HITACHI
DATA SYSTEMS

Hitachi Freedom Storage™
Thunder 9200™
Sun® Solaris® Host Installation Guide

Document Revision Level

Revision	Date	Description
MK-90DF514-0	January 2001	Initial Release

Source Document Revision Level

The following source document was used to produce this Thunder 9200 host installation guide: *Hitachi Disk Array Subsystem Installation Manual (Fibre Channel)*, revision 1.

Preface

The *Hitachi Freedom Storage™ Lightning 9200™ Sun® Solaris® Host Installation Guide* describes and provides instructions for configuring the devices on the Hitachi Thunder 9200™ array subsystem for operation with the Sun® Solaris® operating system. This configuration guide is written for system administrators responsible for disk management (and/or Hitachi field personnel?) and assumes that:

- the user has a background in data processing and understands direct-access storage device subsystems and their basic functions,
- the user is familiar with the Hitachi Lightning 9200™ array subsystem,
- the user is familiar with the Sun® Solaris® operating system and the Sun® SPARCstation, SPARCserver, SPARCcenter, and/or Ultra Series systems, and
- the user is familiar with the UNIX® file system, system commands, and utilities.

Note: The term "9200" refers to the entire Hitachi Lightning 9200™ subsystem family, unless otherwise noted. Please refer to the *Hitachi Lightning 9200™ User and Reference Guide* (MK-90DF504) for further information on the 9200 disk array subsystems.

For further information on the Sun® Solaris® operating system, please consult the Sun® Solaris® online help and/or user documentation, or contact Sun® technical support.

Contents

List of Figures

List of Tables

Chapter 1 Overview of Thunder 9200™ Sun® Solaris® Configuration

1.1 Sun® Solaris® Configuration

This document describes the requirements and procedures for connecting the 9200 subsystem to Sun® Solaris® server and configuring the new 9200 devices for operation with the Sun® Solaris® server operating system. The Hitachi Data Systems representative performs the physical installation of the 9200 subsystem. The user prepares for 9200 subsystem installation, and then configures the new 9200 devices with assistance as needed from the Hitachi Data Systems representative.

Configuration of the 9200 disk devices for Sun® Solaris®® operations includes:

- Setting LUs and file sizes (see Chapter 2),
- Installing the 9200 subsystem (see Chapter 3),
- Creating and formatting partitions (see Chapter 3), and
- Verifying file system operations (see Chapter 3).

1.2 Hitachi Thunder 9200™ Array Subsystem

The Hitachi Freedom Storage™ Thunder 9200™ RAID subsystem supports concurrent attachment to multiple UNIX®-based and PC-server platforms. Please contact your Hitachi Data Systems account team for the latest information on platform support. The 9200 subsystem provides continuous data availability, high-speed response, scaleable connectivity, and expandable capacity for PC server and open-system storage. The 9200 subsystem can operate with multihost applications and host clusters, and is designed to handle very large databases as well as data warehousing and data mining applications that store and retrieve terabytes of data.

Chapter 2 Preparing for New Device Configuration

2.1 Configuration Requirements

Before using the Hitachi Freedom Storage™ 9200, have the host platform and software installed and ready to use.

Note: Hitachi Data Systems recommends applying the latest patches to your operating system. Contact Sun Microsystems for the latest Solaris patch revision information.

You need superuser privileges to perform the system administrator functions described in this chapter. If you do not have superuser privileges, contact your Sun Solaris system administrator.

This document pertains specifically to the Software Configuration of your host platform. Refer to the *Hitachi Freedom Storage 9200 User and Reference Manual* (MK-90DF504) for hardware specific details and information for connecting and configuring your 9200.

> ***Note:*** Hitachi Data Systems plans to support future releases of the Sun® Solaris® operating system. For the latest information on Sun® Solaris® version support, please contact your Hitachi Data Systems account team.

- Fibre-channel adapters. Make sure to install all utilities, tools, and drivers that come with the adapter(s). For information on driver requirements for the adapters, please refer to the user documentation for the adapter or contact the vendor.
  - The 9200 subsystem supports full-speed (100 MB/s) fibre-channel interface, including shortwave non-OFC (open fibre control) optical interface, and multimode optical cables with SC connectors. Do not connect any OFC-type fibre-channel interface to the 9200 subsystem. For information on supported FC adapters (FCAs), optical cables, hubs, and fabric switches, please contact your Hitachi Data Systems account team or the Hitachi Data Systems Support Center.

2.2 Installing the 9200 Subsystem

The 9200 subsystem comes with all hardware and cabling required for installation.

Note: The Hitachi Data Systems representative must use the 9200 Maintenance Manual during all installation activities. Follow all precautions and procedures in the maintenance manual, and always check all specifications to ensure proper installation and configuration.

2.3 Connecting the Single Controller, Multiple Ports Configuration

To use a single controller, multiple ports configuration, the following rules apply:

- Setting alternate links causes a port fault to become a controller fault
- The same LU cannot be installed from different ports when using a single host. Each LU must be installed individually.
- In a host adapter configuration for a 1:1 connection between host and port, the same LU may be installed

2.4 Connecting Dual Controller Multiple Ports

In a dual controller, multiple ports configuration, a one-to-one pattern in connecting ports between the host adapter and the subsystem, or a one-to-many pattern in connecting the ports between the host adapter and the subsystem via a fibre-channel hub or switch is needed. Use the LU mapping function for this to avoid access competition. Such competition could corrupt data. Access contention may be avoided by operating the host server in the cold standby mode. To ensure availability, configuring with a fibre-channel hub or switch provides distribution to several hubs and switches.

Chapter 3 Configuring the New Devices

Configuration of the new 9200 devices for the Sun® Solaris® involve the following activities:

- Setting the logical units
- Setting and Registering Disk Geometry and Partitions Setting the Logical Units
- Creating file systems
- Setting Auto-Mount

Once the 9200 is installed and connected, you must set and recognize the new LUs by adding the 9200 logical devices to the **sd.conf** file (/kernel/drv/sd.conf). See Figure 3.1 and Figure 3.2.The **sd.conf** file includes the SCSI TID and LUN for all LDEVs connected to the Sun® system. After editing the **sd.conf** file, you will halt the system and reboot.

To set and recognize LUNs:

1. Log in as root, and make a back up of the /kernel/drv/sd.conf file by entering the command: `cp -ip /kernel/drv/sd.conf /kernel/drv/sd.conf.standard`
2. Edit the file /kernel/drv/sd.conf
3. Edit the /kernel/drv/sd.conf file as shown in Figure 3.1. Make sure to make an entry (SCSI TID and LUN) for every new device being added to the Sun® system.
4. Exit the vi editor by entering the command: `ESC + :wq`
5. Shutdown the Sun® system by entering the command: `halt`
6. Reboot the Sun® system by entering the command: `boot -r`
7. Log in to the Sun® system as root, and verify that the system recognizes the 9200 by entering the command: `dmesg | more`. Figure 3.3 shows fibre device recognition.
8. Verify that the vendor name, product name, and number of blocks match the values listed in Figure 3.3.

```
# cp -ip /kernel/drv/sd.conf /kernel/drv/sd/conf/standard        ← Copy the /kernel/drv/sd.
:                                                                  conf file.
#
# vi /kernel/drv/sd.conf                                         ← Enter the file (vi shown).
#ident  "@(#)sd.conf    1.8     93/05/03 SMI"
name="sd" class="scsi"
        target=0 lun=0;

name="sd" class="scsi"
        target=1 lun=0;

name="sd" class="scsi"
        target=2 lun=0;

name="sd" class="scsi"                                           ← Add this information for
        target=2 lun=1;                                          ← all assigned target IDs
                                                                   and LUNs. (*See note.)
name="sd" class="scsi"
        target=3 lun=0;

name="sd" class="scsi"
        target=4 lun=0;
#
# halt                                                           ← Enter halt.
Jan 11 10:10:09 sunss20 halt:halted by root
Jan 11 10:10:09 sunss20 syslogd:going down on signal 15
Syncing file systems... done
Halted

Program terminated
Type help for more information
OK

volume management starting.
The system is ready.

host console login: root                                         ← Log in as root.
Password:Hitachi                                                 ← Enter password (not
                                                                   displayed)
Oct 11 15:28:13 host login: ROOT LOGIN /dev/console
Last login:Tue Oct 11 15:25:12 on console
Sun Microsystems inc.  SunOS 5.5       Generic September 1993
#
#
#
```

Figure 3.1 Setting and Recognizing LUNs

```
# vi␣/kernel/drv/sd.conf ↵
# Copyright (c) 1992, by Sun Microsystems, Inc.
#
#ident       "@(#)sd.conf    1.9     98/01/11 SMI"
name="sd" class="scsi" class_prop="atapi"
     target=0 lun=0;

name="sd" class="scsi" class_prop="atapi
     target=0 lun=1;

name="sd" class="scsi" class_prop="atapi"
     target=0 lun=2;

name="sd" class="scsi" class_prop="atapi"
     target=0 lun=3;

name="sd" class="scsi" class_prop="atapi"
     target=0 lun=4;
```

Figure 3.2 Setting and Recognizing LUNs for Solaris 8

```
# dmesg ↵

fca-pci1: JNI Fibre Channel Adapter model FCI-1063
fca-pci1:     SCSI ID 125 / AL_PA 0x1
fca-pci1: Fibre Channel WWN: 200000e069c00466
fca-pci1: FCA Driver Version 2.2.1.HIT.06.01, March 29,2000 for Solaris 2.5,2.6
fca-pci1: All Rights Reserved.
fca-pci1: < Total IOPB space used: 1140160 bytes >
fca-pci1: < Total DMA space used: 4235293 bytes >
fca-pci1: < DMA redzone len 128 bytes >
fca-pci1: Host: Port 000001 (WWN 200000e069c00466)
fca-pci1: Port 0000ef (WWN 50000e1000002df9) available.
fca-pci1: Target 0 Lun 0: Port 0000ef (WWN 50000e1000002df9) ready.
sd45 at pci1242,46431: target 0 lun 0
sd45 is /pci@1f,4000/fibre-channel@5/sd@0,0
WARNING: /pci@1f,4000/fibre-channel@5/sd@0,0 (sd45)
      corrupt label - wrong magic number

      Vendor 'HITACHI', product 'DFXXX', 10240000 512 byte blocks
fca-pci1: Target 0 Lun 1: Port 0000ef (WWN 50000e1000002df9) ready.
sd67 at pci1242,46431: target 0 lun 1
sd67 is /pci@1f,4000/fibre-channel@5/sd@0,1
WARNING: /pci@1f,4000/fibre-channel@5/sd@0,1 (sd67):
      corrupt label - wrong magic number

      Vendor 'HITACHI', product 'DFXXX', 10240000 512 byte blocks
root on /pci@1f,4000/scsi@3/disk@8,0:a fstype ufs
PCI-device: ebus@1, ebus #0
SUNW,envctrltwo0 at ebus0: offset 14,600000
SUNW,envctrltwo0 is /pci@1f,4000/ebus@1/SUNW,envctrltwo@14,600000
su0 at ebus0: offset 14,3083f8
su0 is /pci@1f,4000/ebus@1/su@14,3083f8
su1 at ebus0: offset 14,3062f8
su1 is /pci@1f,4000/ebus@1/su@14,3062f8
keyboard is </pci@1f,4000/ebus@1/su@14,3083f8> major <37> minor <0>
mouse is </pci@1f,4000/ebus@1/su@14,3062f8> major <37> minor <1>
stdin is </pci@1f,4000/ebus@1/su@14,3083f8> major <37> minor <0>
PCI-device: TSI,gfxp@2, gfxp #0
TSI,gfxp0 is /pci@1f,4000/TSI,gfxp@2
TSI: gfxp0 is GFX8P @ 1152x900
stdout is </pci@1f,4000/TSI,gfxp@2> major <88> minor <0>
se0 at ebus0: offset 14,400000
```

Figure 3.3 Verifying New Devices

3.1 Setting the Disk Geometry

Use the **format** command to register Disk Geometry and partition information (see Figure 3.4). Following this, check the /etc/format.dat file (see Figure 3.5).

```
format ↵
Searching for disks...done

c1t0d0: configured with capacity of 4.83GB
c1t0d1: configured with capacity of 4.83GB


AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
          /pci@1f,4000/scsi@3/sd@0,0
       1. c0t8d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
          /pci@1f,4000/scsi@3/sd@8,0
       2. c0t9d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
          /pci@1f,4000/scsi@3/sd@9,0
       3. c1t0d0 <HITACHI-DFXXX-0000 cyl 264 alt 2 hd 50 sec 768>
          /pci@1f,4000/fibre-channel@5/sd@0,0
       4. c1t0d1 <HITACHI-DFXXX-0000 cyl 264 alt 2 hd 50 sec 768>
          /pci@1f,4000/fibre-channel@5/sd@0,1
Specify disk (enter its number): 3 ↵
selecting c1t0d0
[disk formatted]
Disk not labeled.  Label it no        w? y ↵
```

Figure 3.4 Setting and Registering Disk Geometry and Partitions

```
# tail␣/etc/format.dat ↵
#
disk_type = "HITACHI-DFXXX-0000" ¥
       : ctlr = SCSI : ncyl = 264 : acyl = 2 : pcyl = 266 ¥
       : nhead = 50 : nsect = 768 : rpm = 5400

partition = "DF400-LU00" ¥
       : disk = "HITACHI-DFXXX-0000" : ctlr = SCSI ¥
       : 0 =  1, 5030400 : 1 =  132, 5030400 : 2 =  0, 10137600 ¥
       : 6 =  14, 9600000
```

Figure 3.5 Checking /etc/format.dat

3.2 Setting the Partitions

Use the **format menu** to set the partitions (see Figure 3.6). Use the label command to label the disk (see Figure 3.7).

```
FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit
format> partition ↵

PARTITION MENU:
        0      - change `0' partition
        1      - change `1' partition
        2      - change `2' partition
        3      - change `3' partition
        4      - change `4' partition
        5      - change `5' partition
        6      - change `6' partition
        7      - change `7' partition
        select - select a predefined table
        modify - modify a predefined partition table
        name   - name the current table
        print  - display the current table
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit
partition> print ↵
Current partition table (original):
Total disk cylinders available: 264 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders       Size            Blocks
  0       root    wm       0 -   6      131.25MB    (7/0/0)      268800
  1       swap    wu       7 -  13      131.25MB    (7/0/0)      268800
  2     backup    wu       0 - 263        4.83GB    (264/0/0) 10137600
  3 unassigned    wm       0              0         (0/0/0)          0
  4 unassigned    wm       0              0         (0/0/0)          0
  5 unassigned    wm       0              0         (0/0/0)          0
  6        usr    wm      14 - 263        4.58GB    (250/0/0)  9600000
  7 unassigned    wm       0              0         (0/0/0)          0

partition> 0 ↵
Part      Tag    Flag     Cylinders       Size            Blocks
  0       root    wm       0 -   6      131.25MB    (7/0/0)      268800
```

```
Enter partition id tag[root]: ↵
Enter partition permission flags[wm]: ↵
Enter new starting cyl[0]: 1 ↵
Enter partition size[268800b, 7c, 131.25mb, 0.13gb]: 131c ↵
partition> 1 ↵
Part      Tag    Flag     Cylinders       Size            Blocks
  1       swap    wu       7 -  13      131.25MB    (7/0/0)     268800

Enter partition id tag[swap]: ↵
Enter partition permission flags[wu]: ↵
Enter new starting cyl[7]: 132 ↵
Enter partition size[268800b, 7c, 131.25mb, 0.13gb]: 131c ↵
 :
```

Figure 3.6 Setting the Partitions

```
partition> label ↵
Ready to label disk, continue? y ↵

partition> print ↵
Current partition table (unnamed):
Total disk cylinders available: 264 + 2 (reserved cylinders)

Part  Tag        Flag  Cylinders          Size               Blocks
  0   root       wm    1  131     2.40GB               (131/0/0)       5030400
  1   swap       wu    132  262   2.40GB               (131/0/0)       5030400
  2   backup     wu    0  263     4.83GB               (264/0/0)       10137600
  3   unassigned         wm    0          0            (0/0/0) 0
  4   unassigned         wm    0          0            (0/0/0) 0
  5   unassigned         wm    0          0            (0/0/0) 0
  6   usr        wm    14  263    4.58GB               (250/0/0)       9600000
  7   unassigned         wm    0          0            (0/0/0) 0

partition> quit ↵
 :
format> quit ↵
```

Figure 3.7 Labeling the Disk

3.3 Registering the Disk Geometry and Partition Information

Use the **format** command to verify disk geometry and partition information (see Figure 3.8).

```
# format ↵

Searching for disks...done

c1t0d1: configured with capacity of 4.83GB

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
          /pci@1f,4000/scsi@3/sd@0,0
       1. c0t8d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
          /pci@1f,4000/scsi@3/sd@8,0
       2. c0t9d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
          /pci@1f,4000/scsi@3/sd@9,0
       3. c1t0d0 <HITACHI-DFXXX-0000 cyl 264 alt 2 hd 50 sec 768>
          /pci@1f,4000/fibre-channel@5/sd@0,0
       4. c1t0d1 <HITACHI-DFXXX-0000 cyl 264 alt 2 hd 50 sec 768>
          /pci@1f,4000/fibre-channel@5/sd@0,1
Specify disk (enter its number): 3 ↵
selecting c1t0d0
[disk formatted]
 :
format> partition ↵

PARTITION MENU:
        0      - change `0' partition
 :
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit

partition> name ↵
Enter table name (remember quotes): DFXXX-LU00 ↵

partition> quit ↵


FORMAT MENU:
        disk       - select a disk
 :
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit

format> save ↵
Saving new disk and partition definitions
Enter file name["./format.dat"]: /etc/format.dat ↵
format> quit ↵
```

Figure 3.8 Checking Disk Geometry and Partition Information

After verifying the disk's geometry and partition information, use the **tail** command to check contents of the /etc/format.dat file (see Figure 3.9).

```
# tail␣/etc/format.dat ↵
#
disk_type = "HITACHI-DFXXX-0000" ¥
      : ctlr = SCSI : ncyl = 264 : acyl = 2 : pcyl = 266 ¥
      : nhead = 50 : nsect = 768 : rpm = 5400

partition = "DF400-LU00" ¥
      : disk = "HITACHI-DFXXX-0000" : ctlr = SCSI ¥
      : 0 =  1, 5030400 : 1 =  132, 5030400 : 2 =  0, 10137600 ¥
      : 6 =  14, 9600000
```

Figure 3.9 Verifying Disk Geometry and Partition Information

3.4 Creating the File Systems

Use the **newfs** command to create each file system (see Figure 3.10). Use the **mkdir** command to create a directory for mounting the file system (see Figure 3.11). Use the **mount** command to check the file system (see Figure 3.12).

```
# newfs␣/dev/rdsk/c1t0d0s0 ↵
newfs: construct a new file system /dev/rdsk/c1t0d0s0: (y/n)? y ↵
/dev/rdsk/c1t0d0s0:   Total number of cylinders: 5030400 (131 cylinders, 50 trucks, 768
sectors)
     2456.2MB in 66 cylinder groups (2 c/g, 48.00MB/g, 5952 i/g)
Where to back up super-blocks (for fsck -F ufs -O b= #):
 32, 77600, 155168, 232736, 310304, 387872, 465440, 543008, 620576, 698144,
 775712, 853280, 930848, 1008416, 1085984, 1163552, 1241120, 1318688, 1396256,
 1473824, 1551392, 1628960, 1706528, 1784096, 1861664, 1939232, 2016800,
 2094368, 2171936, 2249504, 2327072, 2404640, 2482208, 2559776, 2637344,
 2714912, 2792480, 2870048, 2947616, 3025184, 3102752, 3180320, 3257888,
 3335456, 3413024, 3490592, 3568160, 3645728, 3723296, 3800864, 3878432,
 3956000, 4033568, 4111136, 4188704, 4266272, 4343840, 4421408, 4498976,
 4576544, 4654112, 4731680, 4809248, 4886816, 4915232, 4992800,
# exit
```

Figure 3.10 Creating a File System

```
# mkdir␣/array1
```

Figure 3.11 Creating a Directory for Mounting the File System

```
# mount␣/dev/dsk/c1t0d0s0␣/array1
```

Figure 3.12 Verifying the File System

3.5 Setting and Verifying the Auto-Mount Parameters

To set Auto-Mount, partitions should be registered into the /etc/vfstab file (see Figure 3.13).

```
# vi␣/etc/vfstab ↵
#device               device          mount   FS      fsck    mount   mount
#to mount             to fsck         point   type    pass    at boot options
#
/dev/dsk/c1d0s2       /dev/rdsk/c1d0s2/usr      ufs   1       yes     -
fd            -               /dev/fd   fd    -       no      -
/proc         -               /proc     proc  -       no      -
①                     ②                         ③       ④     ⑤       ⑥
/dev/dsk/c0t8d0s4     -                         -       swap  -       no      -
/dev/dsk/c0t8d0s0     /dev/rdsk/c0t8d0s0        /       ufs   1       no      -
/dev/dsk/c0t8d0s6     /dev/rdsk/c0t8d0s6        /usr    ufs   1       no      -
/dev/dsk/c0t8d0s3     /dev/rdsk/c0t8d0s3        /var    ufs   1       no      -
/dev/dsk/c0t8d0s7     /dev/rdsk/c0t8d0s7        /export/home ufs2       yes     -
/dev/dsk/c0t8d0s5     /dev/rdsk/c0t8d0s5        /opt    ufs   2       yes     -
/dev/dsk/c0t8d0s1     /dev/rdsk/c0t8d0s1        /usr/openwin ufs2       yes     -
swap          -               /tmp      tmpfs -       yes     -
/dev/dsk/c1t0d0s0     /dev/rdsk/c1t0d0s0        /array1 ufs   3       yes     -
/dev/dsk/c1t0d1s0     /dev/rdsk/c1t0d1s0        /array2 ufs   3       yes     -
```

Figure 3.13 Setting and Verifying the Auto-Mounting Parameters

Table 3.1 Auto-Mount Parameters

Parameter #	Name	Enter
①	Device to mount	Block type device filename
②	Device to fsck	Device file name
③	Mount point	Mount directory name
④	fs type	Type of file system (e.g., UFS, AdvFS)
⑤	Fsck pass	Order of performing file system checks
⑥	Mount at Boot	To mount or not to mount at boot

Chapter 4 Troubleshooting

4.1 Troubleshooting

The Hitachi Freedom Storage™ Thunder 9200™ array subsystem provides continuous data availability. For troubleshooting information for the 9200 subsystem, please refer to the *Hitachi Freedom Storage™ Thunder 9200™ User and Reference Manual* (MK-90DF504).

4.2 Calling the Support Center

If you need to call the Hitachi Data Systems Support Center, make sure to provide as much information about the problem as possible, including the circumstances surrounding the error or failure and the exact content of any error messages displayed on the host system(s). Please check the SVP service information messages (SIMs) using Web access, and note the reference codes and severity levels of the recent messages.

The worldwide Hitachi Data Systems Support Centers are:

- Hitachi Data Systems North America/Latin America
  San Diego, California, USA
  1-800-348-4357
- Hitachi Data Systems Europe
  Contact Hitachi Data Systems Local Support
- Hitachi Data Systems Asia Pacific
  North Ryde, Australia
  011-61-2-9325-3300

Appendix A Acronyms and Abbreviations

FCA	FC adapter
LU	logical unit
LUN	logical unit number
MB	megabytes
OFC	open fibre control
PC	personal computer system
R-SIM	remote SIM
RAID	redundant array of independent disks
SCSI	small computer system interface
SIM	service information message