

Installation Information

GR Multipath Driver (Linux Version)

(November, 2004)

Contents

Supported Operating System (OS) Versions.....	1
Installation for Each OS Version.....	2
Supported Storage Systems	3
ETERNUS6000 storage system	3
ETERNUS3000 storage system	3
ETERNUS GR Storage systems	4
Channel Adapter ID and Connection Points.....	5
ETERNUS6000	5
ETERNUS3000 model 50 and ETERNUS GR710 rear view	5
ETERNUS3000 model 80 and 100 rear view	6
ETERNUS3000 model 200, 300, 400, 500, 600 and 700 rear view	6
ETERNUS GR720 and GR730 rear view	7
ETERNUS GR740, GR820 and GR840	7
Connection Requirements.....	8
GR Multipath Driver	8
FC Card.....	8
Assigned-/Non-assigned-CM Type Storage Systems.....	9
Upgrade from Previous Versions	10
Upgrade from V1.0L00 - V1.0L03 to V1.0L04.....	10
Upgrade from V1.0L03 (to which a T00149-01 urgent repair patch is applied) to V1.0L04 ...	10
Notes.....	11

Trademarks

Microsoft®, Windows®2000, and Windows NT® are registered trademarks of Microsoft Corporation in the United States and other countries. Sun, Sun Microsystems, the Sun Logo, Solaris and all Solaris based marks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries, and are used under license. This document uses the abbreviation "Solaris OE" for Solaris™ Operating Environment.

UNIX is a registered trademark of X/Open Company, Ltd. in the United States and in other countries.

All other names of systems and products mentioned in this documentation not marked with ® or ™ remain the property of their respective companies as trademarks or registered marks.

Supported Operating System (OS) Versions

This table shows the versions of Linux supported by GR Multipath Driver.

Supported OS Versions	Kernel Versions
Red Hat Enterprise Linux ES (v. 3) for Itanium	2.4.21-4.0.1.EL 2.4.21-9.EL 2.4.21-15.EL
Red Hat Enterprise Linux AS (v. 3) for Itanium	2.4.21-4.0.1.EL 2.4.21-9.EL 2.4.21-15.EL
Red Hat Enterprise Linux ES (v. 3)	2.4.21-4.0.1.EL 2.4.21-9.EL 2.4.21-15.EL
Red Hat Enterprise Linux AS (v. 3)	2.4.21-4.0.1.EL 2.4.21-9.EL 2.4.21-15.EL
Red Hat Enterprise Linux ES (v. 2.1)	2.4.9-e.12 2.4.9-e.25 2.4.9-e.27
Red Hat Enterprise Linux AS (v. 2.1) (Previously called 'Red Hat Linux Advanced Server 2.1')	2.4.9-e.8 2.4.9-e.12 2.4.9-e.25 2.4.9-e.27

Installation for Each OS Version

For detailed installation procedures, refer to the manual supplied with the product.

Supported Storage Systems

GR Multipath Driver supports the following storage systems.

- ETERNUS6000
- ETERNUS3000
- ETERNUS GR Storage

ETERNUS6000 storage systems

Storage System	GRMPD	Version Level
ETERNUS6000	for Enterprise Model	V1.0L04 or later

ETERNUS3000 storage systems

Storage System	GRMPD	Version Level
ETERNUS3000 model 50	for Entry Model	V1.0L04 or later
	for Standard Model	
	for Enterprise Model	
ETERNUS3000 model 80	for Entry Model	V1.0L04 or later
	for Standard Model	
	for Enterprise Model	
ETERNUS3000 model 100	for Entry Model	V1.0L04 or later
	for Standard Model	
	for Enterprise Model	
ETERNUS3000 model 200	for Standard Model	V1.0L04 or later
	for Enterprise Model	
ETERNUS3000 model 300	for Standard Model	V1.0L04 or later
	for Enterprise Model	
ETERNUS3000 model 400	for Standard Model	V1.0L04 or later
	for Enterprise Model	
ETERNUS3000 model 500	for Standard Model	V1.0L04 or later
	for Enterprise Model	
ETERNUS3000 model 600	for Standard Model	V1.0L04 or later
	for Enterprise Model	
ETERNUS3000 model 700	for Standard Model	V1.0L04 or later
	for Enterprise Model	

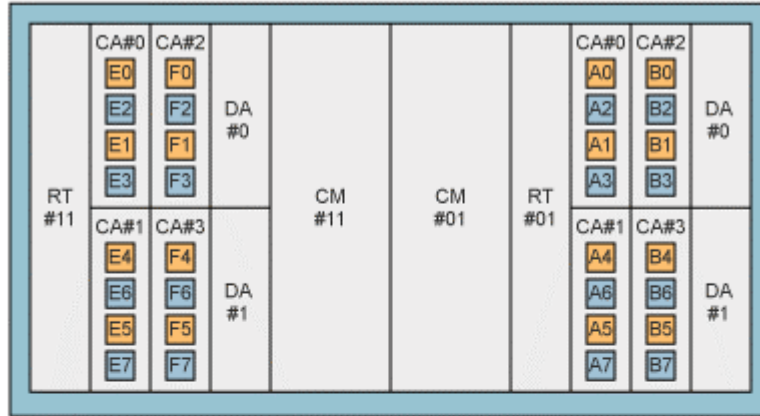
ETERNUS GR Storage Systems

Storage System	GRMPD	Version Level
ETERNUS GR710	for Entry Model	V1.0L02 or later
	for Standard Model	V1.0L02 or later
	for Enterprise Model	V1.0L02 or later
ETERNUS GR720	for Standard Model	V1.0L02 or later
	for Enterprise Model	V1.0L02 or later
ETERNUS GR730	for Standard Model	V1.0L02 or later
	for Enterprise Model	V1.0L02 or later
ETERNUS GR740	for Enterprise Model	V1.0L02 or later
ETERNUS GR820	for Enterprise Model	V1.0L02 or later
ETERNUS GR840	for Enterprise Model	V1.0L02 or later

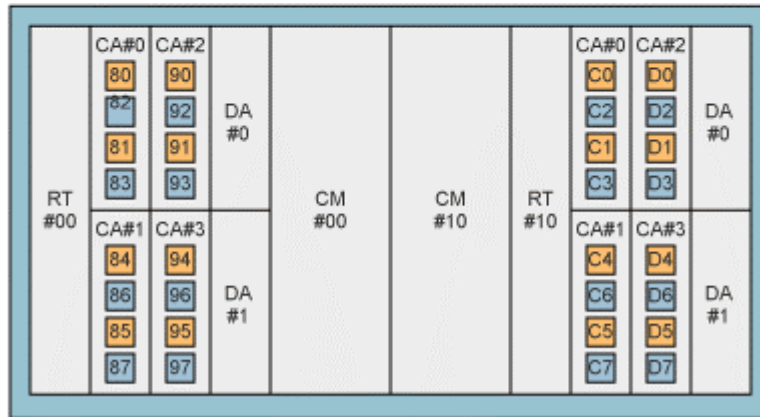
Channel Adapter ID and Connection Points

ETERNUS6000

[Front view]




[Rear view]




■ When using 4Port-CA ■ CAID

ETERNUS3000 model 50 and GR710 rear view

[Rack mount]

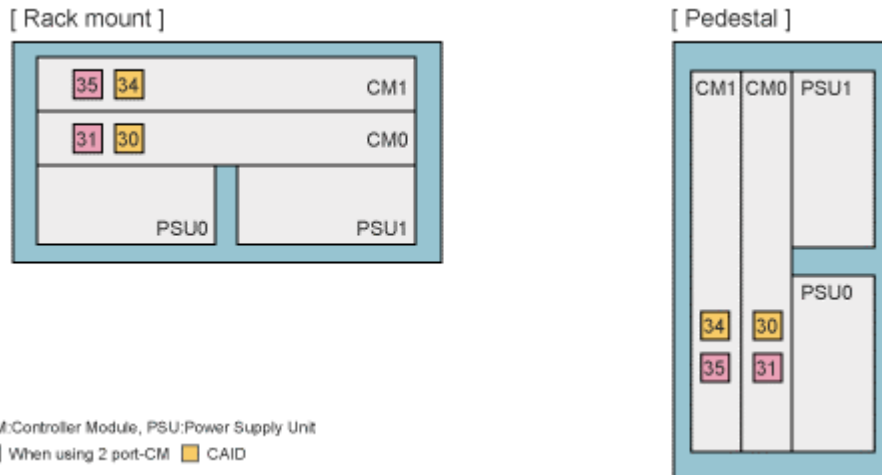


[Pedestal]

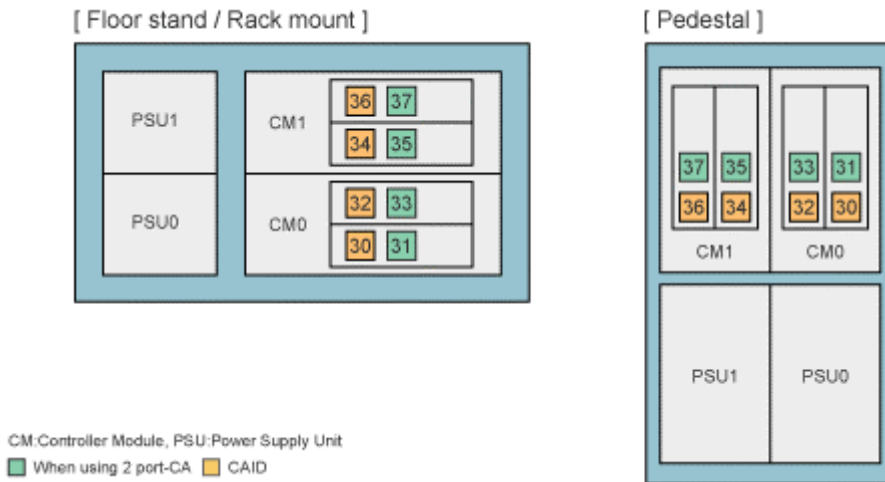


CM: Controller Module, PSU: Power Supply Unit
 ■ CAID


ETERNUS3000 model 80 and 100 rear view



ETERNUS3000 model 200, 300, 400, 500, 600 and 700 rear view




ETERNUS GR720 and GR730 rear view



CM: Controller Module
 CAID

ETERNUS GR740, GR820 and GR840



CM: Controller Module
 CAID

Connection Requirements

The tables below show FC card types, topologies, clustering software and FC card drivers supported by GR Multipath Driver.

They are the same as those supported by GR Device Driver.

Connection Requirements for GRMPD

- FC Card

GRMPD Version Level	FC Card
V1.0L00	PG-FC102
V1.0L01	PG-FC102
V1.0L02	PG-FC102
V1.0L03	PG-FC102, PG-FC105, PG-FC106
V1.0L04	PG-FC102, PG-FC105, PG-FC106, PG-FCD101

- Topology

GRMPD Version Level	Topology
V1.0L00 or later	Point-to-Point, Fabric

- Clustering Software

GRMPD Version Level	Clustering Software
V1.0L01 or later	PRIMECLUSTER

Connection Requirements for FC Card

FC Card	Supported Storage Systems	Topology
PG-FC102	ETERNUS3000	Point-to-Point, Fabric
PG-FC105	ETERNUS6000	
PG-FC106	ETERNUS GR Storage	
PG-FCD101		

Assigned-/Non-assigned CM Type Storage Systems

For ETERNUS3000 and GR Storage, there are two system types: “Assigned-CM” and “Non-assigned-CM.” With Assigned-CM storage systems, the main access path for each LU is assigned to a particular controller. With Non-assigned-CM storage systems, there are no assigned LU access paths as such.

With “Assigned-CM” type storage systems, the path connected to the assigned controller is active. Paths to other controllers are on standby. With “Non-assigned-CM” type storage systems, all paths are active and used for access.

The table below shows the “Assigned-CM” and “Non-assigned-CM” storage systems.

Load balancing/failover performance can differ depending on “Assigned-CM” and “Non-assigned-CM” use and the number of paths employed. For details, refer to the supplied product manual.

Assigned CM type	ETERNUS3000 ETERNUS GR710, GR720 and GR730
Non-assigned CM type	ETERNUS6000 ETERNUS GR740, GR820 and GR840

Upgrade from Previous Versions

Upgrade procedure from previous versions, V1.0L00 - V1.0L03 (to which a T00149-01 urgent repair patch is applied) to V1.0L04 is as follows.

Upgrade from V1.0L00 - V1.0L03 to V1.0L04

Upgrade procedure from V1.0L00 - V1.0L03 to V1.0L04 is as follows.

1. Uninstall the previous version of GRMPD.
 - * For the uninstall procedure, refer to the GRMPD Installation Guides of previous versions.
2. Install the new version: V1.0L04.
 - * For the installation procedure, refer to the GRMPD Installation Guides of V1.0L04.
 - * Do not use rpm -U for such upgrade procedures.

Upgrade from V1.0L03 (to which a T00149-01 urgent repair patch is applied) to V1.0L04

Upgrade procedure from V1.0L03 (to which a T00149-01 urgent repair patch is applied) to V1.0L04 is as follows.

1. Uninstall the previous version of GRMPD.
 - * For the uninstall procedure, refer to the GRMPD Installation Guides of the previous versions
2. Execute the following command through the root (superuser) account.
`# rm -fr /usr/fjsvgmpd (# represents the command prompt)`
3. Install the new version: V1.0L04.
 - * For the installation procedure, refer to the GRMPD Installation Guides of V1.0L04.
 - * Never use rpm -U for such upgrade procedures.

The following is an example statement used when connecting two PG-FC105s to FC-CA0 of CM0 and FC-CA0 of CM1 on an ETERNUS3000 model 100 in FC-AL direct connection. (The driver version is V2.01g)

```
alias scsi_hostadapter1 lpfcdd
alias scsi_hostadapter2 lpfcdd
options lpfcdd lpfc_tgt_queue_depth=40 lpfc_linkdown_tmo=20 lpfc_nodev_tmo=1 lpfc_topology=4
```

- **PG-FC106**

Set up in the same way as PG-FC105.

* For detailed installation procedures for PG-FC106, refer to the document supplied with the PG-FC106 driver.

2. Examples of bootloader setting when using Red Hat Enterprise Linux (v. 3) for Itanium

The following are examples of initial RAM disk creation and bootloader settings for Red Hat Enterprise Linux (v. 3) for Itanium. Make sure you set the proper configuration for your environment. For detailed installation procedures, refer to the installation guide supplied with the product. For detailed information on bootloader settings, refer to the document on elilo.

- **Example of Initial RAM disk (initrd) creation for 2.4.21-9.EL kernel**

```
# /usr/fjsvgmpd/bin/mkinitrd-mpd /boot/efi/efi/redhat/initrd-2.4.21-9.EL-grmpd.img 2.4.21-9.EL
```

- **Example of bootloader (elilo) setting**

Edit the file, /boot/efi/efi/redhat/elilo.conf

When installing GR Multipath Driver

```
image=vmlinuz-2.4.21-9.EL
label=linux
initrd=initrd-2.4.21-9.EL-grmpd.img ← specify the initial RAM disk created
read-only
root=/dev/sda3
```

When uninstalling GR Multipath Driver

```
image=vmlinuz-2.4.21-9.EL
label=linux
initrd=initrd-2.4.21-9.EL.img ← specify the initial RAM disk which existed before GRMPD was
                             implemented
read-only
root=/dev/sda3
```

3. Disk array fixed-recognition-order function

Disk array fixed-recognition-order function is available in GRMPD V1.0L03 (with T00149-01 urgent repair patch applied), V1.0L04 or later. This function makes the system recognize multiple disk arrays in a specific order so that the recognition order remains unchanged each time the system starts. Disk arrays determined as unrecognizable due to a power/path failure are not incorporated in the system. This means that the LUs on the disk array also become unrecognizable*. The purpose of this function is to prevent possible data corruption caused by changing the order the disk arrays are recognized. If such a change does occur, check the disk arrays (their status and connection with servers), restore the failed part, and then reboot the system.

* If critical data is on the unrecognizable LUs, the system/applications may not be able to start.

4. Notes for expansion, reduction and replacement of disk arrays

If any operation listed below is performed on a system where GRMPD (V1.0L03 with T00149-01 urgent repair patch applied), V1.0L04 or later is used, it necessary to boot the system using the variable-recognition-order initrd* and perform the creation procedure for fixed-recognition-order initrd again.

* For detailed creation procedure for fixed-recognition-order initrd, refer to the installation guide for each version and the patch application guide using the following headings for reference.

- Expansion of disk arrays connected to the system
- Reduction of disk arrays connected to the system
- Replacement of disk arrays connected to the system

If there is no variable-recognition-order initrd, follow the procedure below to create a variable-recognition-order initrd.

1. Execute the following command through the root (superuser) account

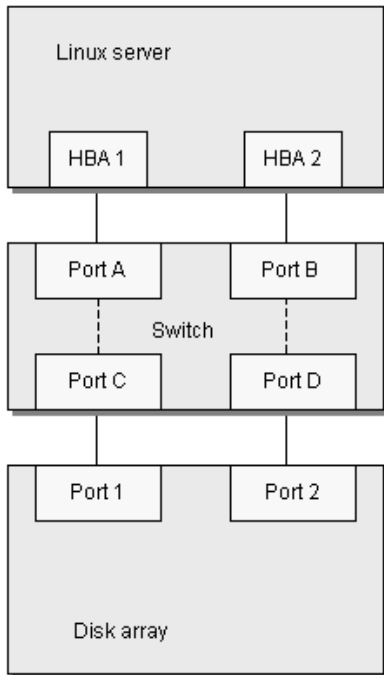
```
# rm -fr /etc/fjsvgmpd/order
```
2. Create variable-recognition-order initrd

```
# /usr/fjsvgmpd/bin/mkinitrd-mpd the name of the variable-recognition-order initrd and the kernel version
```

* For details of mkinitrd-mpd, refer to the installation guide for each version.

5. Path Configuration

With GRMPD for Linux, paths must be connected one-to-one between servers and switches and between switches and disk arrays.



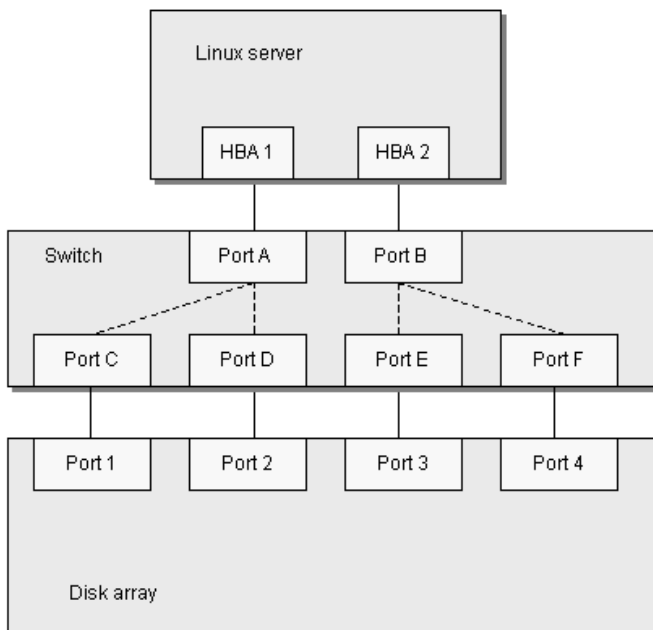
The routes between the Linux server and the disk array are as follows:

- HBA 1 – Port A – Port C – Port 1
- HBA 2 – Port B – Port D – Port 2

This configuration includes two HBAs and two routes to the disk array. Since the numbers of HBAs and routes correspond with each other, a multiple-path configuration can be built correctly.

Note: the dotted lines indicate that Ports A and C, and Ports B and D respectively, correspond with each other.

In a configuration where the path numbers between the server and switch, and between the switch and disk array are different, as shown in the figure below, correct multiple-path configuration is not possible.



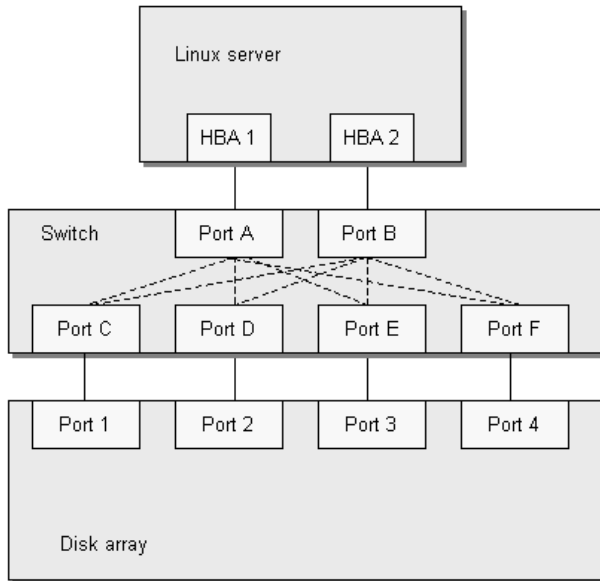
The routes between the Linux server and the disk array are as follows:

- HBA 1 – Port A – Port C – Port 1
- HBA 1 – Port A – Port D – Port 2
- HBA 2 – Port B – Port E – Port 3
- HBA 2 – Port B – Port F – Port 4

This configuration includes two HBAs and four routes to the disk array. With such a topology, correct multiple-path configuration is not possible as the number of HBAs and routes to the disk array do not correspond with each other.

Note: the dotted lines indicate that Ports A and C, Ports A and D, Ports B and E, and Ports B and F, respectively, correspond with each other by zoning.

Without switch based zoning, correct multiple-path configuration cannot be built since the path numbers between the server and switch, and between the switch and disk array are different.

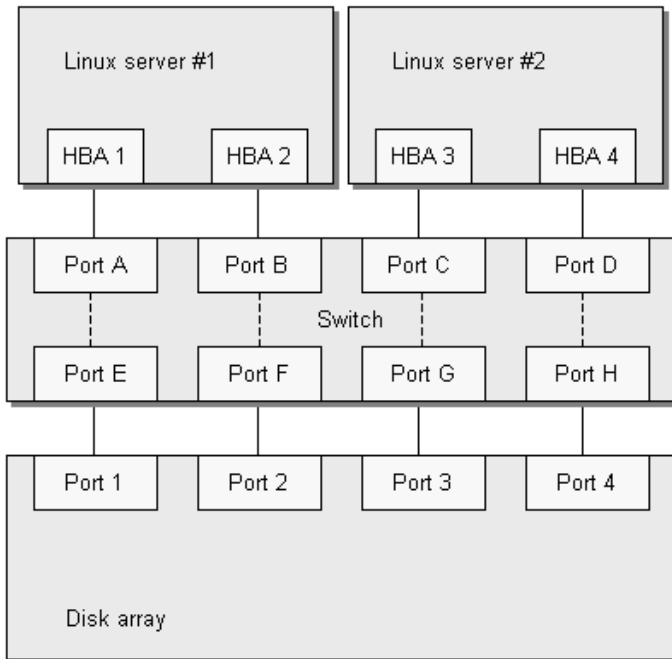


Without switch zoning, the routes between the Linux server and the disk array are as follows:

- HBA 1 – Port A – Port C – Port 1
- HBA 1 – Port A – Port D – Port 2
- HBA 1 – Port A – Port E – Port 3
- HBA 1 – Port A – Port F – Port 4
- HBA 2 – Port B – Port C – Port 1
- HBA 2 – Port B – Port D – Port 2
- HBA 2 – Port B – Port E – Port 3
- HBA 2 – Port B – Port F – Port 4

In this configuration, there are eight routes to the disk array while there are only two HBAs. Since the number of HBAs and routes does not correspond, correct multiple-path configuration is not possible.

The figure below shows a configuration example with two Linux servers and one disk array where correct multiple-path configuration is possible.



The routes between the Linux servers and the disk array are as follows:

Linux server #1

- HBA 1 – Port A – Port E – Port 1
- HBA 2 – Port B – Port F – Port 2

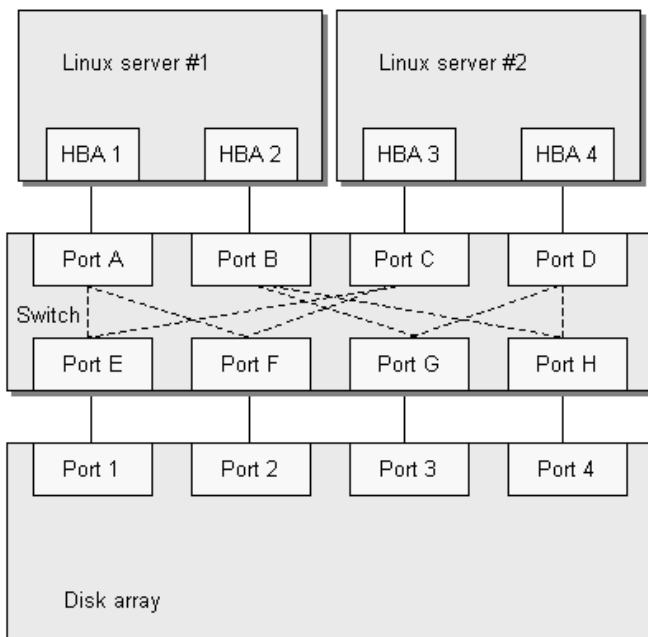
Linux server #2

- HBA 3 – Port C – Port G – Port 3
- HBA 4 – Port D – Port H – Port 4

In this configuration, each server has two HBAs and two routes to the disk array. Since the number of HBAs and routes to the disk array correspond, a correct multiple-path configuration can be built.

Note: the dotted lines indicate that Ports A and E, Ports B and F, Ports C and G, and Ports D and H, respectively, correspond with each other by zoning.

The figure below shows a configuration example with two Linux servers and one disk array where correct multiple-path configuration is not possible.



The routes between the Linux servers and the disk array are as follows:

Linux server #1

- HBA 1 – Port A – Port E – Port 1
- HBA 1 – Port A – Port F – Port 2
- HBA 2 – Port B – Port G – Port 3
- HBA 2 – Port B – Port H – Port 4

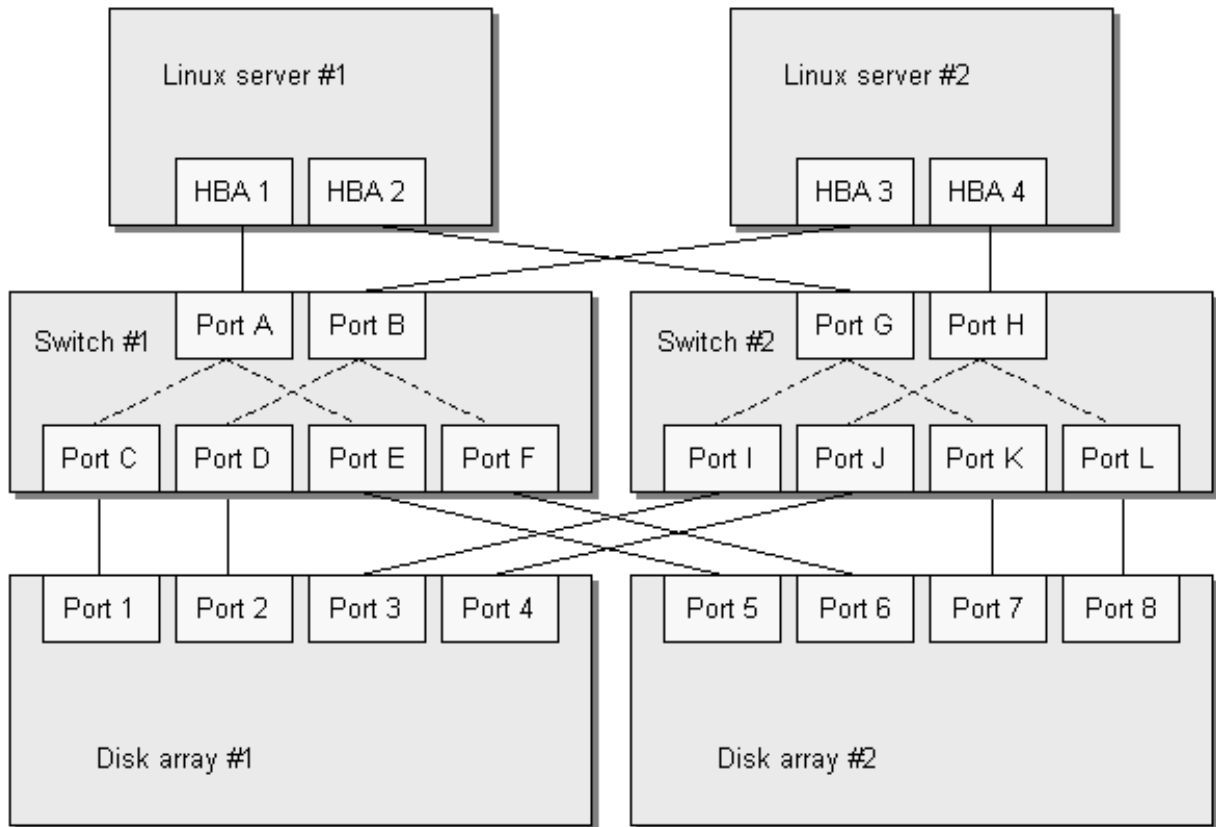
Linux server #2

- HBA 3 – Port C – Port E – Port 1
- HBA 3 – Port C – Port F – Port 2
- HBA 4 – Port D – Port G – Port 3
- HBA 4 – Port D – Port H – Port 4

This configuration includes four routes to the disk array, while each server has only two HBAs. Since the number of HBAs and routes does not correspond, correct multiple-path configuration is not possible.

Note: the dotted lines indicate that Ports A and E, Ports A and F, Ports B and G, Ports B and H, Ports C and E, Ports C and F, Ports D and G, and Ports D and H, respectively, correspond with each other by zoning.

The figure below shows a configuration example with two Linux servers and two disk arrays where correct multiple-path configuration can be built.



The routes between the Linux servers and the disk arrays are as follows:

Between Linux server #1 and Disk array #1:

- HBA 1 – Port A – Port C – Port 1
- HBA 2 – Port G – Port I – Port 3

Between Linux server #1 and Disk array #2:

- HBA 1 – Port A – Port E – Port 5
- HBA 2 – Port G – Port K – Port 7

Between Linux server #2 and Disk array #1:

- HBA 3 – Port B – Port D – Port 2
- HBA 4 – Port H – Port J – Port 4

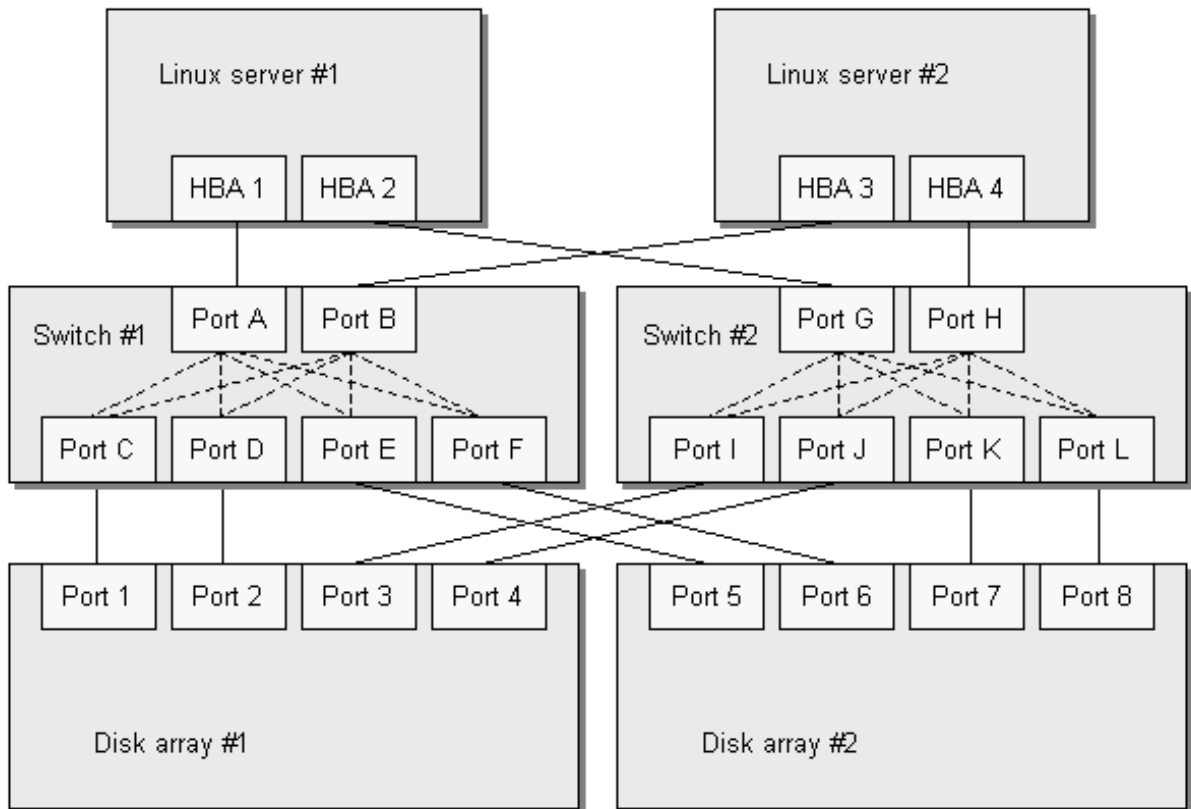
Between Linux server #2 and Disk array #2:

- HBA 3 – Port B – Port F – Port 6
- HBA 4 – Port H – Port L – Port 8

In this configuration, two HBAs and two routes are allocated respectively for each combination of servers and disk arrays. Since the number of HBAs and routes to the disk array correspond, a correct multiple-path configuration can be built.

Note: the dotted lines indicate that Ports A and C, Ports A and E, Ports B and D, Ports B and F, Ports G and I, Ports G and K, Ports H and J, and Ports H and L, respectively, correspond with each other by zoning.

The figure below shows a configuration example with two Linux servers and two disk arrays where correct multiple-path configuration is not possible.



The routes between the Linux servers and the disk arrays are as follows:


- | | |
|--|--|
| Between Linux server #1 and Disk array #1: | Between Linux server #1 and Disk array #2: |
| •HBA 1 – Port A – Port C – Port 1 | •HBA 1 – Port A – Port E – Port 5 |
| •HBA 1 – Port A – Port D – Port 2 | •HBA 1 – Port A – Port F – Port 6 |
| •HBA 2 – Port G – Port I – Port 3 | •HBA 2 – Port G – Port K – Port 7 |
| •HBA 2 – Port G – Port J – Port 4 | •HBA 2 – Port G – Port L – Port 8 |
| Between Linux server #2 and Disk array #1: | Between Linux server #2 and Disk array #2: |
| •HBA 3 – Port B – Port C – Port 1 | •HBA 3 – Port B – Port E – Port 5 |
| •HBA 3 – Port B – Port D – Port 2 | •HBA 3 – Port B – Port F – Port 6 |
| •HBA 4 – Port H – Port I – Port 3 | •HBA 4 – Port H – Port K – Port 7 |
| •HBA 4 – Port H – Port J – Port 4 | •HBA 4 – Port H – Port L – Port 8 |

In this configuration, two HBAs and four routes are allocated respectively for each combination of servers and disk arrays. Since the number of HBAs and routes do not correspond, correct multiple-path configuration is not possible.

Note: the dotted lines indicate that Ports A and C, Ports A and D, Ports A and E, Ports A and F, Ports B and C, Ports B and D, Ports B and E, Ports B and F, Ports G and I, Ports G and J, Ports G and K, Ports G and L, Ports H and I, Ports H and J, Ports H and K, and Ports H and L, respectively, correspond with each other by zoning.

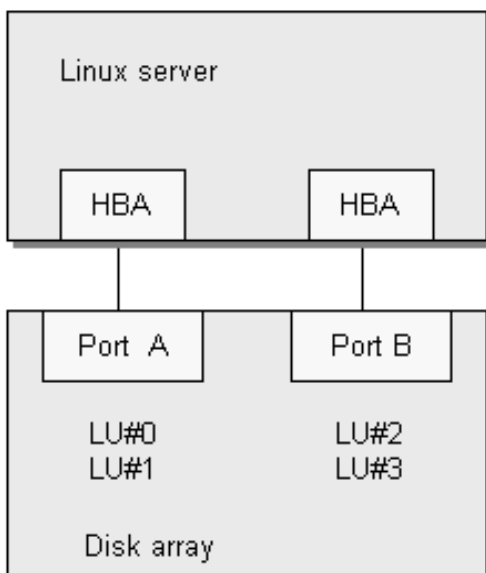
6. LU Configuration

When the server and the disk array are connected using multiple paths, each port must show the same LU configuration to the server.



* An LU configuration where LU#0 and LU# are mapped on both Ports A and B.

In a configuration where one port shows LU#0 and LU#1 while the other port shows LU#2 and LU#3, correct multiple-path configuration is not possible.



* In an LU configuration where LU#0 and LU#1 are mapped on Port A, while LU#3 and LU#4 are mapped on Port B, correct multiple-path configuration is not possible.

About This Installation Information

This Installation Information is devoted to providing technical information and an overview of the basic facilities of GR Multipath Driver. The contents of this document may be modified without any prior notice. Please contact FUJITSU LIMITED if you find any error in descriptions.

FUJITSU LIMITED is not responsible for indemnity that might be caused by the contents in this documentation or any damage related to contents in this documentation.

FUJITSU LIMITED

Global Marketing
Platform Business Group
4-1-1, Kamikodanaka, Nakahara-ku, Kawasaki 211-8588 Japan
E-MAIL: storage-system@fujitsu.com

Web Information Service
"Fujitsu Storage System Web Site" <http://www.storage-system.fujitsu.com/>