



T01

Bladecenter booting from FAStT: Best practices and performance

James Lenaburg

IBM @server xSeries
Technical Conference

Aug. 9 - 13, 2004

Chicago, IL



Bladecenter booting from FASStT: Best practices and performance

Agenda:

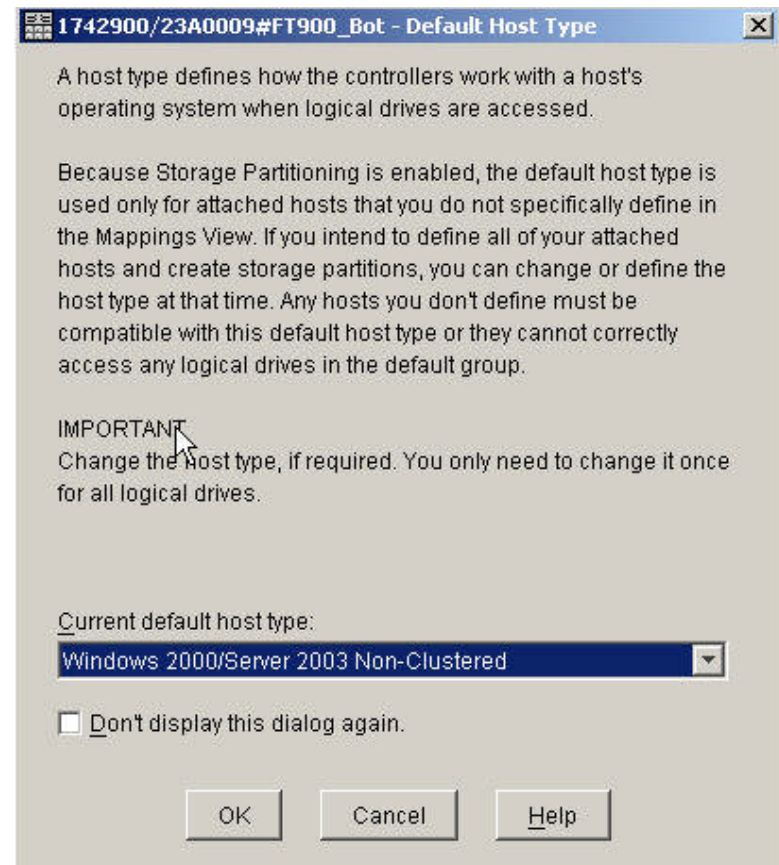
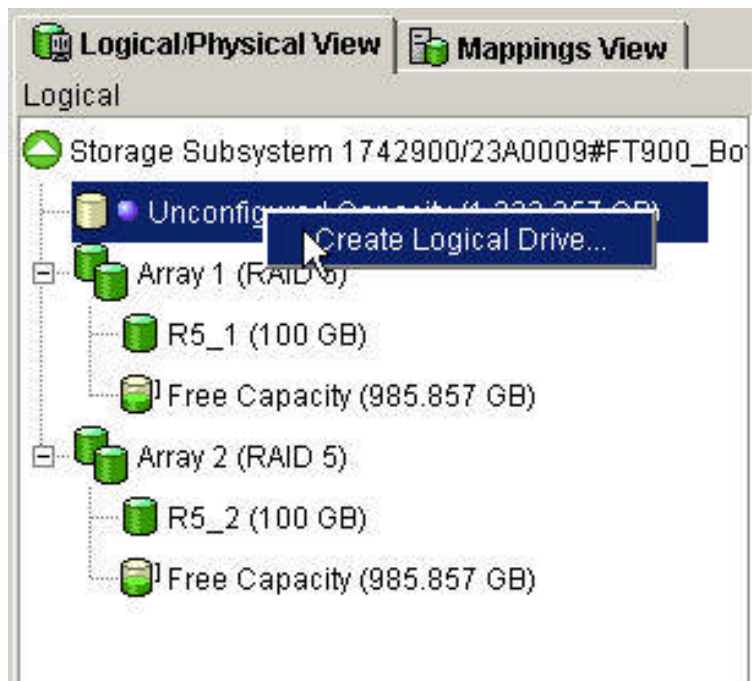
- Zoning recommendations
- How to create LUNs on FASStT and assign them
- How to configure HS20 Fibre Option
- How to Install the OS
- Post installation activities
- Array/LUN performance considerations



Zoning recommendations

- Initially, you can setup single initiator zoning (i.e. Port A to FAStT Controller A, Port B to FAStT Controller B) so the HS20 can see both FAStT controllers.
- For OS installation, you only want the HS20 to see ONE path to the boot LUN.
- Recommend disabling the B path (disable Switch Module 4) and assign all LUNs initially to the A Port and A Controller
- After OS is installed AND RDAC is installed, recommend assigning half of the HS20s to one FAStT controller, and half to the other. Then re-enable Switch Module 4.

How to create LUNs on FASSt and assign them



How to create LUNs on FAStT and assign them - Continued

1742900/23A0009#FT900_Bot - Create Logical Drive Wizard - Specify Logical Drive Parameters

Now you must specify the various parameters for an individual logical drive. From the capacity you previously allocated, indicate exactly how much of that capacity you want to use for the logical drive.

RAID level of array: 1
Maximum logical drive capacity allowed: 67.866 GB

New logical drive capacity: Units:

Name (30 characters maximum):

Advanced logical drive parameters:
 Use recommended settings
 Customize settings (I/O characteristics, controller ownership, logical drive-to-LUN mapping)

< Back Finish Cancel Help

1742900/23A0009#FT900_Bot - Create Logical Drive Wizard - ...

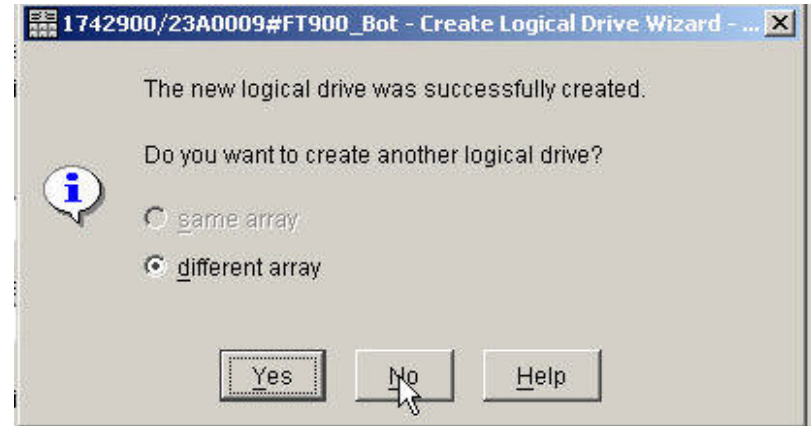
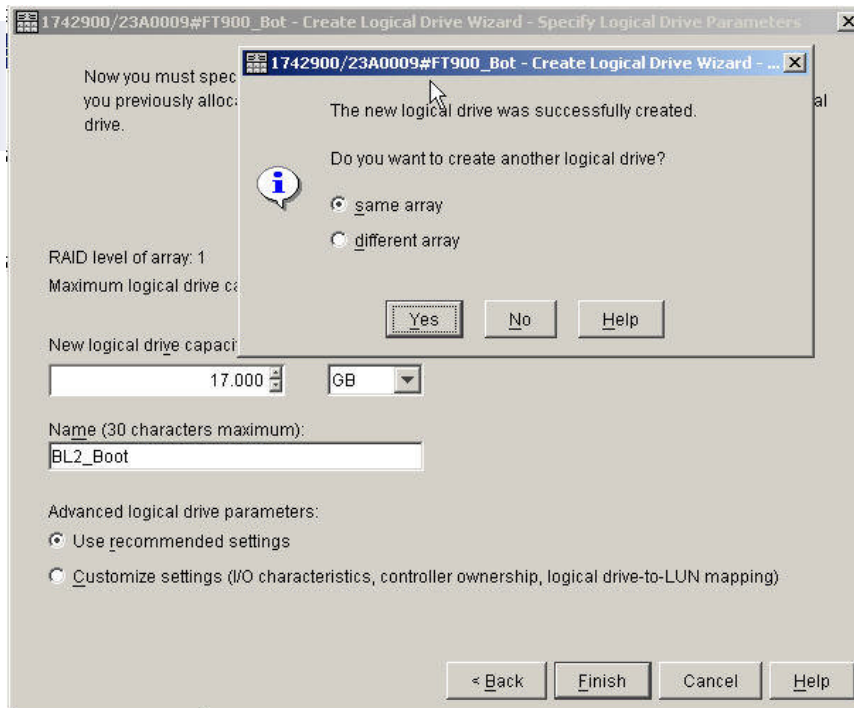
The new logical drive was successfully created.

Do you want to create another logical drive?

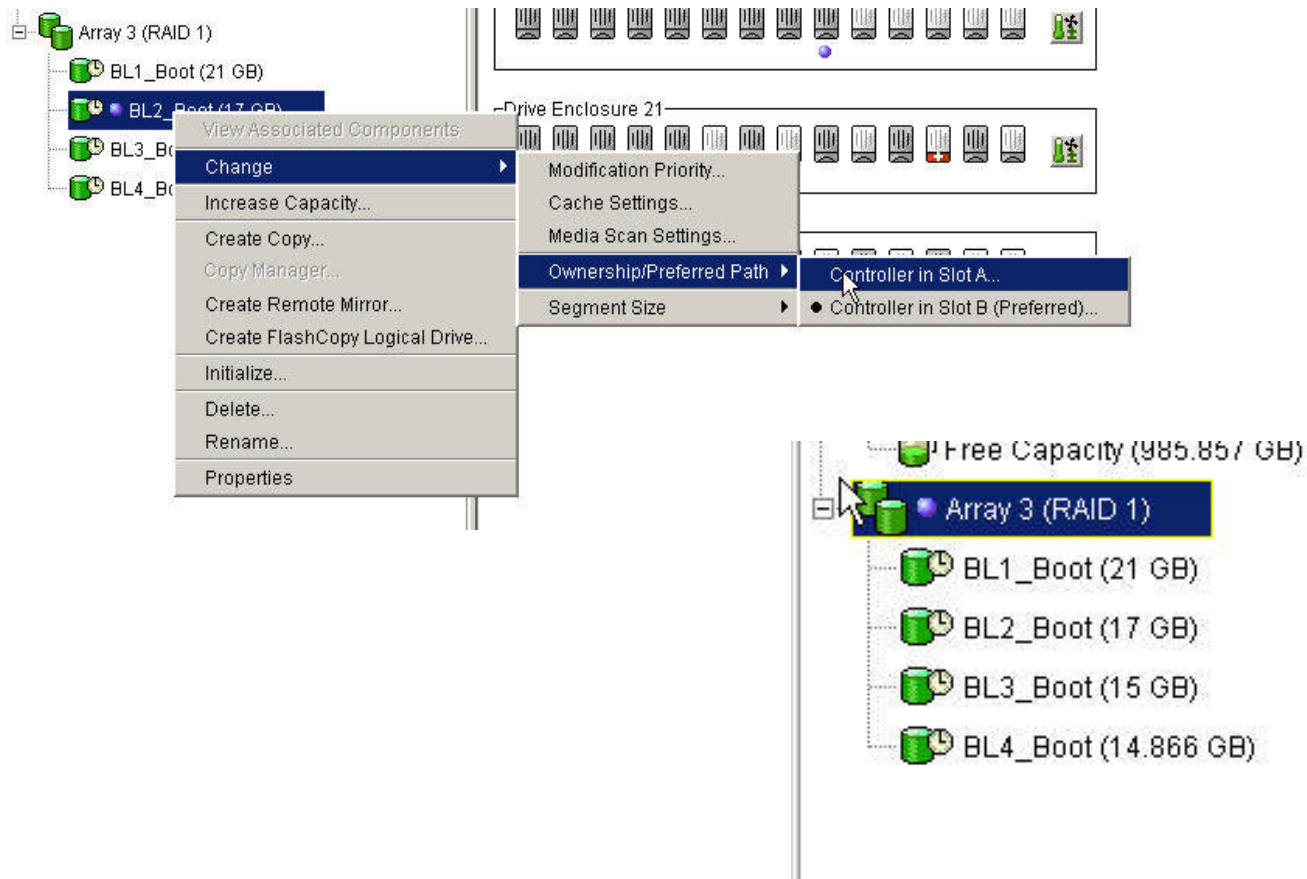
same array
 different array

Yes No Help

How to create LUNs on FASStT and assign them - Continued

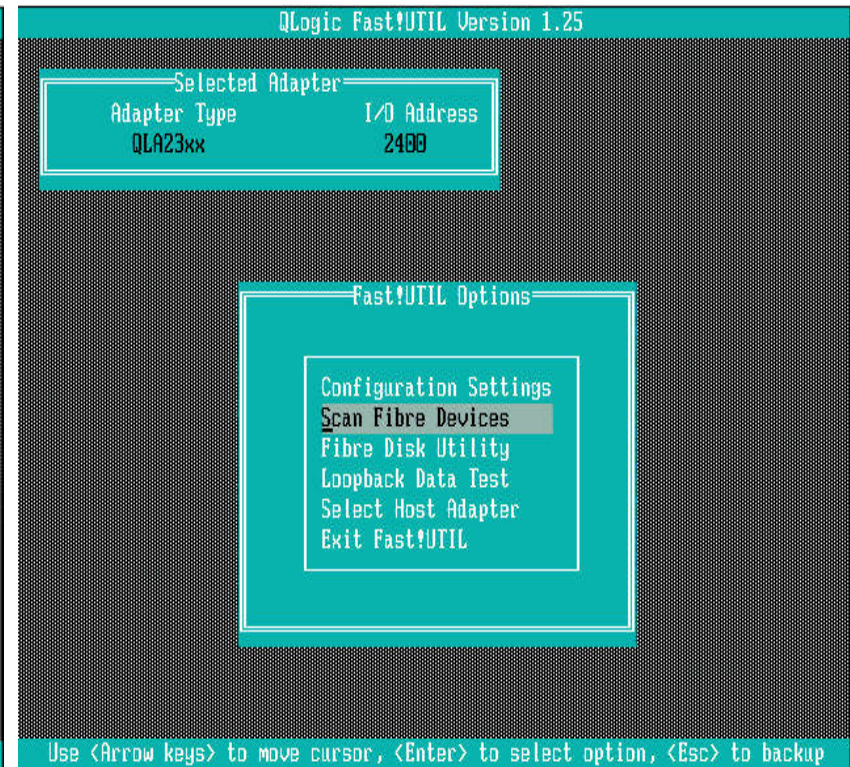



How to create LUNs on FASStT and assign them - Continued



“Tickling the HS20 Fibre Option”

In order for the FASTt to “see” the HBA ports, you need to open the port by doing a Scan Fibre Devices. Do this for both ports.



Assigning the Boot LUNs to the HS20s

1742900/23A0009#FT900_Bot - IBM FASTT Storage Manager 8 (Subsystem Management)

Storage Subsystem View Mappings Array Logical Drive Controller Drive Advanced Help

IBM TotalStorage

Logical/Physical View Mappings View

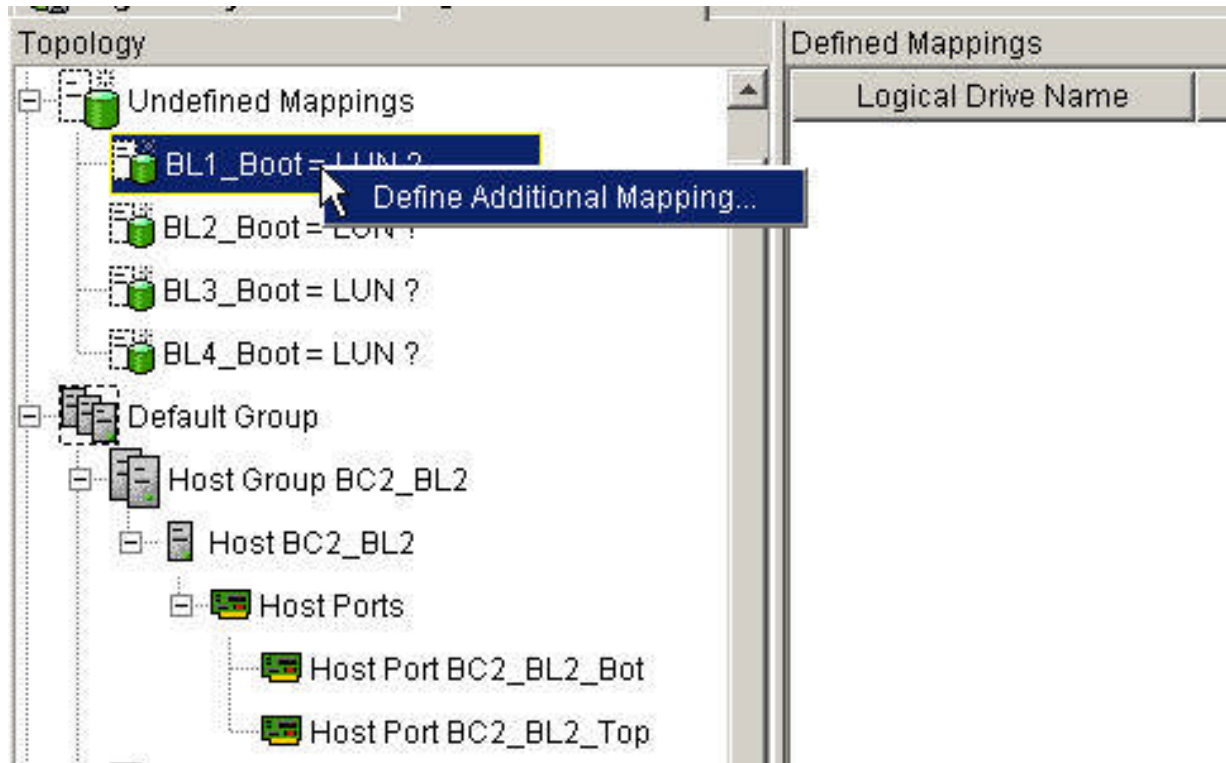
Topology

- Undefined Mappings
 - BL1_Boot = LUN ?
 - BL2_Boot = LUN ?
 - BL3_Boot = LUN ?
 - BL4_Boot = LUN ?
- Default Group
 - Host Group BC2_BL2
 - Host BC2_BL2
 - Host Ports
 - Host Port BC2_BL2_Bot
 - Host Port BC2_BL2_Top
 - Host Group BC2_BL1
 - Host BC2_BL1
 - Host Ports
 - Host Port BC2_BL1_Bot
 - Host Port BC2_BL1_Top
 - Host Group BC2_BL4
 - Host BC2_BL4
 - Host Ports
 - Host Port BC2_BL4_Bot
 - Host Port BC2_BL4_Top
 - Host Group BC2_BL3
 - Host BC2_BL3

Defined Mappings

Logical Drive Name	Accessible By	LUN	Logical Drive Capacity	Type
R5_1	Host Group NATSBA...	1	100 GB	Standard
R5_2	Host Group NATSBA...	2	100 GB	Standard

Assigning the Boot LUNs to the HS20s



Assigning the Boot LUNs to the HS20s

1742900/23A0009#FT900_Bot - Define Additional Mapping

Select a host group or host, logical unit number (LUN), and logical drive to create a logical drive-to-LUN mapping.

Host group or host:

Logical unit number (LUN) (0 to 255):

Logical Drive:

Logical Drive Name	Logical Drive Capacity
Access	
BL1_Boot	21 GB
BL2_Boot	17 GB
BL3_Boot	15 GB
BL4_Boot	14.866 GB

Assigning the Boot LUNs to the HS20s

1742900/23A0009#FT900_Bot - IBM FASTT Storage Manager 8 (Subsystem Management)

Storage Subsystem View Mappings Array Logical Drive Controller Drive Advanced Help

Logical/Physical View Mappings View

Topology

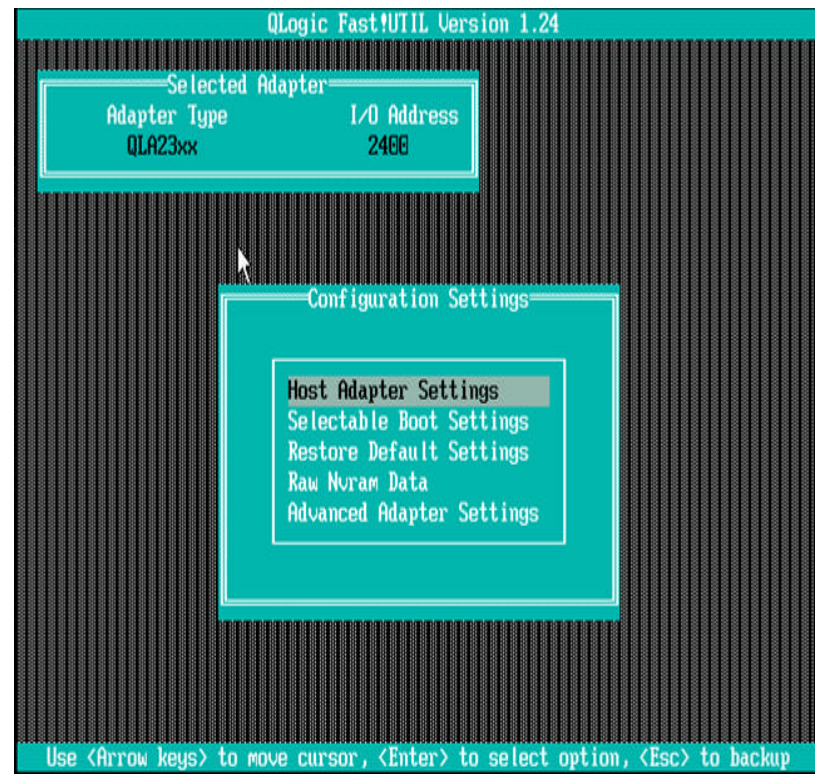

- Storage Subsystem 1742900/23A0009#FT900_Bo
 - Undefined Mappings
 - Default Group
 - Host Group NATSBACK
 - Host Group BC2_BL1
 - Host Group BC2_BL2
 - Host Group BC2_BL3
 - Host Group BC2_BL4

Defined Mappings

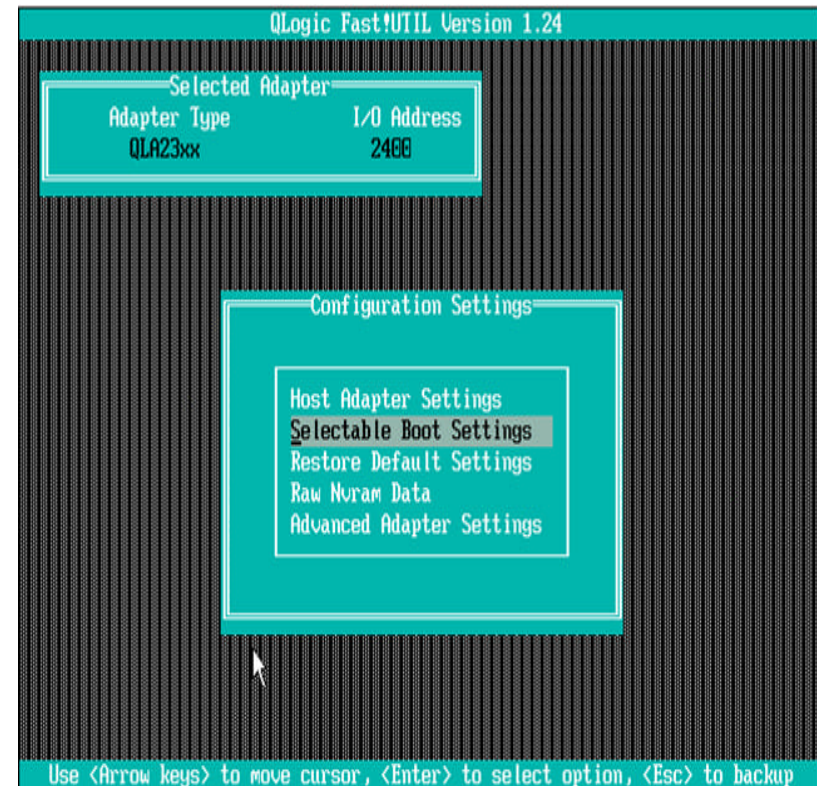
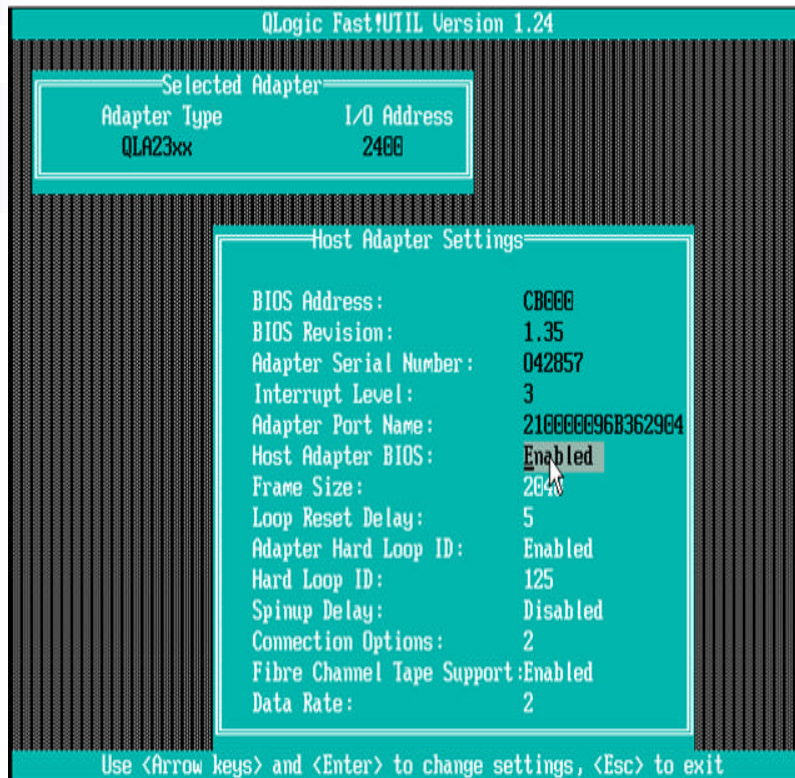
Logical Drive Name	Accessible By	LUN	Logical Drive Capacity	Type
BL1_Boot	Host Group BC2_BL1	0	21 GB	Standard
BL2_Boot	Host Group BC2_BL2	0	17 GB	Standard
BL3_Boot	Host Group BC2_BL3	0	15 GB	Standard
BL4_Boot	Host Group BC2_BL4	0	14.866 GB	Standard
R5_1	Host Group NATSBA...	1	100 GB	Standard
R5_2	Host Group NATSBA...	2	100 GB	Standard

How to configure HS20 Fibre Option

We are going to assign the Boot LUN to HBA Port A



How to configure HS20 Fibre Option



How to configure HS20 Fibre Option

QLogic Fast!UTIL Version 1.24

Selected Adapter

Adapter Type	I/O Address
QLA23xx	2400

Selectable Boot Settings

Selectable Boot: Enabled

(Primary) Boot Port Name,Lun: 0000000000000000, 0

Boot Port Name,Lun: 0000000000000000, 0

Boot Port Name,Lun: 0000000000000000, 0

Boot Port Name,Lun: 0000000000000000, 0

Press "C" to clear a Boot Port Name entry

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

QLogic Fast!UTIL Version 1.24

Select Fibre Channel Device



ID	Vendor	Product	Rev	Port Name	Port ID
128	No device present				
129	IBM	1742-900	0520	200200A0B80F27FB	010000
130	No device present				
131	No device present				
132	No device present				
133	No device present				
134	No device present				
135	No device present				
136	No device present				
137	No device present				
138	No device present				
139	No device present				
140	No device present				
141	No device present				
142	No device present				
143	No device present				

Use <PageUp>/<PageDown> keys to display more devices

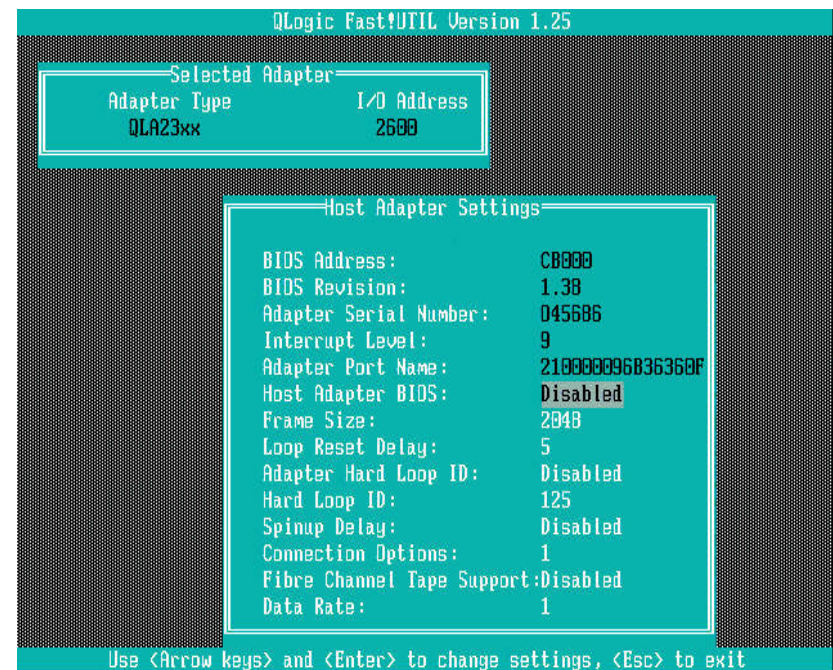

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup

How to configure HS20 Fibre Option

Now we want to assign the Boot LUN to port B also, but we will disable this path during the initial OS install



How to configure HS20 Fibre Option





How to configure HS20 Fibre Option

QLogic FastUTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2500

Host Adapter Settings	
BIOS Address:	CB000
BIOS Revision:	1.38
Adapter Serial Number:	045686
Interrupt Level:	9
Adapter Port Name:	210000096B36360F
Host Adapter BIOS:	Enabled
Frame Size:	2048
Loop Reset Delay:	5
Adapter Hard Loop ID:	Disabled
Hard Loop ID:	125
Spinup Delay:	Disabled
Connection Options:	1
Fibre Channel Tape Support:	Disabled
Data Rate:	1

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

QLogic FastUTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2500

Configuration Settings
Host Adapter Settings
Selectable Boot Settings
Restore Default Settings
Raw Nvram Data
Advanced Adapter Settings

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup



How to configure HS20 Fibre Option

QLogic Fast!UTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2600

Selectable Boot Settings	
Selectable Boot:	Disabled
(Primary) Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Press "C" to clear a Boot Port Name entry	

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

QLogic Fast!UTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2600

Selectable Boot Settings	
Selectable Boot:	Enabled
(Primary) Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Press "C" to clear a Boot Port Name entry	

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit



How to configure HS20 Fibre Option

QLogic Fast!UTIL Version 1.25

Select Fibre Channel Device

ID	Vendor	Product	Rev	Port Name	Port ID
128	No device present				
129	IBM	1742-900	0520	200300A0B0F27FA	020200
130	No device present				
131	No device present				
132	No device present				
133	No device present				
134	No device present				
135	No device present				
136	No device present				
137	No device present				
138	No device present				
139	No device present				
140	No device present				
141	No device present				
142	No device present				
143	No device present				

Use <PageUp/PageDown> keys to display more devices

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup

QLogic Fast!UTIL Version 1.25

Selected Adapter

Adapter Type	I/O Address
QLA23xx	2600


Selectable Boot Settings

Selectable Boot:	Enabled
(Primary) Boot Port Name,Lun:	200300A0B0F27FA, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0
Boot Port Name,Lun:	0000000000000000, 0

Press "C" to clear a Boot Port Name entry

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

How to configure HS20 Fibre Option





Optional Settings If Needed

QLogic Fast!UTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2400

Host Adapter Settings	
BIOS Address:	C8000
BIOS Revision:	1.38
Adapter Serial Number:	045430
Interrupt Level:	3
Adapter Port Name:	210000096036360E
Host Adapter BIOS:	Enabled
Frame Size:	2048
Loop Reset Delay:	5
Adapter Hard Loop ID:	Disabled
Hard Loop ID:	125
Spinup Delay:	Enabled
Connection Options:	1
Fibre Channel Tape Support:	Disabled
Data Rate:	1

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

QLogic Fast!UTIL Version 1.25

Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2400

Host Adapter Settings		
Option	Data Rate	C8000
0	1GB/S	1.38
1	2GB/S	045430
2	Auto Select	3
		210000096036360E
Adapter Hard Loop ID:	Disabled	Enabled
Hard Loop ID:	125	2048
Spinup Delay:	Enabled	5
Connection Options:	1	
Fibre Channel Tape Support:	Disabled	
Data Rate:	1	

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

QLogic Fast!UTIL Version 1.25

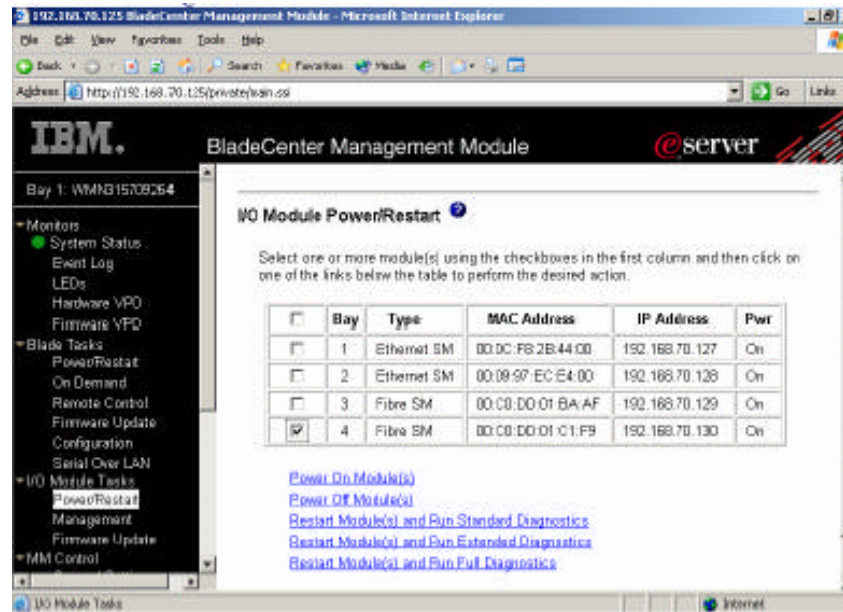
Selected Adapter	
Adapter Type	I/O Address
QLA23xx	2400

Host Adapter Settings			
Option	Type of Connection	Adapter Hard Loop ID:	Disabled
0	Loop only	Hard Loop ID:	125
1	Point to point only	Spinup Delay:	Enabled
2	Loop preferred, otherwise point to point	Connection Options:	1
		Fibre Channel Tape Support:	Disabled
		Data Rate:	1

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

Disabling Path B

- After making sure all of your LUNs are on FASStT controller A, then you should disable Switch Module B, so there is only one path to the Boot LUN



BladeCenter Management Module

IO Module Power/Restart

Select one or more module[s] using the checkboxes in the first column and then click on one of the links below the table to perform the desired action.

<input type="checkbox"/>	Bay	Type	MAC Address	IP Address	Pwr
<input type="checkbox"/>	1	Ethernet SM	00:0C:F8:2B:44:00	192.168.70.127	On
<input type="checkbox"/>	2	Ethernet SM	00:09:9F:EC:E4:00	192.168.70.128	On
<input type="checkbox"/>	3	Fibre SM	00:C0:DD:01:BA:AF	192.168.70.129	On
<input checked="" type="checkbox"/>	4	Fibre SM	00:C0:DD:01:C1:F9	192.168.70.130	On

[Power On Module\(s\)](#)
[Power Off Module\(s\)](#)
[Restart Module\(s\) and Run Standard Diagnostics](#)
[Restart Module\(s\) and Run Extended Diagnostics](#)
[Restart Module\(s\) and Run Full Diagnostics](#)



How to Install the OS

- You must use the HS20 W2003 Qlogic Driver diskette during the install (Press F6 for boot disk)
- During the POST sequence you will notice the following text:
Drive letter C: is moved to the Drive letter D:
Loop ID 129,0 is installed as Drive C:
- This indicates that the Blade is now booting from the LUN 0 associated to the first HBA adapter



How to Install the OS

```
Broadcom NetXtreme Ethernet Boot Agent v3.1.15
Copyright (C) 2008-2002 Broadcom Corporation
All rights reserved.

QLogic Corporation
QLA2312 PCI Fibre Channel ROM BIOS Version 1.29      Subsystem Vendor ID: 1014
Copyright (C) QLogic Corporation 1993-2002. All rights reserved.
www.qlogic.com

Press <CTRL-Q> for FastUTIL
ISP2312 Firmware Version 3.01.12
QLogic adapter using IRQ number 3
QLogic adapter using IRQ number 9


-----
Drive Letter C: is Moved to Drive Letter D:
LODP ID 129,0 is Installed As Drive C:
-----
Device Device  Adaptor Port  Lun  Vendor  Product  Product
Number Type   Number  ID  Number  ID      ID      Revision
  88  Disk     0      018F80  0      IBM     1742    052B
ROM BIOS Installed
```



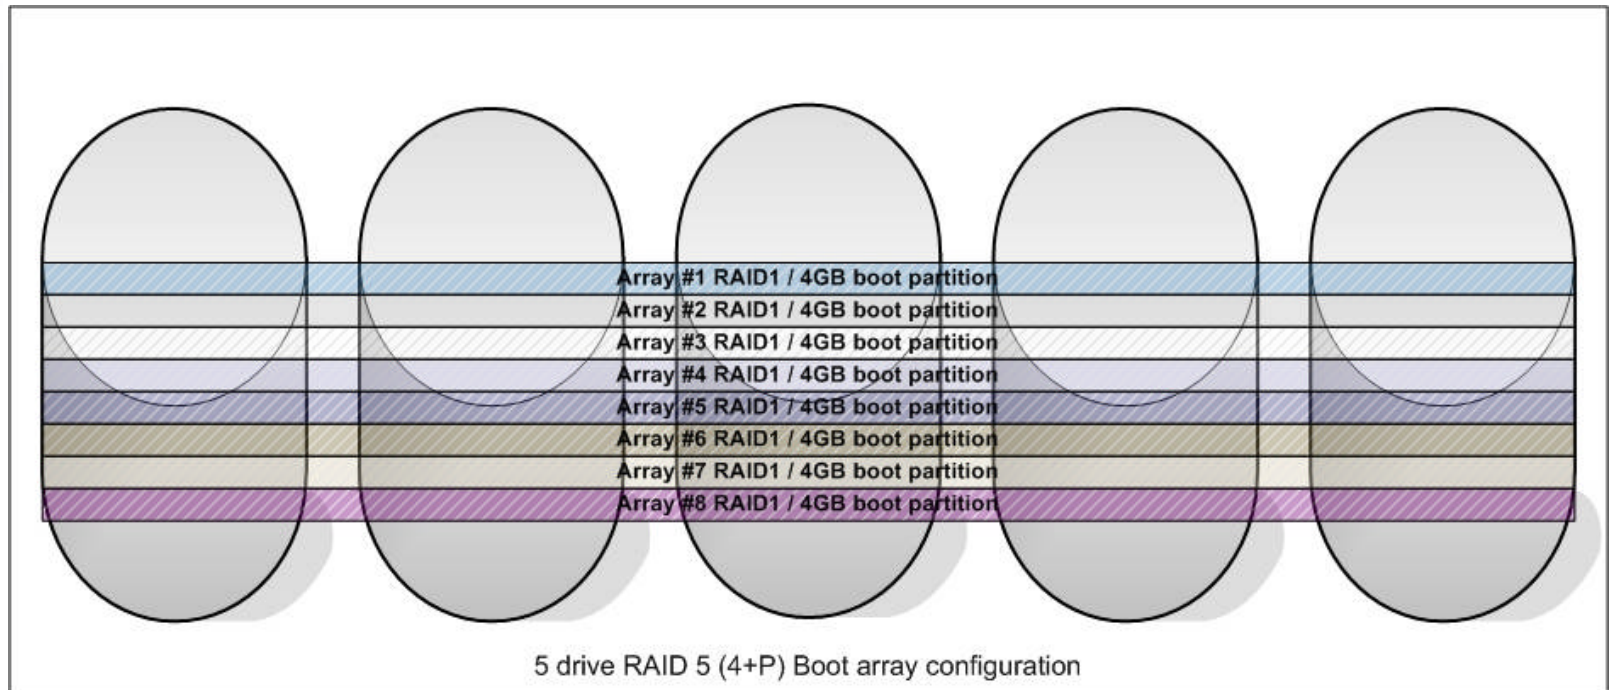
Post installation activities

- Install RDAC 8.42 for W2003
- Re-enable Switch Module 4

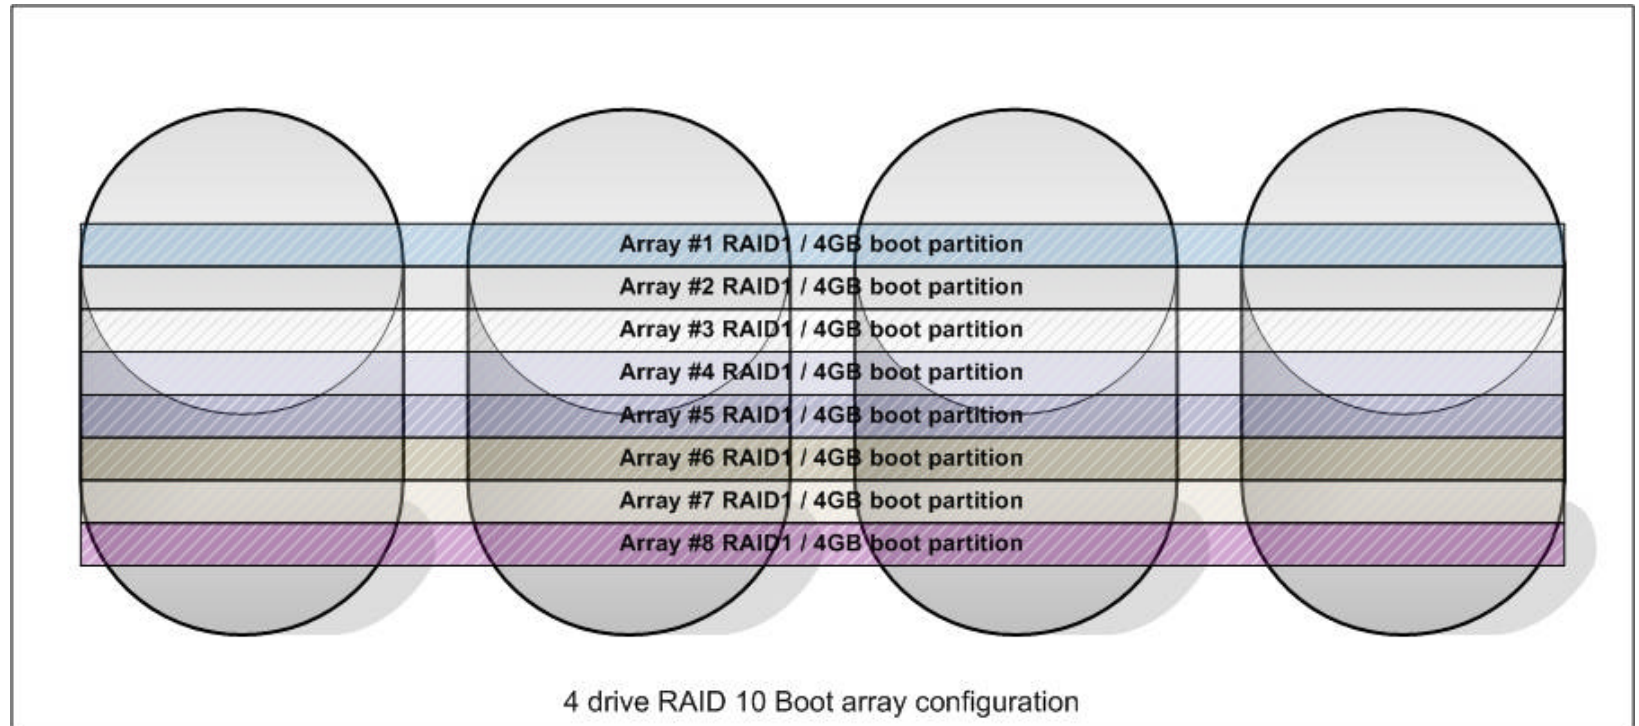
Array/LUN performance considerations



Array/LUN performance considerations



Array/LUN performance considerations





















Array/LUN performance considerations

- Helpcenter “supports” 4 LUNs per array
- ATS tested 8 LUNs per 2 drive R1 array on FT200 – Good performance
- ATS tested 8 LUNs per 5 drive R5 array on FT200 – Good performance
- ATS tested 8 LUNs per 4 drive R10 array on FT200 – Good performance
- Some customers using 14 LUNs per 4 drive R10 array on FT900 with no problems

Simultaneous FT900 14 Blade Boot

FAST900_Bot_SANBOOT_Test - Performance Monitor

Devices	Total I/Os	Read Percentage	Cache Hit Percentage	Current KB/second	Maximum KB/second	Current IO/second	Maximum IO/second
CONTROLLER IN SLOT A	87,324	82.8	41.5	19,305.0	28,083.6	1,032.0	1,762.0
CONTROLLER IN SLOT B	0	0.0	0.0	0.0	0.0	0.0	0.0
 Logical Drive BL10_Boot	1,855	99.4	40.8	124.0	1,658.6	28.8	144.8
 Logical Drive BL11_Boot	1,819	99.4	42.7	115.2	1,663.6	26.6	198.0
 Logical Drive BL12_Boot	1,554	99.5	49.0	256.8	1,643.6	35.4	218.8
 Logical Drive BL13_Boot	1,377	99.9	49.9	0.0	1,756.0	0.0	152.4
 Logical Drive BL14_Boot	1,270	100.0	51.0	470.4	1,003.2	32.2	71.8
 Logical Drive BL1_Boot	6,316	75.2	97.7	0.0	16,239.6	0.0	505.2
 Logical Drive BL2_Boot	8,779	86.0	33.8	180.0	15,396.2	29.8	515.8
 Logical Drive BL3_Boot	15,961	71.6	46.1	3.2	13,773.2	0.6	819.4
 Logical Drive BL4_Boot	15,957	71.7	41.6	12.2	15,606.8	2.2	841.2
 Logical Drive BL5_Boot	7,050	84.5	34.7	225.6	5,074.0	41.4	312.0
 Logical Drive BL6_Boot	8,416	85.7	38.7	370.6	6,180.8	88.4	568.4
 Logical Drive BL7_Boot	3,845	99.3	37.6	13,129.8	13,129.8	417.6	417.6
 Logical Drive BL8_Boot	7,322	91.4	30.6	3,139.6	16,629.6	183.2	486.8
 Logical Drive BL9_Boot	5,803	97.4	31.3	1,277.6	14,327.6	145.8	447.4
 Logical Drive W2003_Master_RD...	0	0.0	0.0	0.0	0.0	0.0	0.0
 Logical Drive W2K_Master_RDA...	0	0.0	0.0	0.0	0.0	0.0	0.0
STORAGE SUBSYSTEM TOTALS	87,324	82.8	41.5	19,305.0	28,083.6	1,032.0	1,762.0

Start: 7/23/04 2:38:24 PM Stop: Time Monitored: 00:09:57

Buttons: Stop, Update, Settings..., Save As..., Close, Help



Gotchas

- Easiest to implement if all boot LUNs on the same controller
- You can setup for boot LUNs to be owned by alternate controllers – BUT
- You must set servers booting from CTRL-B to have CTRL-B LUN listed first in Boot Settings order
- You may see LUNs move at boot time on some servers (I observed this), but once the servers are up, and RDAC has been properly loaded, the LUN will move back to the right controller.



Suggestions

- Smaller boot LUN (4GB-6GB)
- Create a swap LUN for SWAP file
- Data LUNs for data files
- Use Volumecopy to save master LUNs OR
- Use RDM to save Boot LUN image for deployment



Update 8/11/04

- Assign BOTH controllers to the HBA Ports under Assignable Boot Setting screen.
- Put the controller that is the preferred path for the boot LUN as the first entry, put the alternate controller as the second entry
- Choose a host port type with ADT Enabled OR run the ADT Enable script. Windows2000/2003 Non-Clustered (DMP support) has ADT Enabled
- Now, the HBA will be able to use both of the paths to the boot LUN and move the LUN if needed.
- (i.e.) If you have the boot LUN on the preferred path to the A Controller, and you lose the path to the A controller for some reason, then the HBA BIOS will initiate I/O down the path and signal the FAStT controller to move the LUN to the other path, and it will then try booting from the alternate path.

Questions?

