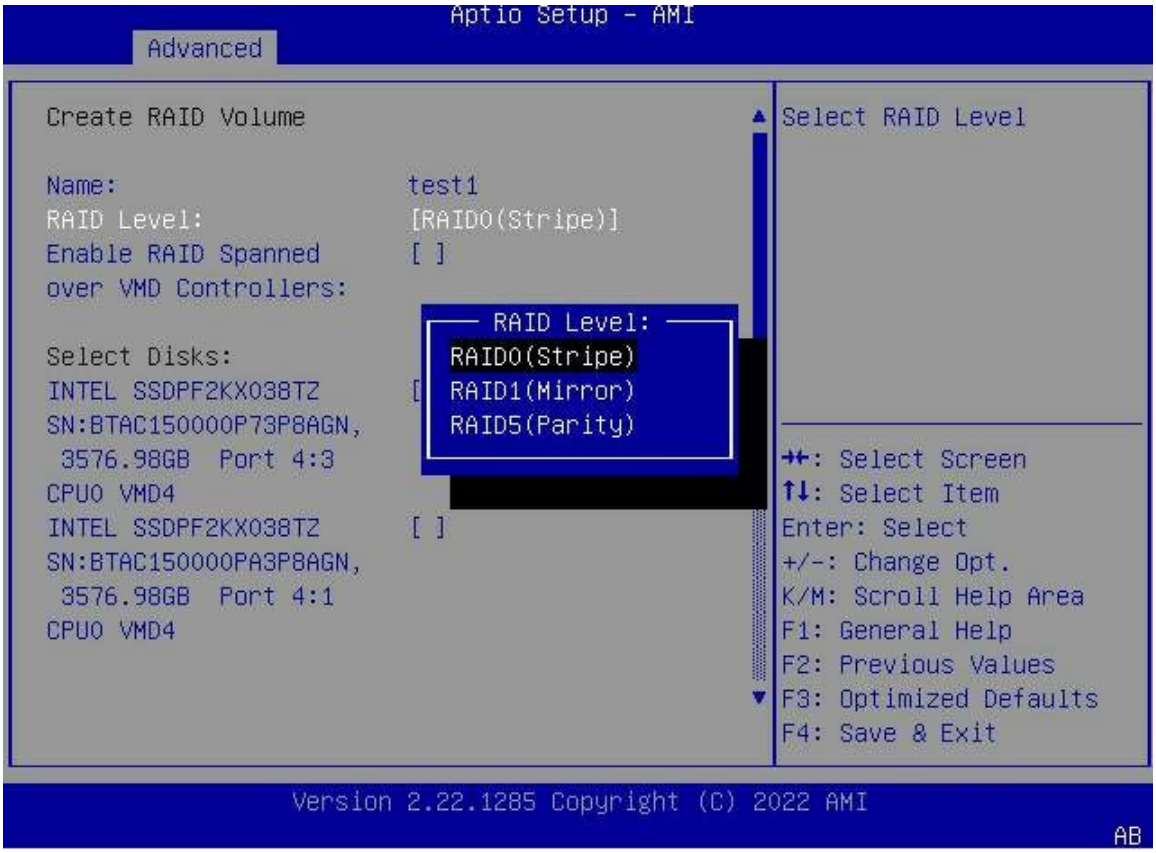

6. (Optional) When the NVMe SSDs used to create the RAID volume are distributed on different CPUs or VMDs, use the arrow keys to select **Enable RAID Spanned over VMD Controllers**, press **Enter**, and select **X** from the displayed shortcut menu to enable the **Enable RAID Spanned over VMD Controllers** option, see Figure 5-15.

Figure 5-15 Enabling the Enable RAID Spanned over VMD Controllers Option

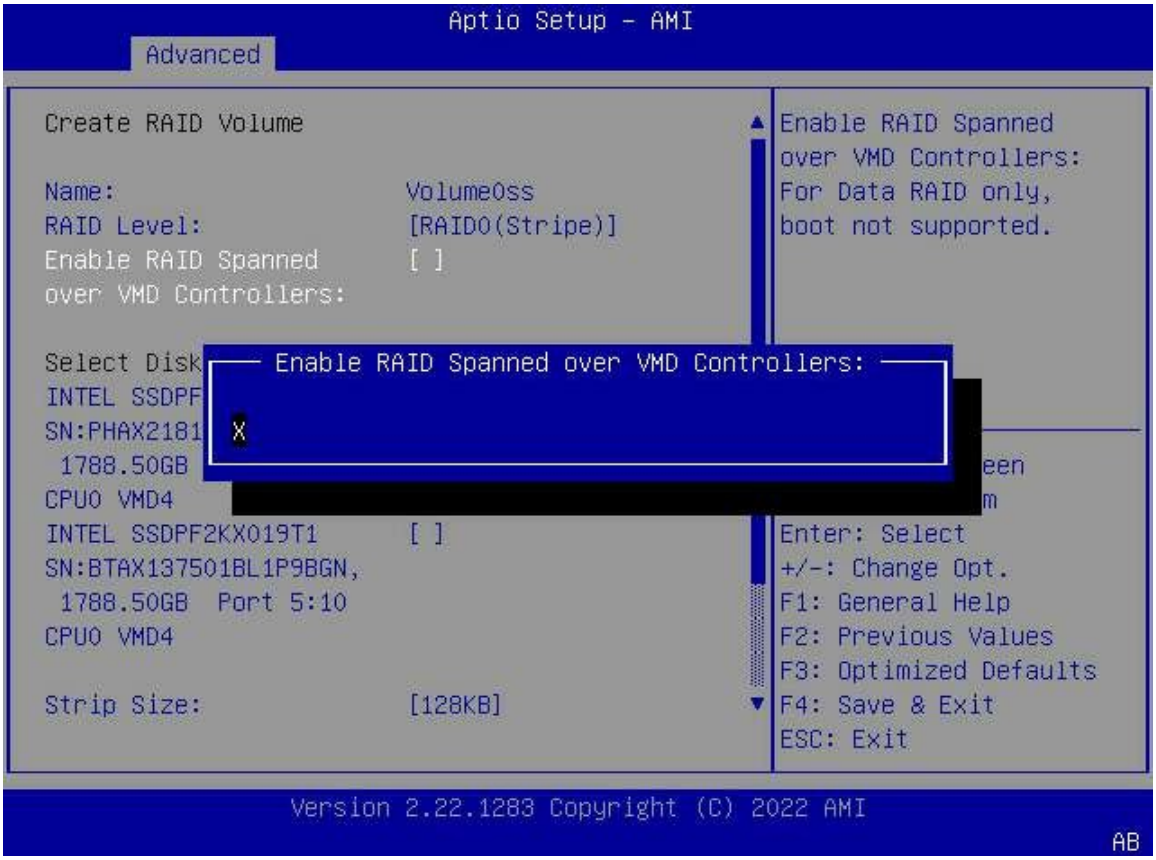

For the distribution of NVMe SSDs on different VMDs and different CPUs on a 12-disk back-plane with default wiring, see Figure 5-16.

Figure 5-16 Distribution of NVMe SSDs on Different VMDs and Different CPUs

CPU1 VMD3	CPU1 VMD3	CPU1 VMD4	CPU0 VMD4
CPU1 VMD3	CPU1 VMD4	CPU1 VMD4	CPU0 VMD4
CPU1 VMD3	CPU1 VMD4	CPU0 VMD4	CPU0 VMD4

7. Use the arrow keys to select the NVME SSD for creating the RAID volume from the **Select Disks** list, press **Enter**, and select **X** from the displayed shortcut menu, see Figure 5-17.

Figure 5-17 Selecting an NVMe SSD

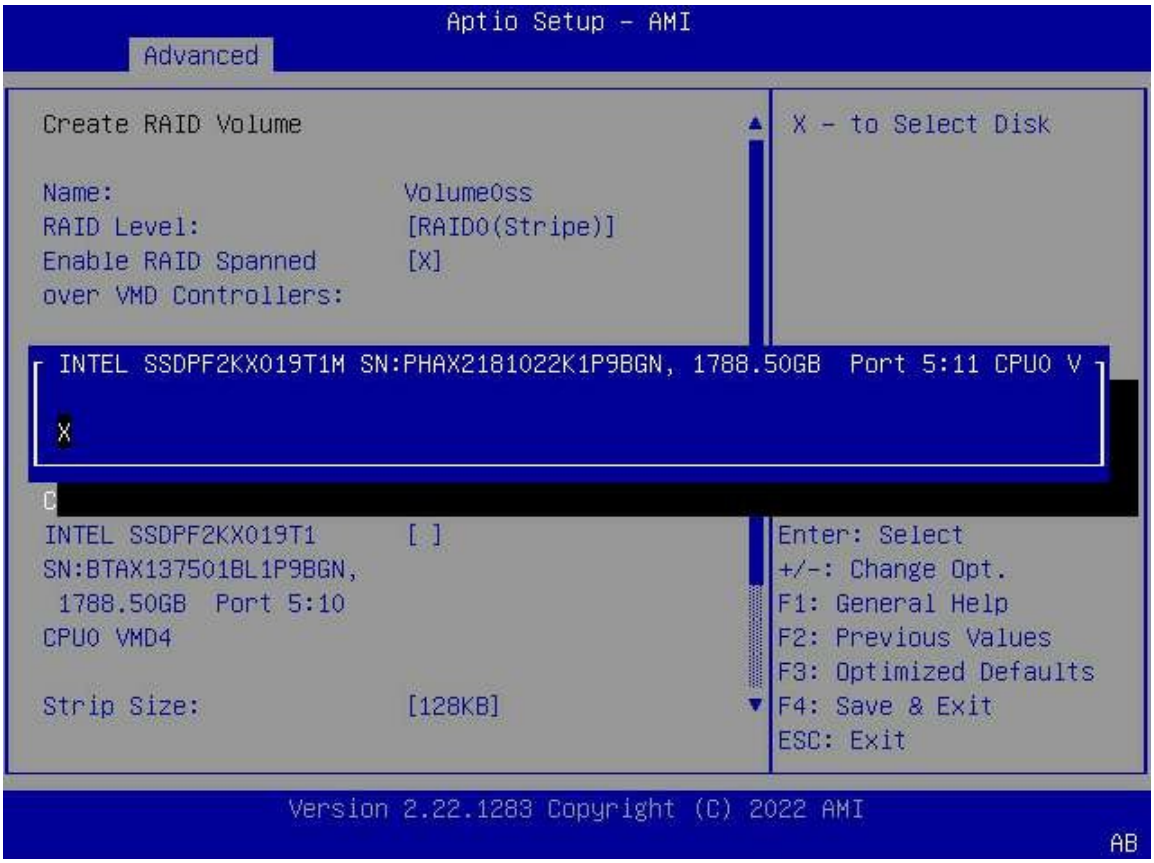

8. Use the arrow keys to select **Strip Size**, and then press **Enter**. From the displayed shortcut menu, select the desired RAID stripe size, see Figure 5-18.

Figure 5-18 Selecting Stripe Capacity

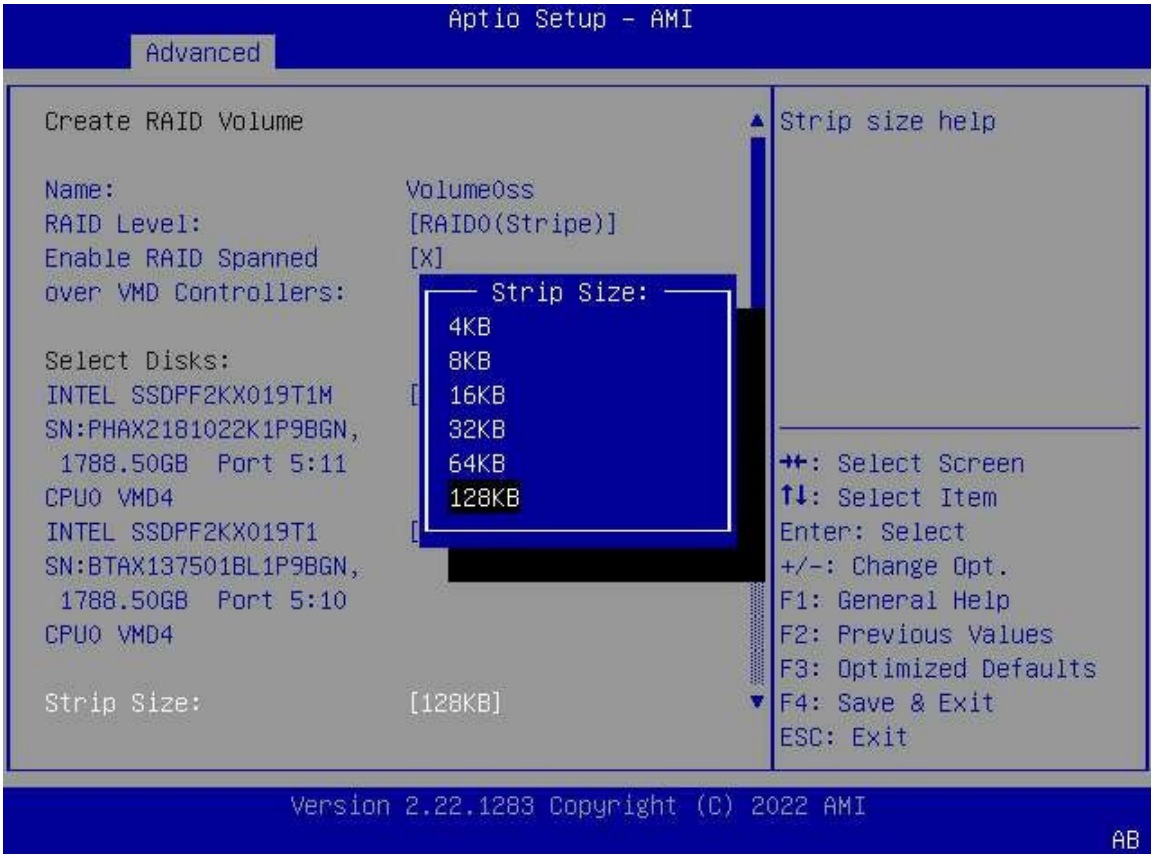

Note

If you create RAID 1, there is no **Strip Size** option.

9. Use the arrow keys to select **Capacity(GB)**, and then press **Enter**. From the displayed shortcut menu, select the desired capacity for the RAID logical disk, see Figure 5-19.

Figure 5-19 Selecting Capacity for RAID

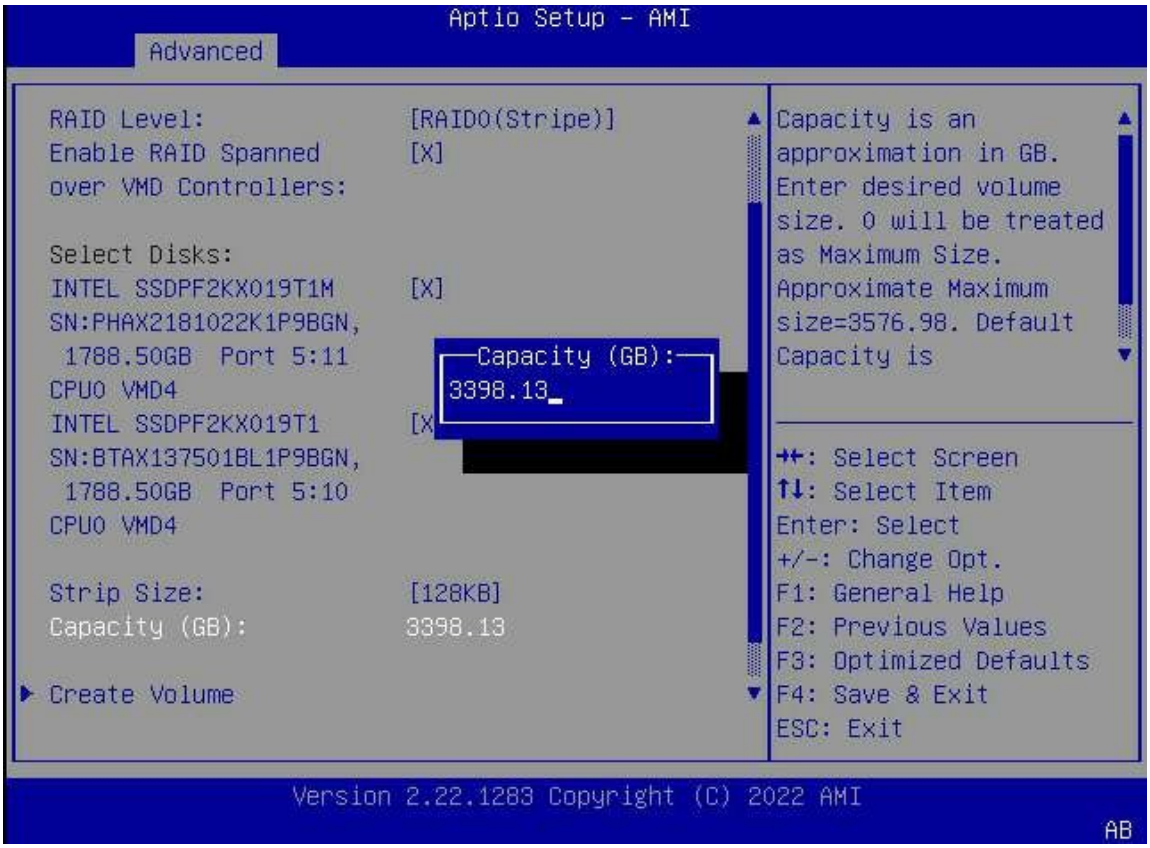

10. Use the arrow keys to select **Create Volume**, and then press **Enter**. The confirmation screen for RAID volume creation is displayed.

Note

If the capacity of the selected disks for creating a RAID volume is different, a warning message is displayed. If **Yes** is selected, the system automatically selects the minimum disk capacity to create a RAID volume.

11. Use the arrow keys to select **Yes**, and then press **Enter** to create the RAID volume. After successful creation, the **All Intel VMD Controllers** screen is displayed, see .

Figure 5-20 RAID Volume Created Successful

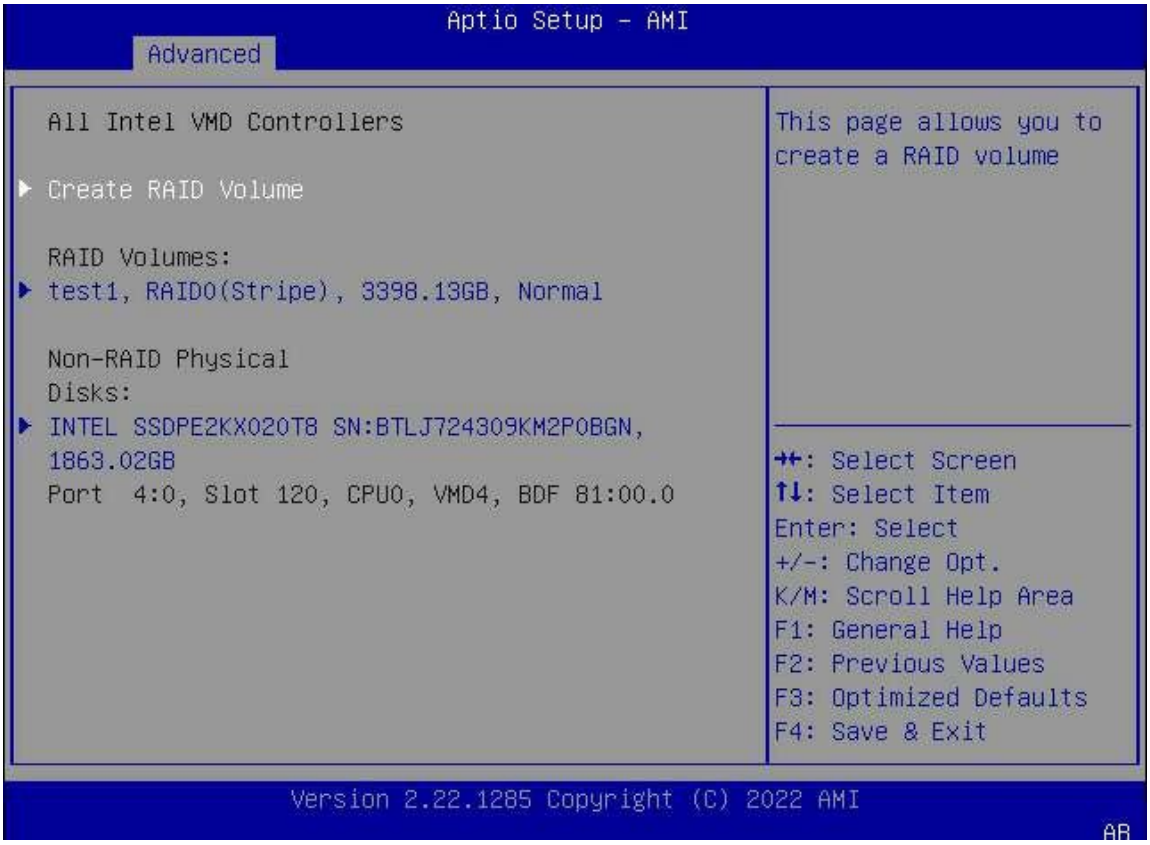

5.2.4 Querying RAID Volume Information

Abstract

This procedure describes how to query the information about a RAID volume created on a RAID controller card, including the RAID volume name, level, capacity, and status.

You can query RAID volume information in the following two ways:

- BIOS screens
- Linux OS screens

Prerequisite

- The VMD function is successfully enabled. For details, refer to “5.2.2 Enabling the VMD Function”.
- A RAID volume is created successfully. For details, refer to “5.2.3 Creating a RAID Volume”.

Steps

- Querying RAID Volume Information on BIOS Screens
  1. On the **Aptio Setup** screen, use the arrow keys to select **Advanced > Intel(R) Virtual RAID on CPU > All Intel VMD Controllers**, and press **Enter**. The **All Intel VMD Con-**

trollers screen is displayed. The existing RAID volumes are displayed in the **RAID Volumes** area, see Figure 5-21.

Figure 5-21 All Intel VMD Controllers Screen

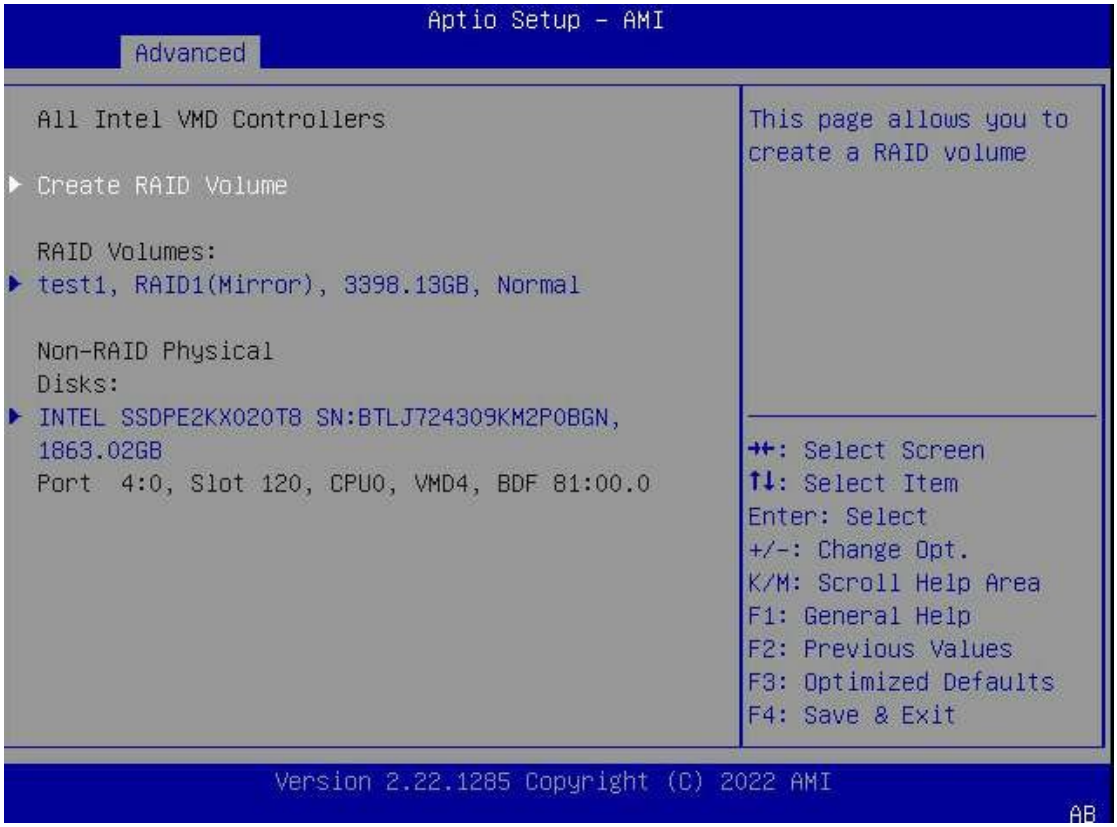

2. Use the arrow keys to select the RAID volume to be queried, and press **Enter**. The RAID volume information is displayed, see Figure 5-22. For a description of the parameters, refer to Table 5-10.

Figure 5-22 RAID Volume Information

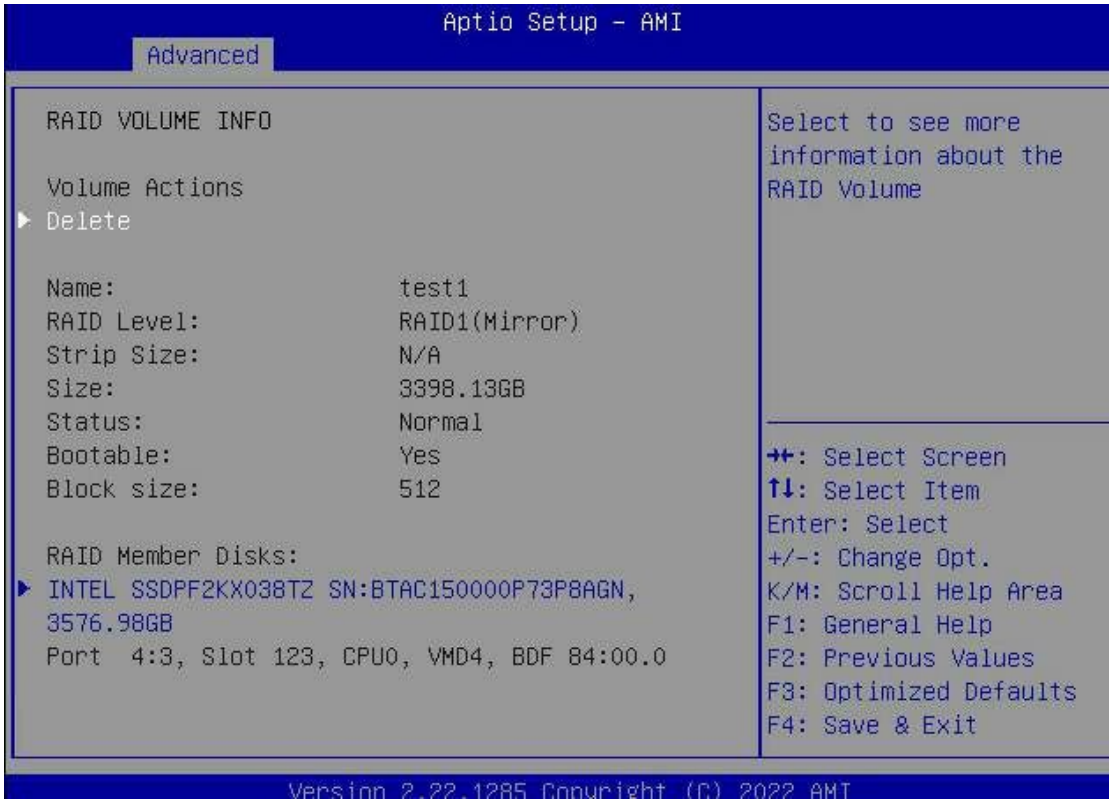

Table 5-10 Parameter Descriptions for RAID Volume Information

Parameter	Description
Name	Name of the RAID volume
RAID Level	Level of the RAID volume
Size	Strip size, namely, the size of the stripe data block written into each disk
Size	Stripe size of a disk group, namely, the data block set written on all disks
Status	Disk group status
Bootable	Indicates that the RAID volume is bootable
Block Size	Size of each data block
RAID Member Disks	Member disks of the RAID volume

- Querying RAID Volume Information on Linux OS Screens

  1. Enter the **`cat /proc/mdstat`** command in the Linux operating system to display the status of all running RAID volumes, see Figure 5-23.

Figure 5-23 RAID Volume Status

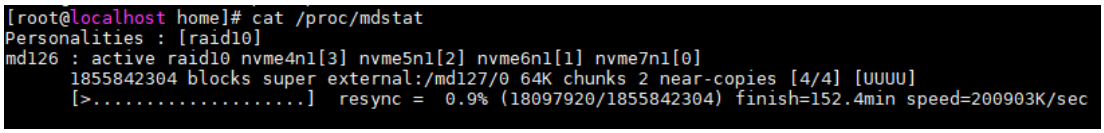

For a description of the RAID volume status parameters, refer to Table 5-11.

Table 5-11 Parameter Descriptions for RAID Volume Status

Parameter	Description
Personalities : [raid10]	Created RAID volume.
md126	Logical drive letter of the RAID volume.
active / inactive	Whether the RAID volume is activated.
nvme4nl[3] nvme5nl[2] nvme6nl[1] nvme7nl[0]	NvmeXnl[Y] indicates the location of the logical disk in the RAID volume.
1855842304 Blocks super external	The available space of the RAID volume is 387072 blocks.
64k chunks	The size of chunks is 64 K.
[4/4] [UUUU]	The RAID volume requires four logical disks, and the four logical disks are operating properly. If a logical disk is not operating properly, "_" is displayed in the U position.
resync=0.9%	Rebuild completion percentage of the current RAID volume.
finish=152.4 Min	Expected rebuild completion time of the RAID volume.
speed=200903k/sec	RAID volume rebuild speed.

2. Enter the **`mdadm - -detai -platform`** command to display the RAID volume information created under the **Intel(R) Virtual RAID on CPU** option, see Figure 5-24.

Figure 5-24 RAID Volume Information

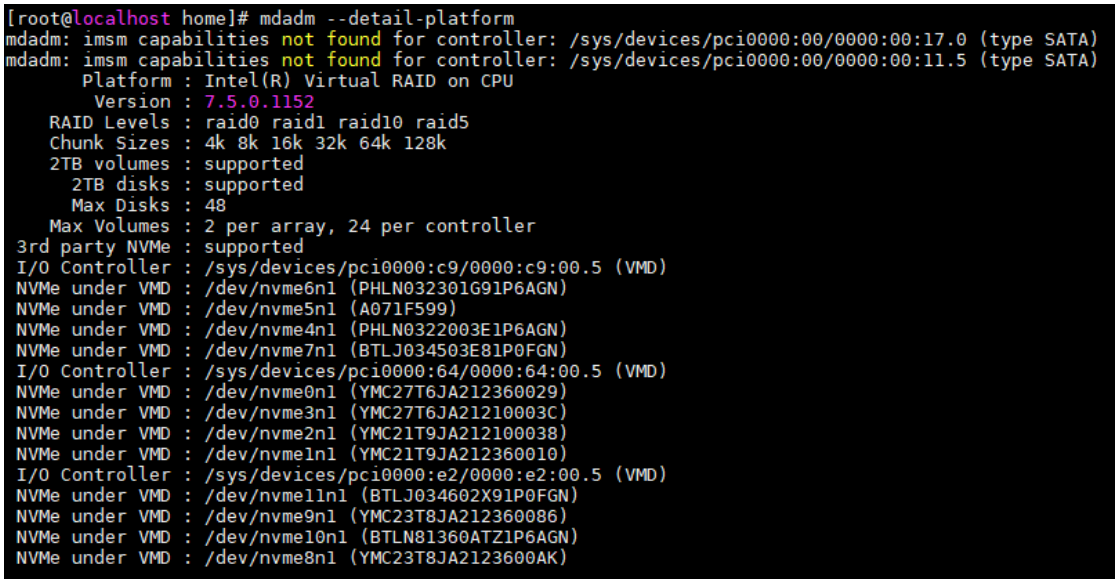

For a description of the RAID volume information parameters, refer to Table 5-12.

Table 5-12 Parameter Descriptions for RAID Volume Information

Parameter	Description
Version	Version number of the NVMe VROC Key
RAID Levels	RAID level supported by the NVMe VROC Key
Chunk Sizes	Data block size supported by the NVMe VROC Key
2TB volumes	Supported, indicating that a 2 TB RAID volume is supported
2TB disks	Supported, indicating that a 2 TB disk is supported
Max Disks	Maximum number of disks supported
Max Volumes	→ Maximum number of volumes supported by each RAID level → Maximum number of volumes supported by each RAID controller card
3rd party NVME	Supported, indicating that the NVMe VROC Key supports third-party NVMe SSDs except for those of the Intel brand

3. Enter the **`lsblk`** command to query the drive letter of the RAID logical disk, see Figure 5-25.

Figure 5-25 Drive Letter

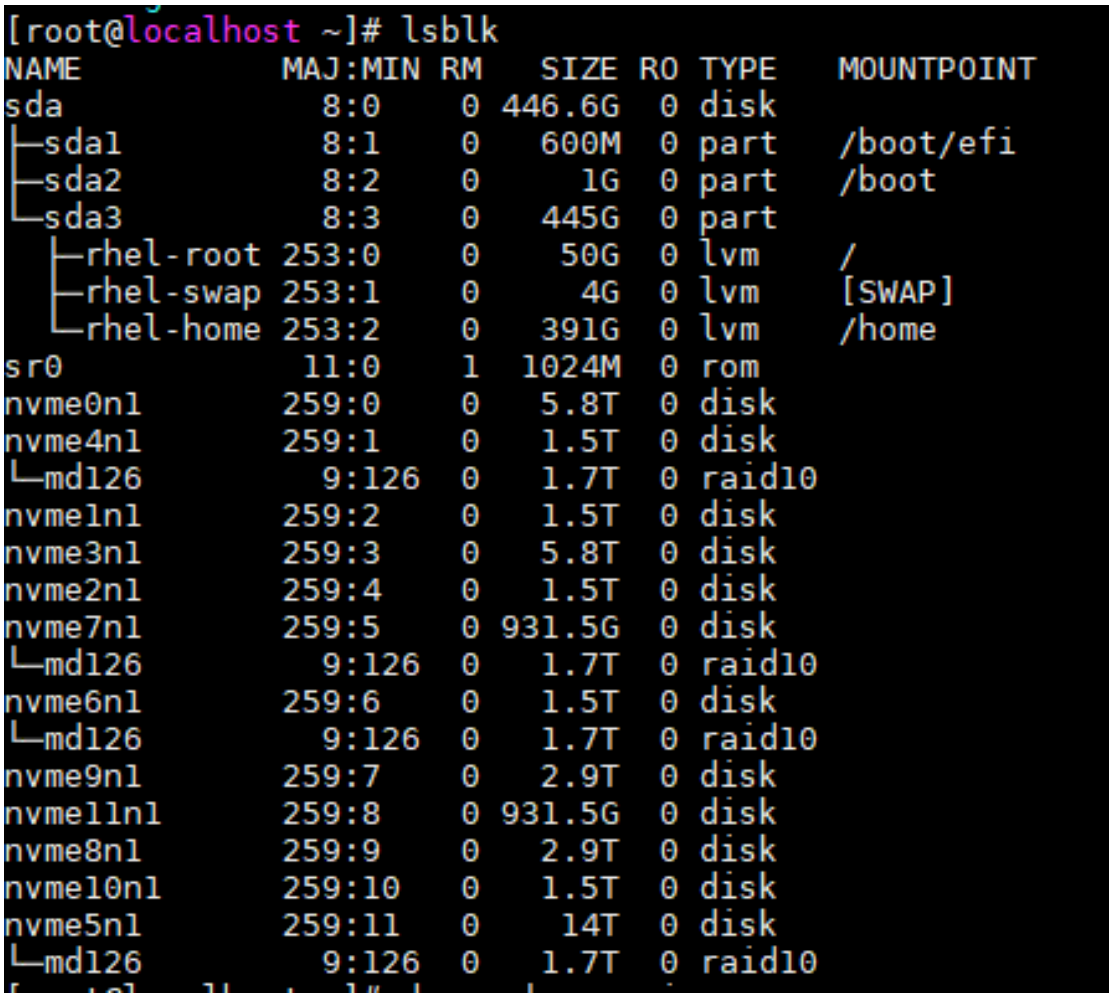

4. Enter the **`mdadm -D /dev/md126`** command to query the drive letter of the md126 volume, see Figure 5-26.

Figure 5-26 Drive Letter of md126

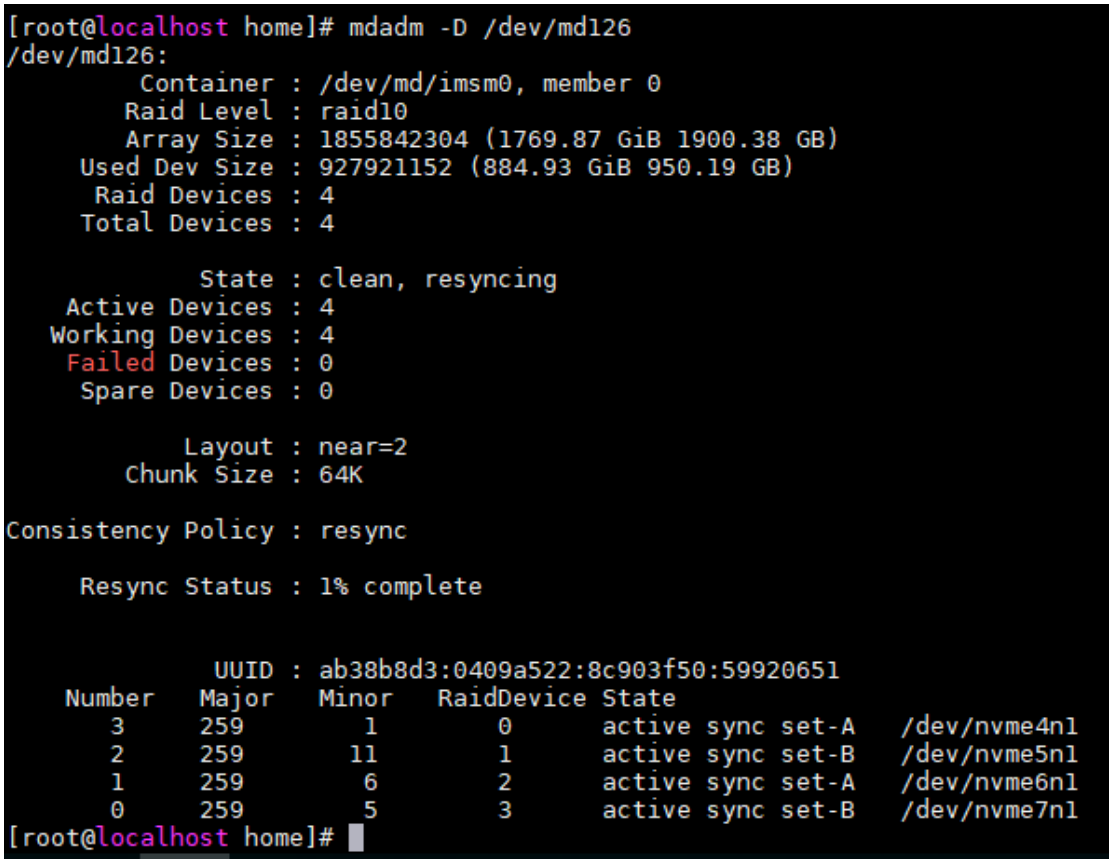

For a description of the drive letter information of md126, refer to Table 5-13.

Table 5-13 Parameter Descriptions for Drive Letter Information of md126

Parameter	Description
Container	Logical disk
Raid Level	RAID Level
Array Size	RAID array size
Used Dev Size	Available capacity of the RAID array
Raid Devices	Number of disks for creating the RAID array
Total Devices	Number of all disks
State	RAID status
Active Devices	Number of active disks
Working Devices	Number of operating disks
Failed Devices	Number of damaged disks
Spare Devices	Number of hot spare disks

Parameter	Description
Layout	RAID check rule parameter
Chunk Size	Data block size
Consistency Policy	If resync is displayed, the data synchronization is completed.
Resync Status	RAID data synchronization progress
Number	Number of a member disk in the RAID array

5.2.5 Deleting a RAID Volume

Abstract

When a server no longer needs a RAID volume, you can delete the RAID volume to release the disk space.

Notice

The data that is lost during deletion of the RAID volume cannot be restored. Therefore, you must make sure that you have backed up important data before deleting the volume.

Prerequisite

A RAID volume is created successfully. For details, refer to "5.2.3 Creating a RAID Volume".

Steps

1. On the **Aptio Setup** screen, use the arrow keys to select **Advanced > Intel(R) Virtual RAID on CPU**, and press **Enter**. The **Intel(R) Virtual RAID on CPU** screen is displayed, see Figure 5-27.

Figure 5-27 Intel(R) Virtual RAID on CPU Screen

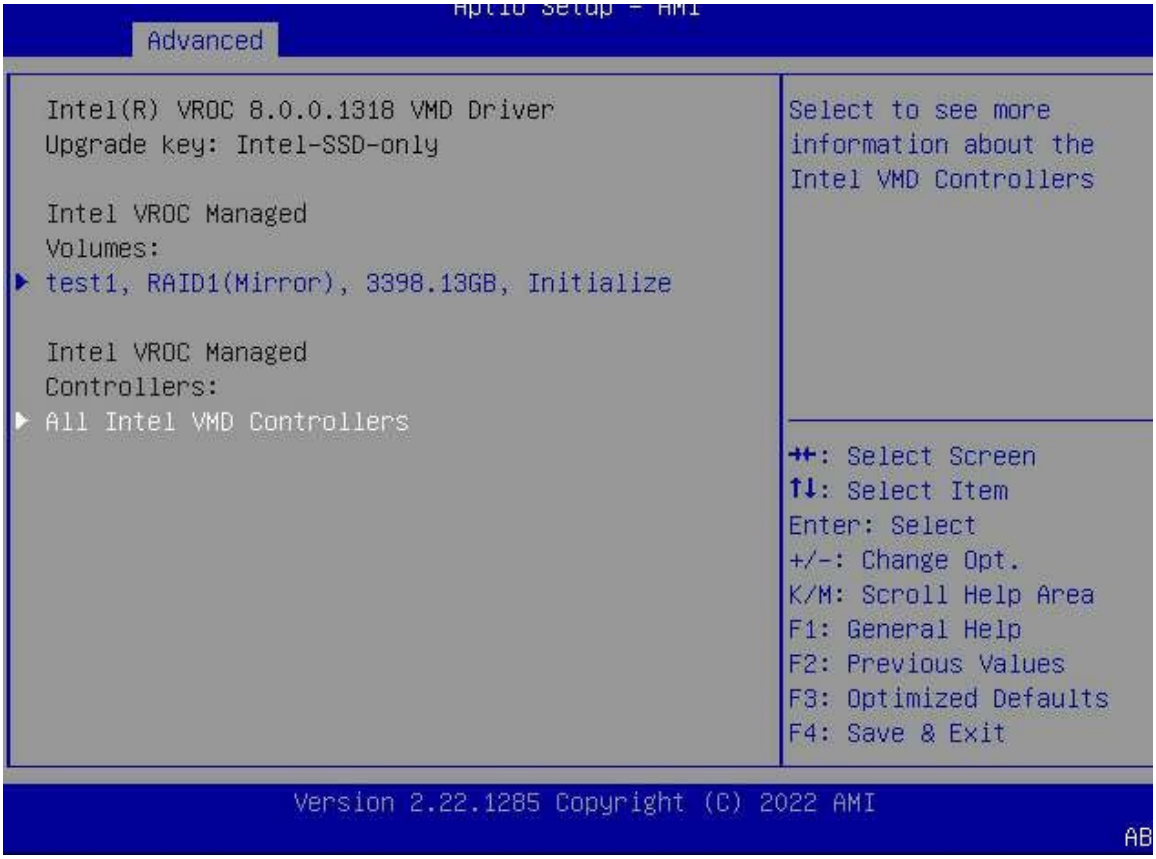

2. Use the arrow keys to select **All Intel VMD Controllers**, and press **Enter**. On the displayed **All Intel VMD Controllers** screen (see Figure 5-28), the existing RAID volumes are displayed in the **RAID Volumes** area.

Figure 5-28 All Intel VMD Controllers Screen

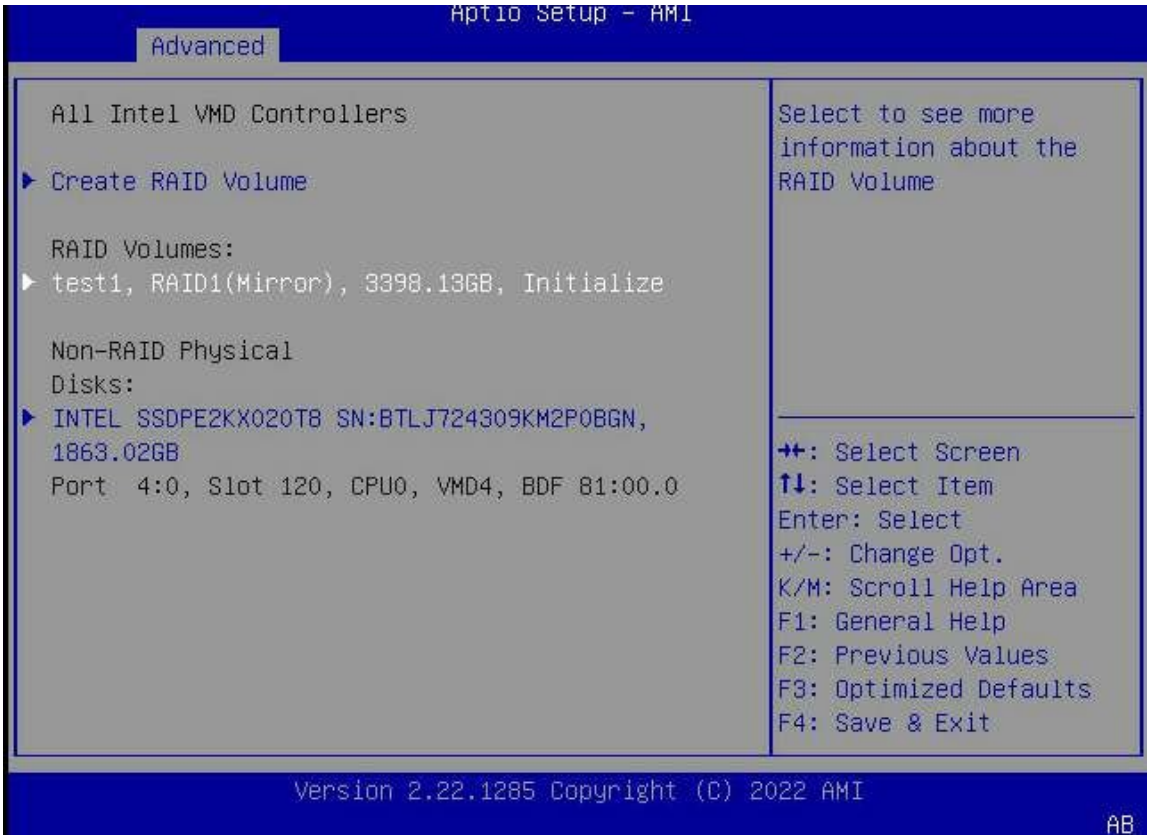

3. Use the arrow keys to select the RAID volume to be deleted, and then press **Enter**. On the displayed **RAID VOLUME INFO** screen, the information about the RAID volume to be deleted is displayed, see Figure 5-29.

Figure 5-29 RAID VOLUME INFO Screen

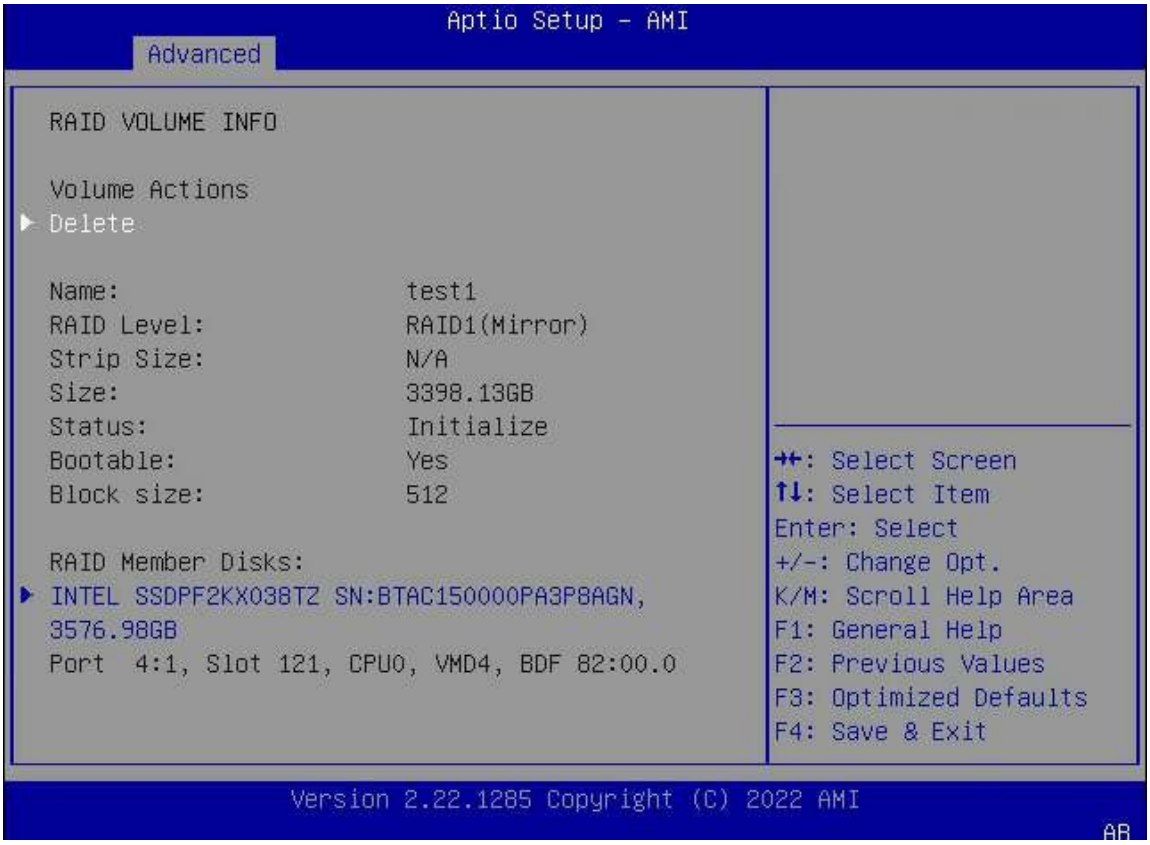

4. Use the arrow keys to select **Delete**, and then press **Enter**. The confirmation screen for RAID volume deletion is displayed, see Figure 5-30.

Figure 5-30 Confirmation Screen for RAID Volume Deletion

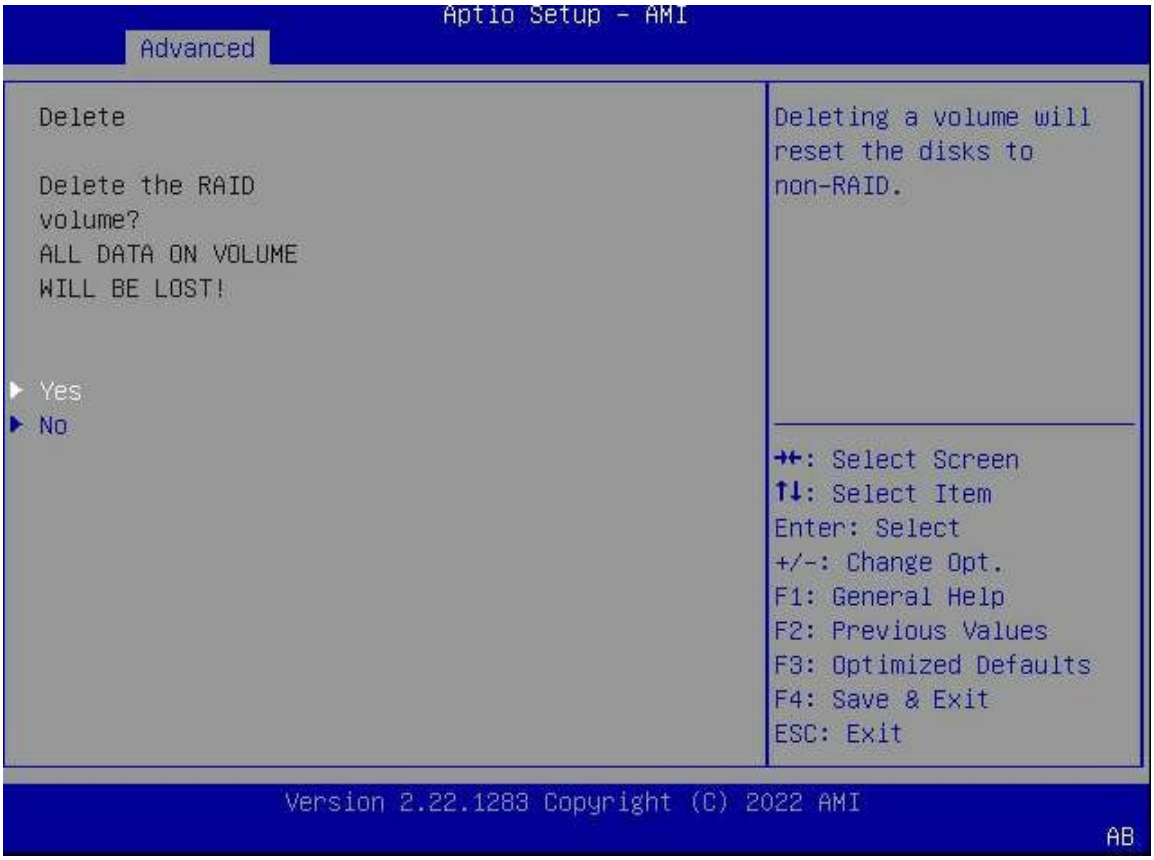

5. Use the arrow keys to select **Yes**, and then press **Enter** to delete the RAID volume. After successful deletion, the **All Intel VMD Controllers** screen is displayed, see Figure 5-31. The RAID volume disappears.

Figure 5-31 All Intel VMD Controllers Screen

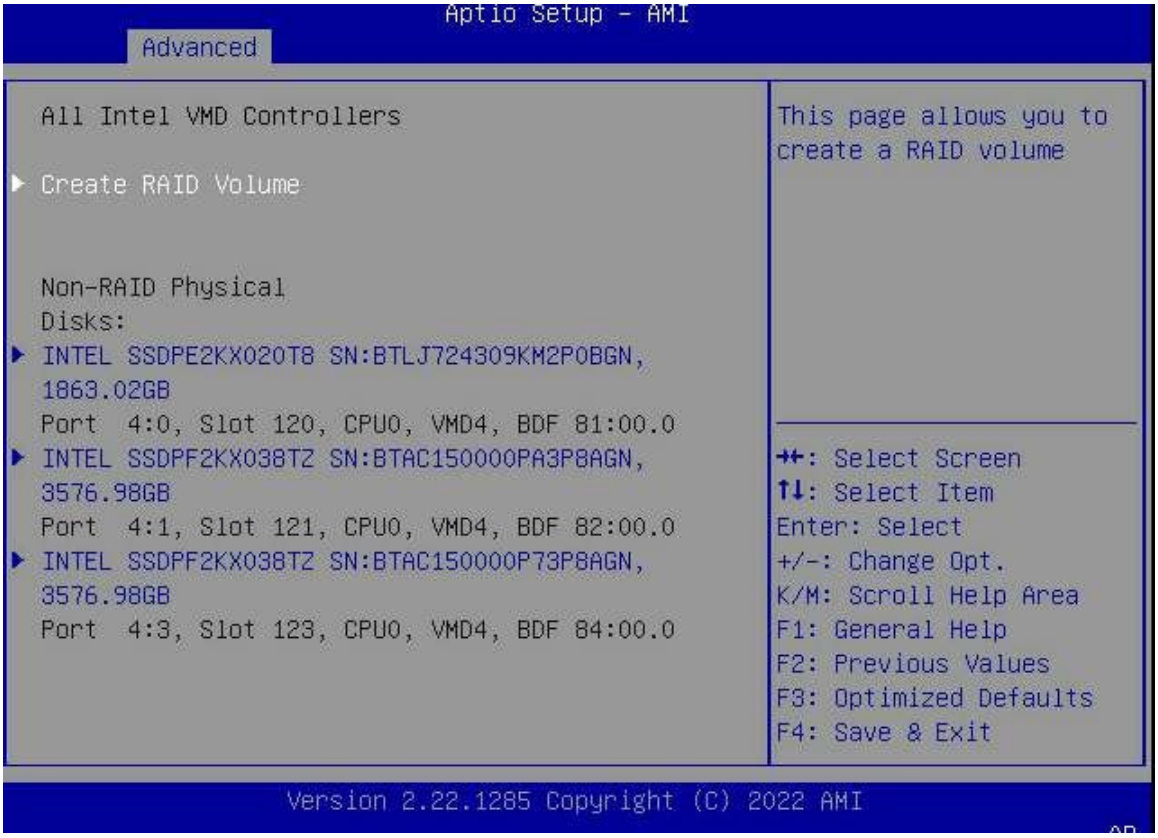

5.3 Common Operations

5.3.1 Setting the Boot Mode to Legacy

Abstract

You can set the boot mode of BIOS to **Legacy** in accordance with the actual operation requirements.

Prerequisite

The operation terminal is already connected to the real-time desktop of the server through the remote console of the management software.

Steps

Notice

Restarting the server may cause system service interruption.

1. Restart the server. The screen as shown in Figure 5-32 is displayed.

Figure 5-32 BIOS Restart Screen

2. During the POST process, press **F2**/**DEL** as prompted. The BIOS configuration utility is started, see Figure 5-33.

Figure 5-33 BIOS Configuration Utility Screen

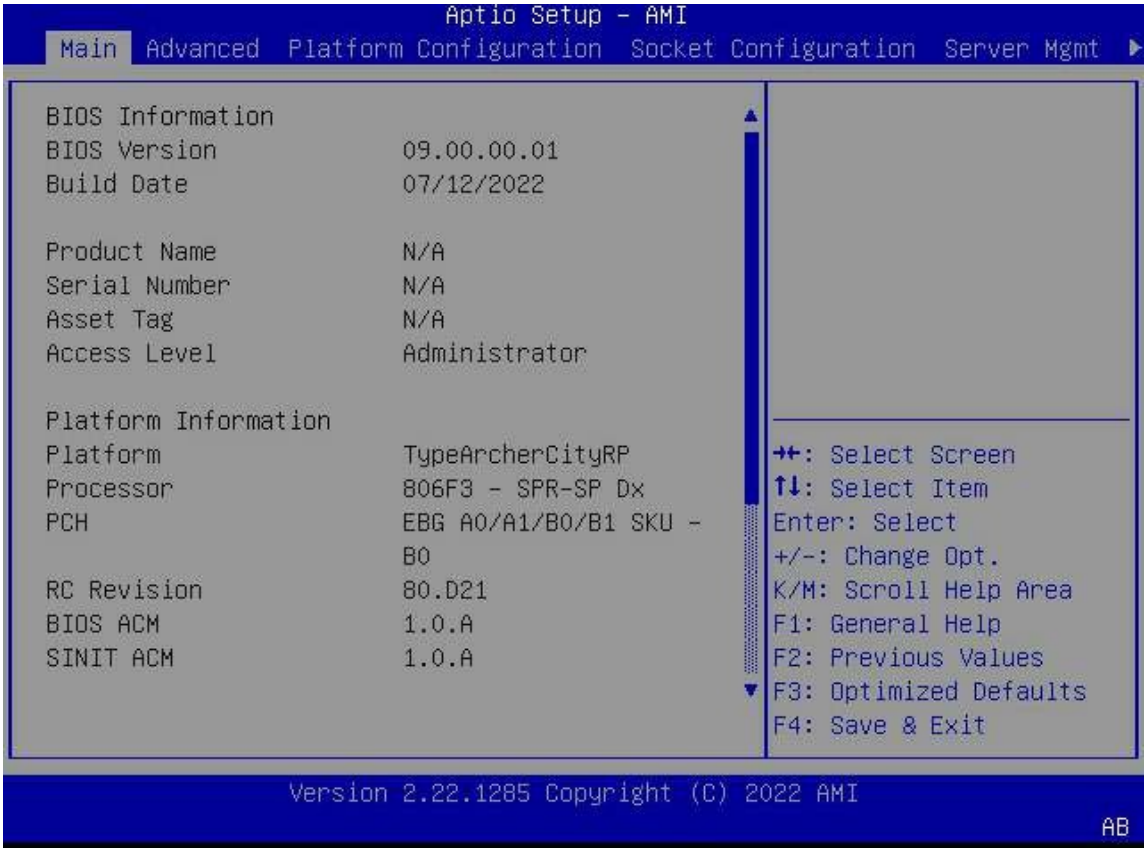

3. Use arrow keys to select **Boot > Boot option filter**, and then press **Enter**. The **Boot option filter** dialog box is displayed, see Figure 5-34.

Figure 5-34 Boot Option Filter Dialog Box

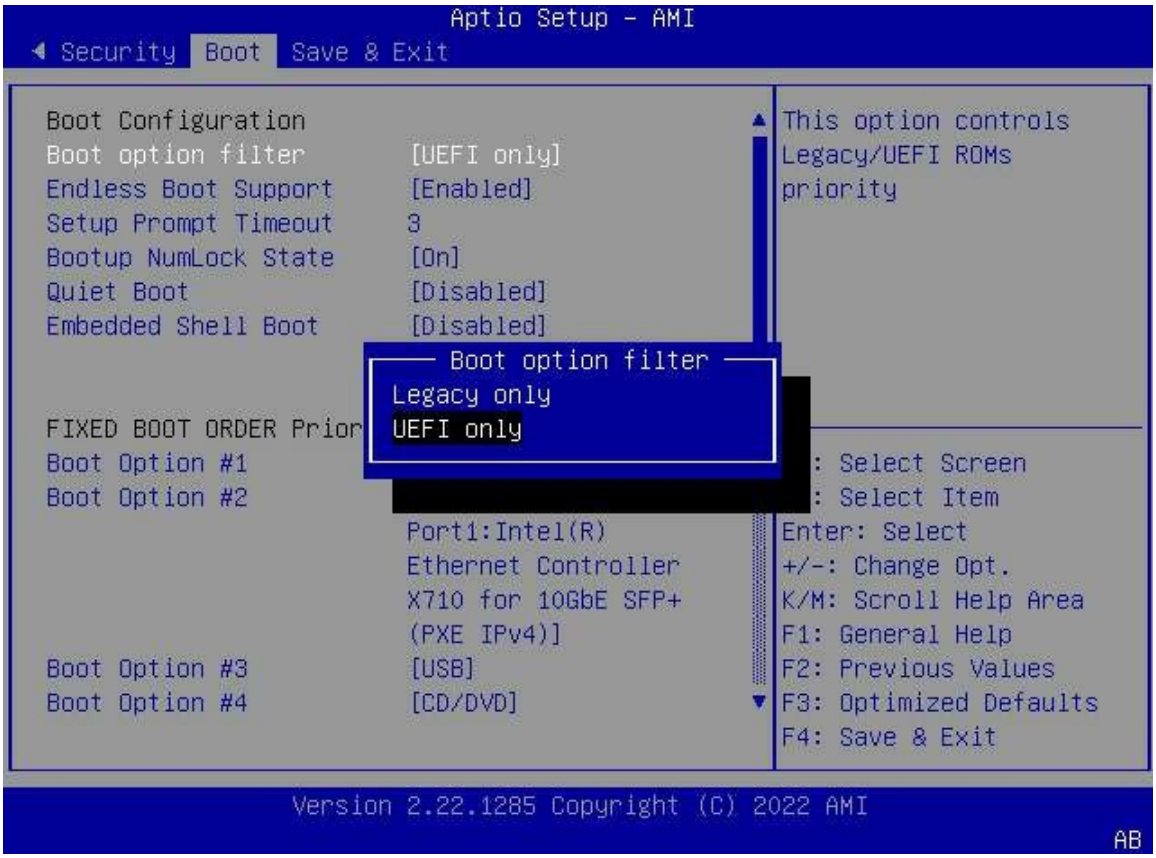

4. Use the arrow keys to select **Legacy only**, and then press **Enter** to set the boot mode of the BIOS to **Legacy**, see Figure 5-35.

Figure 5-35 Setting the Boot Mode to Legacy

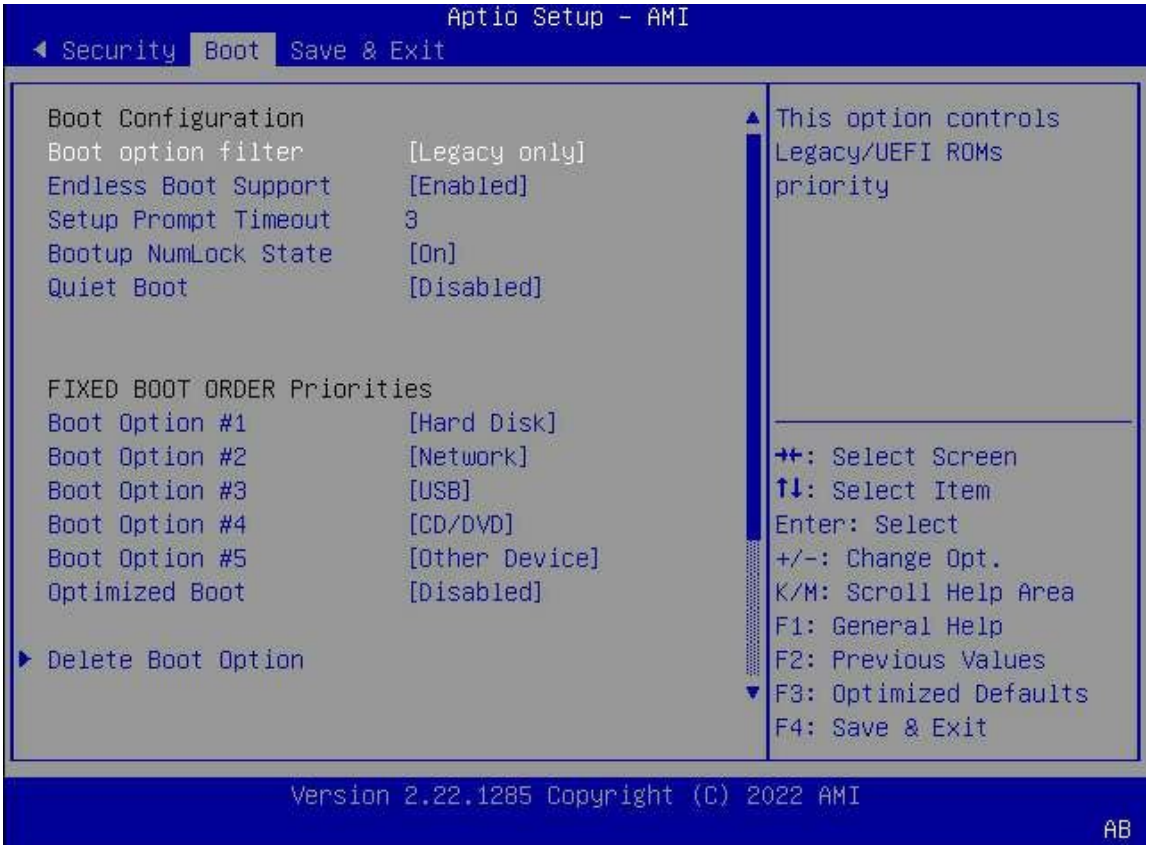

5. Press **F4** or use the arrow keys to select **Save & Exit > Save Changes and Exit**, and then press **Enter**. The **Save & Exit Setup** dialog box is displayed, see Figure 5-36.

Figure 5-36 Save & Exit Setup Dialog

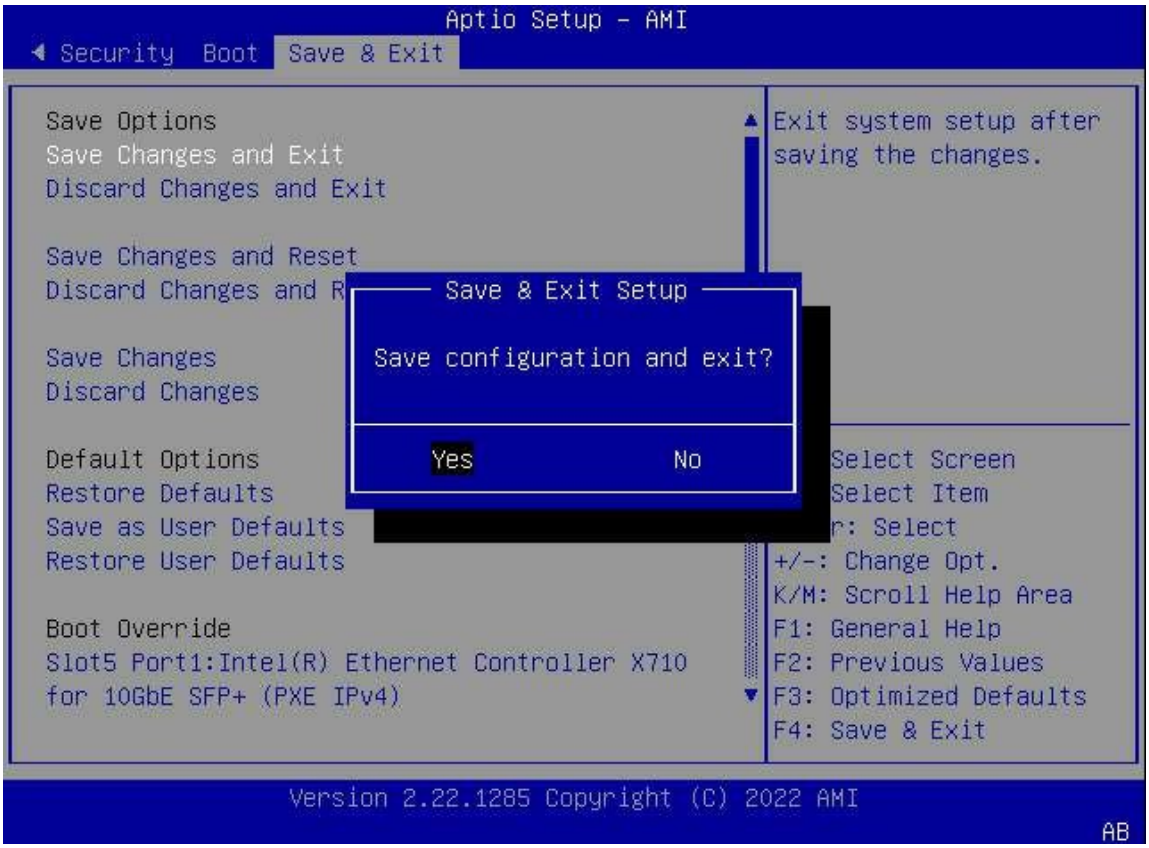

6. Use the arrow keys to select **Yes**, and press **Enter**. The boot mode is set to **Legacy**, and the server is restarted automatically.

5.3.2 Setting the Boot Mode to UEFI

Abstract

You can set the boot mode of BIOS to UEFI in accordance with the actual operation requirements.

Prerequisite

The operation terminal is already connected to the real-time desktop of the server through the remote console of the management software.

Steps

Notice

Restarting the server may cause system service interruption.

1. Restart the server. The screen as shown in Figure 5-37 is displayed.

Figure 5-37 BIOS Restart Screen

2. During the POST process, press **F2**/**DEL** as prompted. The BIOS configuration utility is started, see Figure 5-38.

Figure 5-38 BIOS Configuration Utility Screen

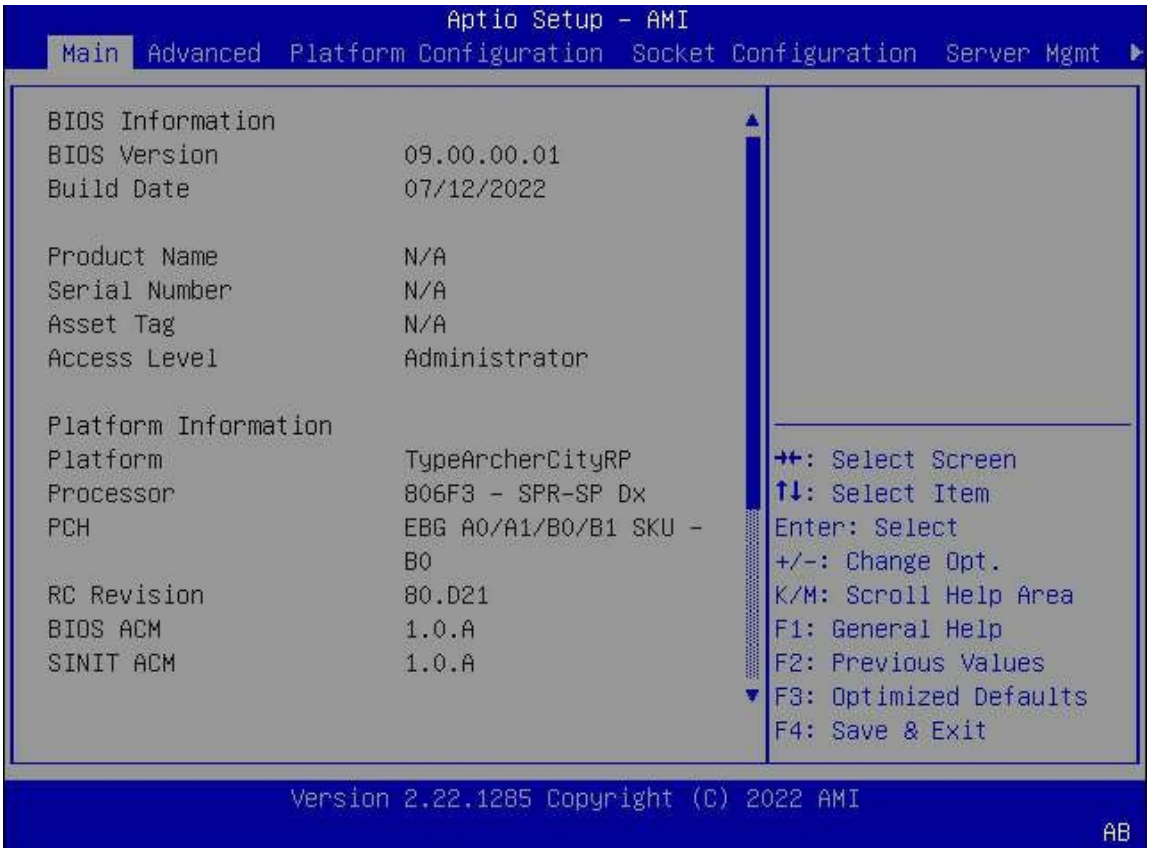

3. Use arrow keys to select **Boot > Boot option filter**, and then press **Enter**. The **Boot option filter** dialog box is displayed, see Figure 5-39.

Figure 5-39 Boot Option Filter Dialog Box

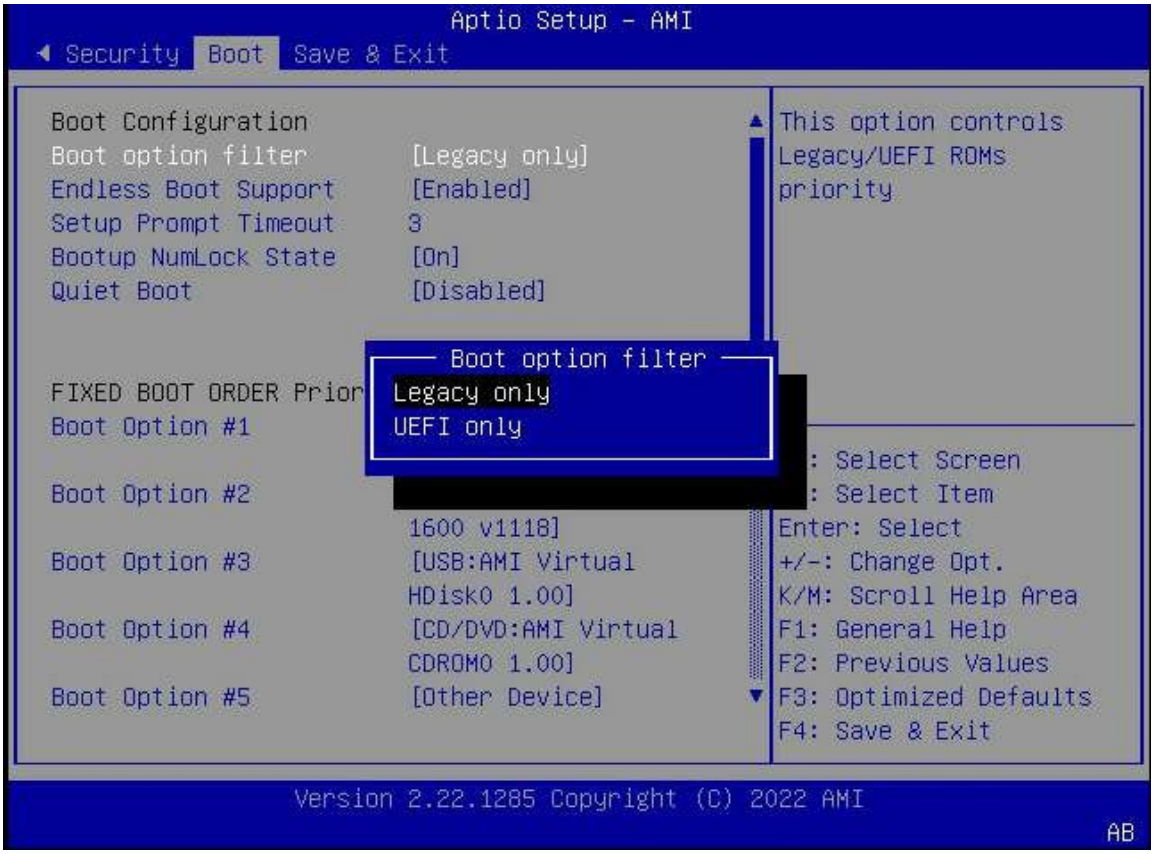

4. Use the arrow keys to select **UEFI only**, and then press **Enter** to set the boot mode of the BIOS to **UEFI**, see Figure 5-40.

Figure 5-40 Setting the Boot Mode to UEFI

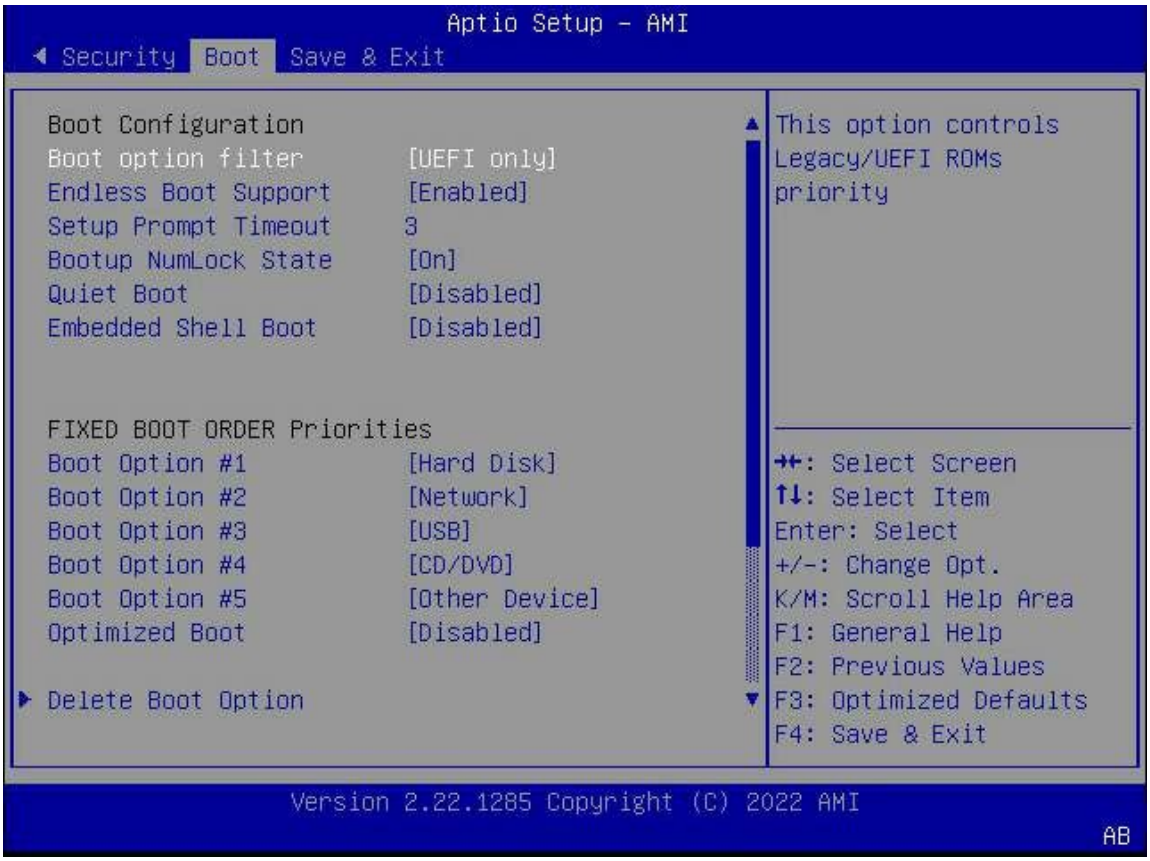

5. Press **F4** or use the arrow keys to select **Save & Exit > Save Changes and Exit**, and then press **Enter**. The **Save & Exit Setup** dialog box is displayed, see Figure 5-41.

Figure 5-41 Save & Exit Setup Dialog

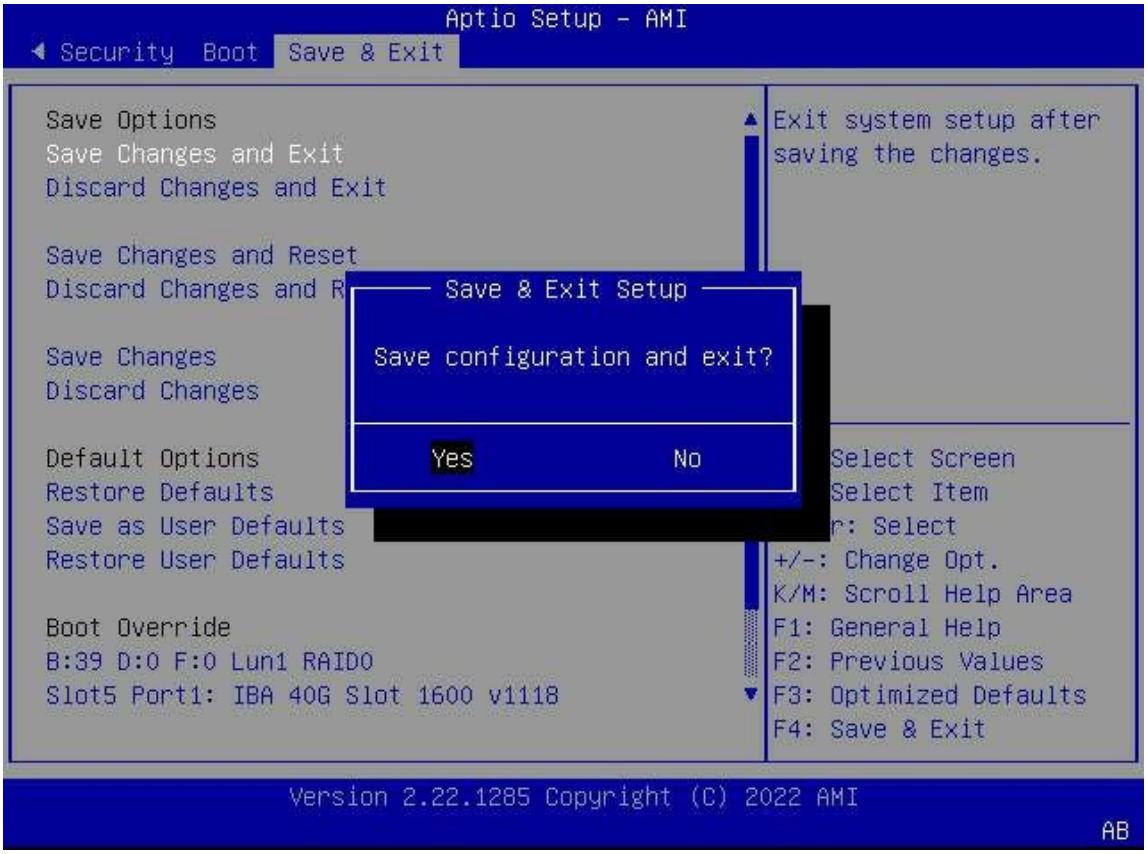

6. Use the arrow keys to select **Yes**, and press **Enter**. The boot mode is set to **UEFI**, and the server is restarted automatically.

Glossary

BBU

- Battery Backup Unit

BIOS

- Basic Input/Output System

BMC

- Baseboard Management Controller

CLI

- Command Line Interface

CPU

- Central Processing Unit

CRC

- Cyclic Redundancy Check

DG

- Disk Group

HBA

- Host Bus Adapter

HDD

- Hard Disk Drive

I/O

- Input/Output

ID

- Identification

IO

- Input & Output

JBOD

- Just a Bunch of Disk

LD

- Logical Disk

NAS

- Network Attached Storage

NVMe

- Non-Volatile Memory Express

OS

- Operating System

BIOS

- Basic Input/Output System

PC

- Personal Computer

PCIe

- Peripheral Component Interconnect Express

PMC

- PCI Mezzanine Card

POST

- Power-On Self-Test

RAID

- Redundant Array of Independent Disks

SAS

- Serial Attached SCSI

SATA

- Serial ATA

SSD

- Solid State Drive

SSH

- Secure Shell

UEFI

- Unified Extensible Firmware Interface

VD

- Virtual Drive

VMD

- Volume Management Device

VROC

- Virtual RAID on CPU