

Chapter 15

-

Fibre Channel Mass Storage (FCMS)

INDEX

Fibre Channel Basics	3
Topologies	4
Point-to-Point (P2P)	4
Fabric.....	5
Arbitrated Loop.....	5
FC HBAs, Software and Patches	8
Fibre Channel Kernel Driver & Tunables	10
Fibre Channel Addressing in HP-UX	12
Peripheral Device Addressing (PDA)	14
Logical Unit Addressing (LUN)	16
Volume Set Addressing (VSA).....	18
Example: Private Loop with FC-SCSI Mux	21
Example: Private Loop versus Public Loop.....	23
Troubleshooting Utilities fcmsutil, tdutil, tdlist, tddiag	26
fcmsutil.....	26
tdutil, tdlist, tddiag	29
How to Replace Disks at Tachyon TL/TS/XL2 HBAs	29
Fibre Channel Storage Devices	31
Glossary	34
Additional Information	36

Fibre Channel Basics

Fibre Channel is a communication protocol that has been developed to satisfy the needs of a growing demand for data throughput. The advantages are as follows:

- different channels and protocols over one physical interface
- high bandwidth (200MB/s or higher)
- flexible setup (topology)
- connections across great distances (a couple of kilometers)
- supports different speeds, media and connections

generally spoken fibre channel combines the advantages of channel and networking technology. A channel is a limited, direct, structured and predictable mechanism for data transmission between a few participants. A channel is typically used in situations where a high transfer rate is needed (e.g. connection of peripheral devices like disks, tape drives, printer or workstations). Well known channel protocols are e.g. SCSI or HIPPI.

Networks are non-structured and unpredictable. They are able to adapt automatically to changing environments and allow a higher number of participants. Hence a higher administration effort (mostly in software) is necessary in order to establish a connection between two points in the network. Thus networks are slower than channels. Well known network protocols are Ethernet, Token Ring or FDDI.

In this chapter we are talking solely about the usage of fibre channel in order to address mass storage devices. This is called FCMS (fibre channel mass storage)

FC is used as the transmission media for SCSI connections. The SCSI protocol is somehow encapsulated.

The essential advantages of FC compared to SCSI are

	SCSI	FC
Distance	25m	100km
Speed	20MB/s	200MB/s
Number of devices	15	126 ports in AL. Much more with fabric
Reliability		extremely low error rate
Robustness		no HW damages by pulling cables

NOTES:

The maximum supported distance between two FC hubs/switches without using special HW like DWDMs or Power-GBICs is 10km. The maximum distance from a device port (N or NL) to another port (F, FL or NL) is 500m for 1Gbit/s and 300m for 2Gbit/s. The cable attenuation needs to be measured in all cases.

FC media	9µm single mode fibre	long wave (10km, 1300nm laser)
	50µm multimode fibre	short wave (500m, 780nm laser)
	62,5µm multimode fibre	FDDI (175m) (possibly mass storage)

The cables should not be bended too strong because they could break. The minimal radius

is 3cm. According to the *can of coke rule* the limitational radius when winding a FC cable is the one of a can of coke.

Topologies

A physical connection at a port consists of two separate fibers because FC is a serial protocol. The fibers for receive (RX) and transmit (TX) are bundled in one cable, for HP mostly of orange color.

Basically you are able to distinguish between the different topologies by two criteria:


- 1) do we have a loop?
- 2) is there a switch connected?

loop	fabric	topology
yes	no	private (arbitrated) loop
yes	yes	public loop
no	no	direct point-to-point
no	yes	switched point-to-point (*)

(*): also known as direct fabric attach (DFA)

Point-to-Point (P2P)

In the p2p case (no loop) two FC devices are directly interconnected with each other. The transmitting diode of one device is connected to the receiving diode of the other and vice versa. The full bandwidth can be used for data transmission. The initialization of the link (login) is quite easy. HP did only implement the switched P2P topology. P2P without fabric can be regarded as a two port loop, i.e. the overhead of the arbitrated loop protocol is used even if there are only 2 ports.




advantages	disadvantages
dix, full bandwidth for the link	high expense in terms of hardware
	no scalability

Fabric

Using the fabric topology you can have up to 2^{24} (approx. 16,7 million) nodes in a meshed configuration.

The advantage of this topology is the fact that multiple devices can communicate with each other simultaneously, each granted the full bandwidth. The use of FC switches is mandatory.



If a *N_Port* registers at a switch it is given a native address-ID (*S_ID*) by the switch (details in the FC addressing part). Further features of a fabric are multicast server, broadcast server, alias server, quality of service facilitator and directory server. Some switches have *FL_Ports* in order to run arbitrated loops.


advantages	disadvantages
high scalability, i.e. many devices (up to 16 mio)	higher initial cost compared to AL
multiple devices communicate simultaneously	slightly lower throughput compared to AL
loss of one component does not interrupt the link	
full bandwidth for each switch port	
performance only minor dependant to length	

NOTE:

Fabric mode is supported as of UX 11.00 Tachyon TL driver version \geq B.11.00.03 and FCMS patch [PHKL_21381](#).

Arbitrated Loop

The Arbitrated Loop (AL) topology is able to connect up to 126 ports that share the total bandwidth. All nodes act as loop devices.



AL is not a „token passing“ protocol. If a port within an AL likes to transmit data it needs to arbitrate the loop in order to obtain the control over the loop. It sends the ARBx signal (arbitrate primitive; x=AL_PA of the ports). When this ARBx signal returns the port has the control over the loop. After it sent the OPN signal (open primitive) to the target port, a quasi point to point connection between the two ports has been established. The other ports act as

repeater.

If more than one port arbitrates the loop at the same time, the port with the lower AL_PA (i.e. higher loop id) wins. Only if this port abandons the control over the loop, the other ports will be able to arbitrate the loop again. In contrast to a token passing protocol there is no limitation in how long a port can have control over the loop. An access fairness algorithm may be used optionally in order to give all other nodes the possibility to arbitrate the loop before the same node gets access again..

Like most of the ring topologies the setup of an AL is simplified by using **hubs**, because standardized cables can be used. A hub is able to detect connected or disconnected devices. A faulty device or a defect cable does no longer tear down the whole network.

Advantages	Disadvantages
cost effective	all nodes share the total bandwidth
scalability	up to 126 nodes per loop
	if one component fails a new initialization is needed
	performance is strong dependant of length and number of nodes

There are private und public loops.

A **private Loop** corresponds to the known implementation of the FC-AL. The addressing is done with the der AL_PA, an 8 bit identifier. All nodes in the loop can talk to each other but not to nodes outside of the loop.

In a **public Loop** the nodes register with their *World Wide Name* at the FL-port of the switch and get a 3 byte identifier assigned, i.e. they are capable of the **fabric login**. The lower byte corresponds to the AL_PA for the communication within the loop whereas the upper 2 bytes characterize the connected switchport. The name server table of the switch stores the mapping between world wide names and 3 byte identifiers. This enables public loop devices to communicate with nodes outside of the loop.

LoopID versus AL_PA

The *Arbitrated Loop Physical Address* (AL_PA) is used for addressing within an arbitrated loop. This hexadecimal value (10 bit) is determined dynamically during initialization of the loop. The equivalent to the AL_PA is the decimal LoopID (8 bit). The table below shows the relation between LoopID and AL_PA:

Loop ID	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
00-0F	EF	E8	E4	E2	E1	E0	DC	DA	D9	D6	D5	D4	D3	D2	D1	CE
10-1F	CD	CC	CB	CA	C9	C7	C6	C5	C3	BC	BA	B9	B6	B5	B4	B3
20-2F	B2	B1	AE	AD	AC	AB	AA	A9	A7	A6	A5	A3	9F	9E	9D	9B
30-3F	98	97	90	8F	88	84	82	81	80	7C	7A	79	76	75	74	73
40-4F	72	71	6E	6D	6C	6B	6A	69	67	66	65	63	5C	5A	59	56
50-5F	55	54	53	52	51	4E	4D	4C	4B	4A	49	47	46	45	43	3C
60-6F	3A	39	36	35	34	33	32	31	2E	2D	2C	2B	2A	29	27	26
70-7F	25	23	1F	1E	1D	1B	18	17	10	0F	08	04	02	01	00	--

Example: determine the AL_PA for **LoopID 22** (decimal):

0d22 = 0x16

LoopID can be found in 2nd row (0x10–0x1F)

7th column holds the appropriate **AL_PA: 0xC6**

The priority grows with increasing LoopID and decreasing AL_PA respectively.

So the highest priority has loop ID 126 (or 0x7E which corresponds to AL_PA 0x00). This id is reserved for a switchport (FL_Port) that may be connected to the loop. The loop ids below 126 are dynamically assigned. These so called soft addresses are used to address e.g. HBAs. Storage devices like disk arrays or tape libraries have to get fixed/hard coded loop ids because the loop id determines the hardware path and therefor the name of the device file.

FC HBAs, Software and Patches

A Host Bus Adapters (HBA) incorporates processors to perform protocol conversion and I/O operations to off-load these duties from the host CPU. HP supports HBAs with the following 3 FC processors:

Generation	Chip name	Topologies	FC speed	Fibre Optic Connector
1	Tachyon	FC-AL	256Mbps/1Gbps	GLM SC connector
2	Tachyon TL/TS	FC-AL-2, P2P	1Gbps	GBIC SC connector
4	Tachyon XL2	FC-AL-2, P2P	1 Gbps/2 Gbps	SFP LC connector

NOTES:

GLM: Gigabit Link Module

GBIC: Gigabit Interface Converter

SFP: Small Form Factor Pluggable LC connector

The Tachyon XL2 adapter is available as of Application Release Dec 2001 (DART 55).

Common features of all adapters are

- Supported FC cables:
 - 50.0/125 (50.0µm core diameter; up to a length of 500m)
 - 62.5/125 (62.5µm core diameter; up to a length of 175m)
- Non-OFC, i.e. sends light (with less power) even with no cable attached
- Follows the ANSI standards for Standard Fibre Channel (X3T11)
- Fulfills the regulations of FTZ-1046 (VDE Level B)

Features that are different are

Tachyon adapter	Tachyon TL/TS adapter	Tachyon XL2 adapter
shortwave GLM (Gigabit Link Module), interconn. for SC Duplex connector	shortwave GBIC (GigaBit Interface Converter), interconn. for SC Duplex connector	SFP (Small Form Factor Pluggable), interconn. for LC duplex conn.
PCI interface	PCI-2X interface	PCI-4X interface
supports private loops (FC-AL) only	supports private and public loops (FC-AL-2)	supports private and public loops (FC-AL-2)
Annual failure rate: 5.0% (A3740A), 3.4% (A3404A, A3636A), 2.0% (A3591B)	Annual failure rate: 1.7% (A5158A), 1.4% (A6684A, A6685A)	Annual failure rate: 1.0% (A6795A)

NOTES:

- Tachyon TL/TS/XL2 adapter are backwards compatible and come with new features
- Tachyon TL/TS/XL2 adapter utilize the CPUs 66% less compared to Tachyon, this results in higher system performance.
- The Tachyon adapter does not support "fabric mode".
- The Tachyon XL2 adapter has OLA/R support for UX 11.11 and later

HP supports the following HBAs:

Product	FC Chip	Bus Type	Systems	OS	Lab Name	Discont. Date
A3404A	Tachyon	HSC	K-Class (except K100)	10.20, 11.00, 11.11	Baby Hugo	09/2002
A3591B	Tachyon	HSC	Dx20/Dx30/Dx70/Dx80/Dx90 R3x0/A180	10.20, 11.00, 11.11	Mombasa	06/2002
A3636A	Tachyon	HSC	T600	10.20, 11.00, 11.11	Baby Jade	05/2002
A3740A	Tachyon	PCI	A/L/N/V	11.00	KnightLite	10/2001
A5158A	Tachyon TL	PCI	Ax00/L/N/V/SD/rp8400 rx4610/rx9610	11.00, 11.11, 11.20(*)	Downy	12/2002
A6684A	Tachyon TL	HSC	Dx20/Dx30, A180 (**) Dx70/Dx80/Dx90, R3x0	10.20, 11.00, 11.11	Bounce	tbd
A6685A	Tachyon TL	HSC	Kx20/x50/x60/x70/x80	10.20, 11.00, 11.11	Bounce	tbd
A6795A	Tachyon XL2	PCI	Ax00/L/N/SD/rp8400	11.00, 11.11 (64 bit only)	Doubledowny	tbd

NOTE (*): only Private Loop is supported with UX 11.20, no fabric.

NOTE ()**: these systems are supported since Application Release Dec 01 (DART 55)
 for UX 10.20: since TL driver B.10.20.03
 for UX 11.00: since TL driver B.11.00.10
 for UX 11.11: since TL driver B.11.11.09
http://wtec.cup.hp.com/~hpux/io/current_issues/articles/100506485440883.html (HP internal)

The drivers below can be found on the application CD-ROMs/DVDs or at <http://software.hp.com> (below *drivers*).

To enable the fibre channel protocol you do always need the FCMS driver as a foundation:

UX 10.20:

J3630BA.FCMS B.10.20.61 Fibre Channel Mass Storage Driver

For UX 11.X the FCMS driver is already included in Core-OS:

FCMassStorage B.11.XX Fibre Channel Mass Storage

The driver is patched by the Fibre Channel Mass Storage (FCMS) patch:

UX 10.20 [PHSS_23581](#) (or newer)
 UX 11.00 [PHKL_23939](#) (or newer)
 UX 11.11 [PHKL_23626](#) (or newer) **PHKL_34161** as of 10/23/2006)

The FCMS driver enables the Tachyon chip only. In order to operate the Tachyon TL/TS or XL2 chip you additionally need to install the corresponding driver.

UX 10.20:

A6684A.FC-Tachyon-TL B.10.20.03 HSC Tachyon TL Fibre Channel
 A6685A.FC-Tachyon-TL B.10.20.03 HSC Tachyon TL Fibre Channel

UX 11.00:

Starting with Application Release Dec 2001 (DART 55) all four TL/TS/XL2 adapters, A6684A/A6685A/A5158A/A6795A, share a common driver:

```
A6795A.FC-TACHYON-TL    B.11.00.10    PCI Tachyon TL/TS/XL2 Fibre Channel
```

The other bundles (A6684A, A6685A and A5158A) contain exactly the same product.

UX 11.11:

Starting with Application Release Dec 2001 (DART 55) all four TL/XL2 adapters, A6684A/A6685A/A5158A/A6795A, share a common driver. It is already included in Core-OS (version B.11.11.09 as of Dec01):

```
FibrChan1-00.FC-TACHYON-TL    B.11.11.09    PCI/HSC FibreChannel;Supptd
                               HW=A6684A,A6685A,A5158A,A6795A
```

The appropriate patch for the Tachyon TL/TS/XL2 driver family is:

UX 11.00 [PHSS_23996](#) (or newer) this patch is only for <= B.11.10.09
 UX 11.11 [PHSS_24121](#) (or newer [PHSS_31326](#)) this patch is only for <= B.11.11.06

Check the current support state here: http://techcom.cup.hp.com/dir_fcms/adapters.html

Fibre Channel Kernel Driver & Tunables

The following **kernel drivers** are FC related. Some are for legacy Tachyon adapters only, some are for Tachyon TL/TS/XL2 adapters only and some are valid for all.

Driver	Tachyon	Tachyon TL/TS/XL2	Explanation
fcms	✓	✓	
fcpmux	✓	✓	if a FC-SCSI mux is used
fcpararray	✓	✓	FCP Array Interface
fcpcdio	✓	✓	
fcpdev	✓	✓	FCP Device Interface
fcTl	✓		Tachyon driver
fcTl_fcp	✓		
fcTl_cntl	✓		Fibre Channel Mass Storage Cntl
fcpc	✓		FCP Protocol Adapter
td		✓	Tachyon TL/TS/XL2 driver

There are the following **kernel tunables** are FC related. Do not modify them unless you are 100% sure of what you're doing.

Tunable	Tachyon	Tachyon TL/TS/XL2	Explanation
fcplarge_config	✓	✓	determines whether additional memory should be allocated for FC ports. A value of 0 (default) will handle up to 64 FC ports. A value of 1 indicates that more the 64 ports may be handled concurrently. Only for Arbitrated Loop!
num_tachyon_adapters	✓		specifies how many Tachyon FCP adapters are installed in the system so that an appropriate amount of memory can be allocated for them at system start-up if the system does not provide I/O virtual addressing. Default of 0 lets the system decide.
max_fcp_reqs	✓		specifies the maximum number of concurrent FCP requests allowed on any Tachyon FCP installed in the machine. Default is 512.

See <http://docs.hp.com/hpux/onlinedocs/os/KCparam.FiberchannelOverview.html> for details.

To determine which version of the drivers is installed:

The FCMS driver:

```
# what /stand/vmunix | grep fcms
    $Revision: libtd.a:    vw: fcms    selectors: CUP11.11_BL2001_1
003 'AR1201-11i' -- ameen_2g_merged_11i 'r11.11' 'cup_td_2g_1111' Thu Oct 11 11
:53:47 PDT 2001 $
    fcms.c $Date: 2001/01/04 10:24:48 $Revision: r11.11/1 PATCH_11.11 (PHKL
_22874)
    fcms_cdio.c $Date: 2000/12/13 15:47:27 $Revision: r11.11/1 PATCH_11.11
(PHKL_22874)
```

and additionally for Tachyon TL/TS/XL2 adapters:

```
# what /stand/vmunix | grep td
    libtd.a HP Fibre Channel Tachyon TL/TS/XL2 Driver B.11.11.09 (AR1201) /
ux/kern/kisu/TL/src/common/wsio/td_glue.c: Oct 11 2001, 11:52:36
    $Revision: libtd.a:    vw: fcms    selectors: CUP11.11_BL2001_1
003 'AR1201-11i' -- ameen_2g_merged_11i 'r11.11' 'cup_td_2g_1111' Thu Oct 11 11
:53:47 PDT 2001 $
```

Fibre Channel Addressing in HP-UX

The HP-UX addressing scheme assigns fixed and hierarchical hardware paths to the devices, based on the SCSI-2 specification. 14 bits are divided as follows:

Bus (7 bit = 128 busses)
 Target (4 bit = 16 targets)
 Lun (3 bit = 8 luns)

The HW path to a SCSI disk would look like this:

```
[SCSI-HBA].[Bus].[Target].[Lun]
```

But Fibre Channel requires much more devices to be mapped:

- Every Loop can have up to 126 FC devices connected
- SAN address space is 2^{24} FC ports (N_Port ID is 24 bit)
- SCSI-3 standard allows up to 2^{64} Luns for each target (i.e. per FC port)

Therefore the FC address space has to be adapted to the HP-UX structures. This is done by expanding the HW path by 3 more fields of 8 bit each, adding 24 bit address space.

The HW path of a FC device comprises the following three parts:

```

          additional FC fields      standard SCSI fields
[FC-HBA].[Domain].[Area].[Port].[Bus].[Target].[Lun]

```

- **HBA (host bus adapter) part**
It depends on the slot at the host where the FC adapter card is seated.
- **Fibre Channel part (N_Port ID)**
This part comprises of the new 8 bit fields **Domain**, **Area** and **Port**. In general these fields are used to hold the **N_Port ID**. This 3 byte value can be regarded as the equivalence to the fixed 6 byte Media-Access Control identifier (MAC address) in the LAN world. In case of a fabric topology (a switch is present) the switch assigns the N_port ID dynamically during login and stores the information it's private name server table.

Domain:

The Domain ID usually identifies the instance of the switch in a fabric environment. For private loops the value is 8.

Area:

In a fabric environment, the Area ID is generally associated with a physical port on a switch. For private loops, the value is 0.

Port:

When interpreting the hardware path for FCMS devices, arrays are defined as having addressable controllers, and non-arrays are defined as not having addressable

controllers. For hardware paths associated with non-array LUNs, the Port ID is set to the value 255 (255 for direct connect controller). For hardware paths associated with array LUNs, the Port ID is set to the Loop ID, for a private loop topology, AL_PA for public loop topology or 0 for direct fabric attach.

Domain and Area field hold the first two bytes of the N_Port ID. The third byte either comes from the Port field or from Bus and Target field of the SCSI part, depending on the addressing mode. The examples some pages ahead will help to understand this.

Topology	Fibre Channel Part		
	Domain_ID (8 bit)	Area_ID (8 bit)	Port_ID (8 bit)
Private Loop	Protocol Adapter defined in FC layer 4 (FC_4) 8 for Mass Storage 5 for Networking	always 0 for private loop	either LoopID or 255 in PDA addressing mode (see below)
Public Loop	usually identifies the Domain ID (i.e. instance) of the switch (at least for Brocade). 5 and 8 are reserved for private loop.	the physical connector (slot) on the switch. For Brocade subtract 16 from the Area ID to get the slot number	either AL_PA or 255 in PDA addressing mode (see below)
Direct Fabric Attach	<i>same as public loop</i>	<i>same as public loop</i>	either 0 or 255 in PDA addressing mode (see below)

NOTE:

For the Brocade switches you can obtain the Domain ID of the switch either in the config menu of the front panel or when logged on via telnet using the `switchshow` command (user: admin, password: password).

- **(virtual) SCSI part**
This part reflects the connected devices (DLTs, LUNs etc.) in a manner that HP-UX is able to access them like SCSI. The SCSI part consists of the fields Bus, Target and Lun.

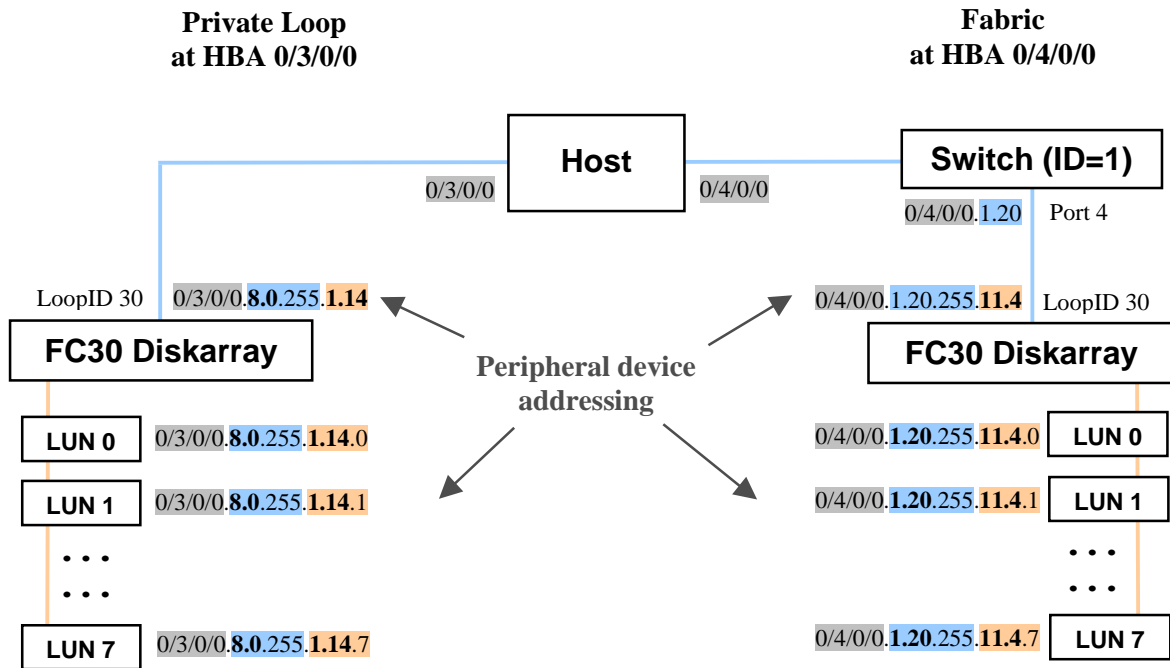
This **grey blue orange** color scheme will be used throughout the whole chapter in order to identify HBA, FC- and SCSI part.

HP-UX differentiates between three **addressing modes**:

- Peripheral Device Addressing (PDA)
- Volume Set Addressing (VSA)
- Logical Unit Addressing (LUN)

The driver sends out an inquiry to get the type of the device and it's capabilities. Which addressing mode to use when depends on the topology and the device.

Peripheral Device Addressing (PDA)



This addressing mode can be identified by the fact that the Loop ID (for private loop) or AL_PA (for public loop) is coded in the fields `bus` and `target`, whereas the SCSI devices are coded in the `LUN` field. The upper 4 bits of the 8 bit Loop ID/AL_PA go to the bus field, the lower 4 bit go to the target field. So given the HW path you can calculate the Loop ID/AL_PA: $16 * \text{Bus} + \text{Target}$. Only the 3 bit of the Lun field are left for addressing Luns, resulting in $2^3 = 8$ possible Luns.


There is no translation at the devices behind the FC controller. This is why the `port` field is always set to 255 (i.e. direct attach). So the 255 in the port field is an indicator for PDA addressing.

PDA is the standard addressing scheme. It is used whenever there's no need to exceed 8 Luns.

Example private loop:

Loop ID = 30 = 0x1E \Rightarrow bus=0x1 = 1, target=0xE = 14

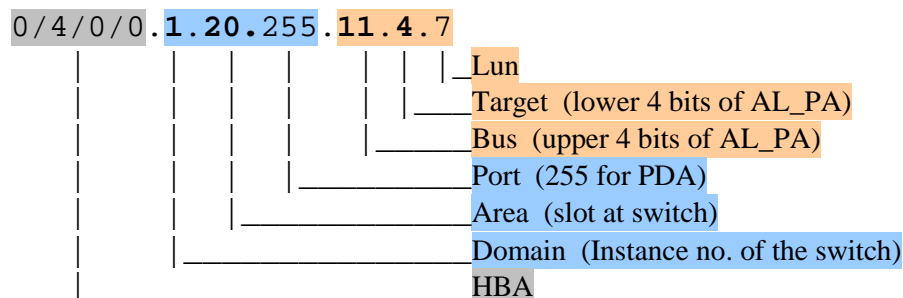
This results in the following HW path:



Example public loop:

Loop ID = 30 = 0x1E ==> AL_PA = 0xB4 => bus=0xB = 11, target=0x4 = 4

This results in the following HW path:



The N_Port ID is coded in the Domain field, the Area field and the AL_PA:

Domain= 1 = 0x01, Area = 20 = 0x14, AL_PA = 0xB4 ==> N_Port ID=0x0114B4

PDA addressed devices are:

- FC controller at XP, EMC, VA, FC60 disk arrays and FC-SCSI Mux
- FC4/2 bridge, FC 2/1 bridge
- Galactica DLT libraries
- FC10 disk devices
- Luns in Galaxy diskarray and Hitachi Data System (HDS) diskarray

NOTE:

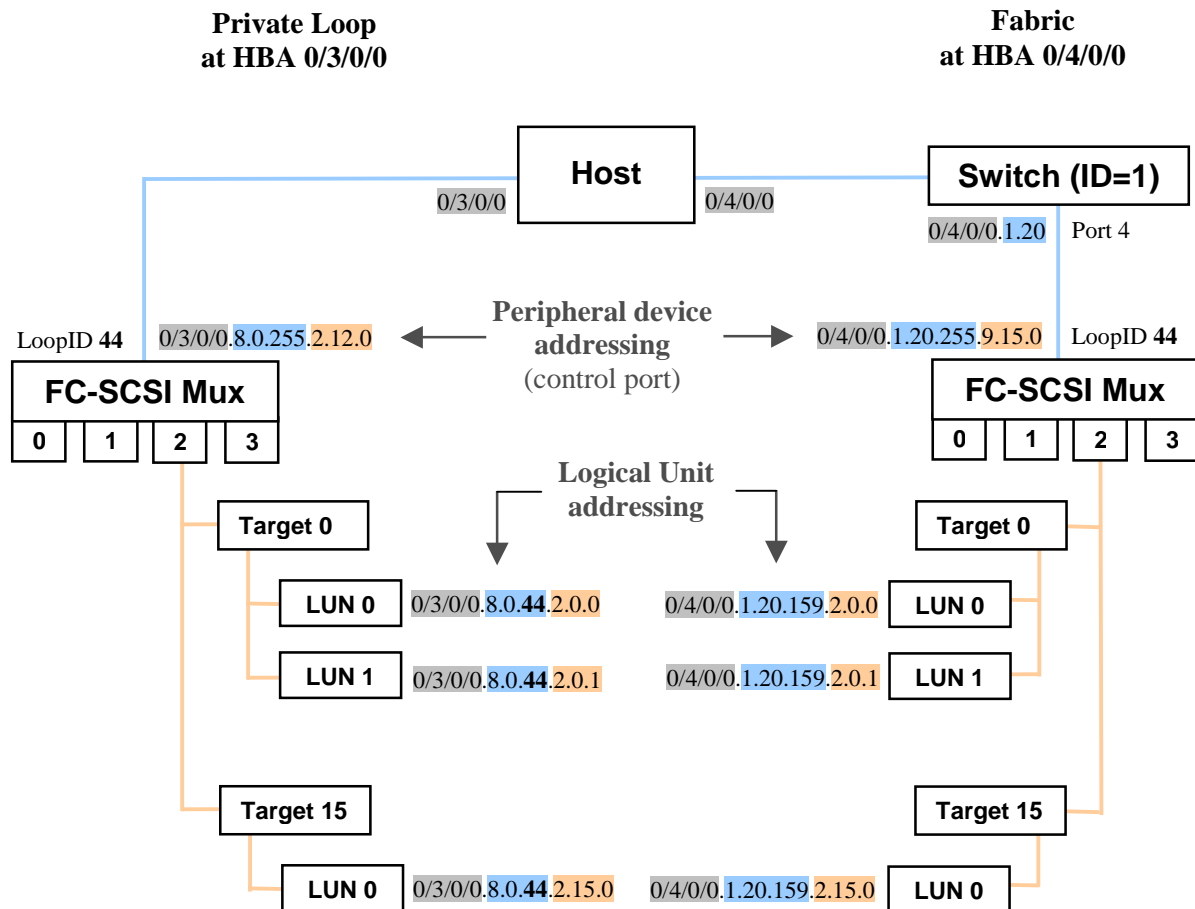
Galaxy has no addressable controller like XP- or EMC-diskarray

NOTE:

although the HDS is identical to HP's XP it cannot be addressed using VSA.

See http://wtec.cup.hp.com/~hpux/io/current_issues/articles/9962369007867.html (HP internal)

Logical Unit Addressing (LUN)



With LUN addressing mode the Loop-ID/AL_PA can be found in the Port field. Bus, Target and Lun fields are used to address SCSI devices as usual.

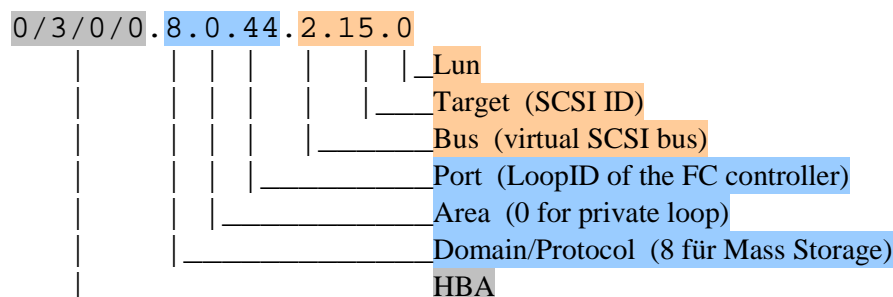
This addressing mode is used because the bridge does a conversion from FC to SCSI and therefore the SCSI devices are more structured. The `port` field holds at the SCSI site the Loop-ID/AL_PA of the previous component (FC-port of the FC-SCSI bridge). The `bus` field holds the instance of the SCSI bus whereas the fields `target` (SCSI ID) and `LUN` (LUN no.) are addressed like in the SCSI world.

This addressing mode is required whenever more than 16 targets exist. The driver decides on SCSI inquiry whether LUN addressing will be used.

Example private loop:

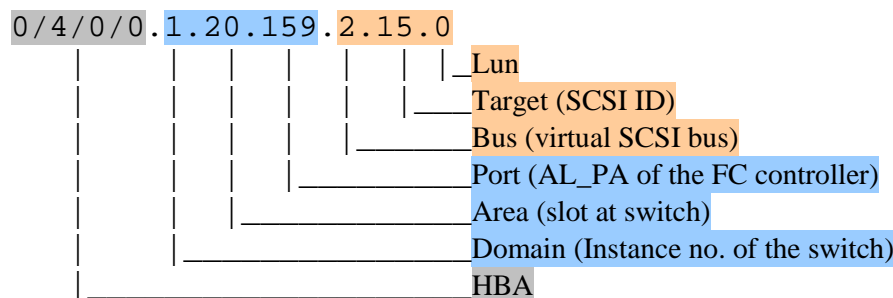
Loop ID of the FC controller of the bridge = **44**

This results in the following HW path:

**Example public loop:**

Loop ID of the FC controller of the bridge = 44 = 0x2C ==> AL_PA = 0x9F = **159**

This results in the following HW path:



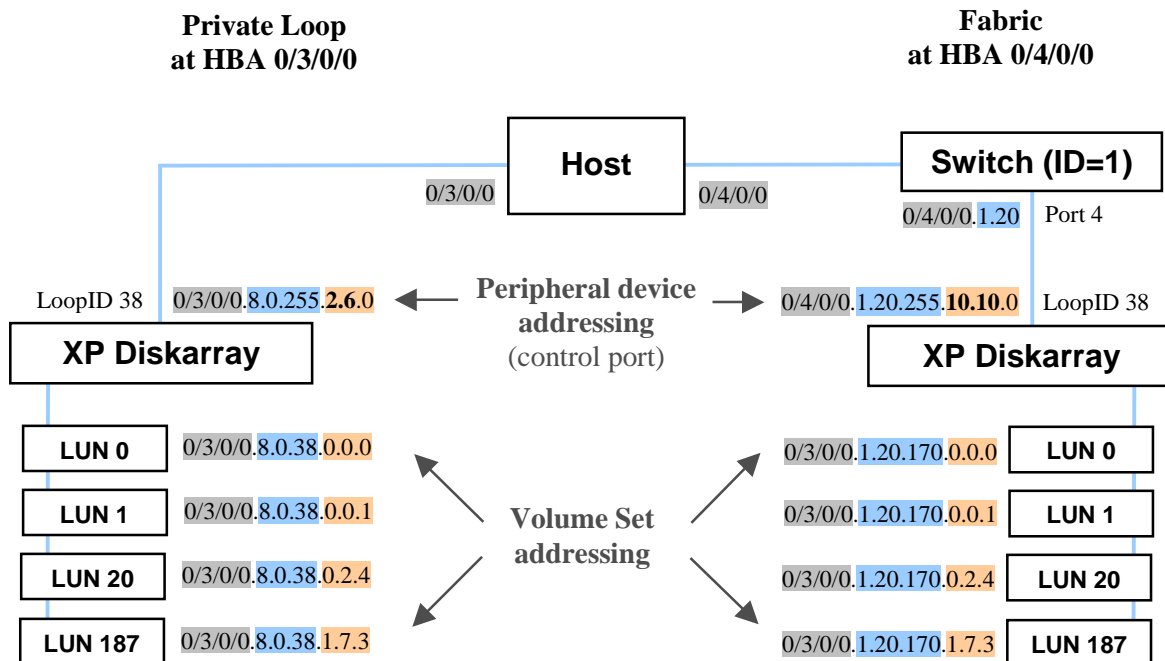
The N_Port ID is coded in the Domain field, the Area field and the AL_PA:

Domain= 1 = 0x01, Area = 20 = 0x14, AL_PA = 0x9F ==> **N_Port ID=0x01149F**

Examples for LUN addressed devices:

- SCSI interfaces of the FC-SCSI Mux and the devices behind this bus
- HP AutoRAID 12H
- XP diskarray in private loop

Volume Set Addressing (VSA)



This addressing mode was introduced to overcome the limitation of 8 Luns for disk arrays in [PDA addressing](#), where each Volume is mapped to a SCSI Lun of a single SCSI target. The whole 14 bits of Bus, Target and Lun field are now used for the addressing of the Luns, resulting in a maximum of $2^{14} = 16384$ addressable Luns.

Like in LUN addressing the `port` field holds the Loop-ID/AL_PA of the previous component on the FC site (FC-Port of the disk array). The fields `bus`, `target` und `Lun` are used for the mapping of the luns, that are configured on the disk array. The driver decides on SCSI inquiry wether VSA addressing will be used.

Example:

The Lun number (or Volume ID) generated on an XP disk array is **187**.

$$\begin{aligned} \text{Vol ID} = 187 &= (\text{Bus Target Lun}) \\ &= (0000001 \ 0111 \ 011)_2 \\ \implies \text{Bus} &= (0000001)_2 = 1, \text{ Target} = (0111)_2 = 7 \text{ und Lun} = (011)_2 = 3 \end{aligned}$$

typically you calculate like this:


$$\begin{aligned} \text{Bus} &= \text{Vol ID DIV } 2^7 \\ \text{Target} &= (\text{Vol ID MOD } 2^7) \text{ DIV } 2^3 \\ \text{Lun} &= (\text{Vol ID MOD } 2^7) \text{ MOD } 2^3 \end{aligned}$$

and vice versa, starting with bus and target: and lun:

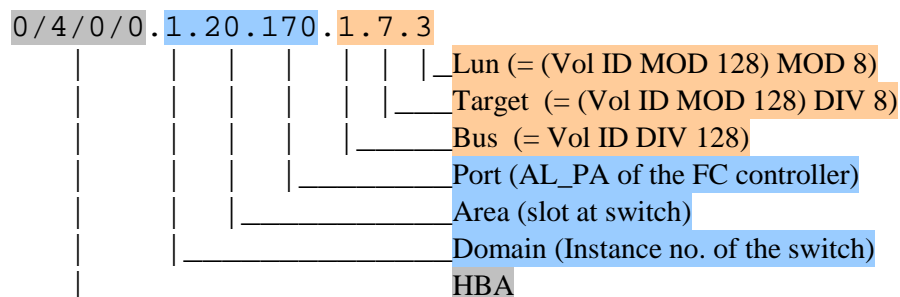
$$\text{Vol ID} = 128 * \text{Bus} + 8 * \text{Target} + \text{Lun}$$

Example private loop:

Loop ID of the FC controller of the XP = 38.

**Example public loop:**

Loop ID of the FC controller of the XP = 38 = 0x26 ==> AL_PA = 0xAA = 170.



The N_Port ID is coded in the Domain field, the Area field and the AL_PA:

Domain= 1 = 0x01, Area = 20 = 0x14, AL_PA = 0xAA ==> **N_Port ID=0x0114AA**

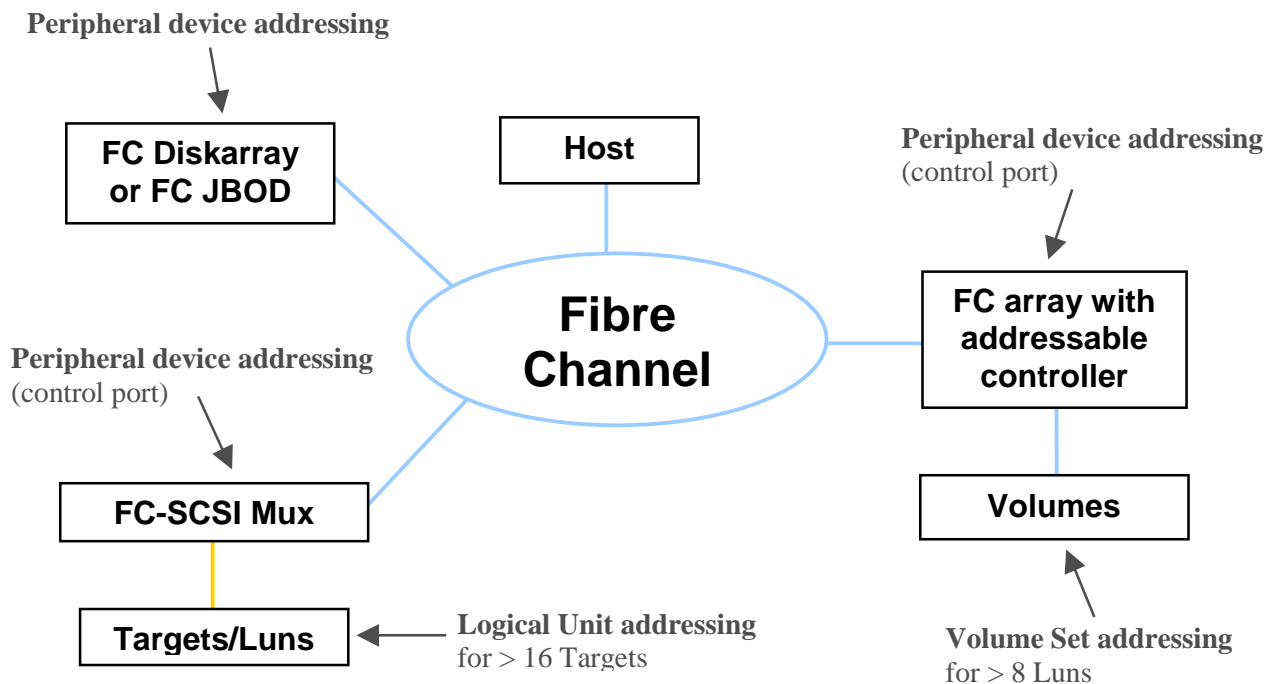
According to PDA addressing the AL_PA is coded in the bus and target field of the HW path of the XP's control port:

AL_PA = 0xAA ==> target = 0xA = 10, bus = 0xA = 10 ==> 0/4/0/0.1.20.255.10.10.0.

VSA addressed devices are:

- XP-, VA-, FC60- and EMC Diskarray Volumes

The following picture summarizes the addressing modes explained above:



some basic rules to remember resulting from the above:

- If the **Domain ID is 8** the Area ID will be 0 and we are in a private loop, i.e. without a switch.
If the **Domain ID is not 0** it holds the ID of the attached switch and the Area ID will hold the switchport (minus 16 for Brocade) where the device is connected. The device is operating in a Public-Loop or in direct fabric attach mode (Port ID = 0)
- If the **Port ID is 255** the device is directly attached to the host over a hub or a switch and operates in PDA mode.
If the **Port ID is 0** the device operates in **direct fabric attach mode** (if the device supports that) otherwise the Port ID holds the Loop ID for private loop or AL_PA for public loop.
- switch means fabric, hub means (private) loop


The utility “SAN Toolbox” is able to do all the above conversions. Its a Windows executable:
http://hprnt06.grc.hp.com/central_san_cc/san_toolbox/index.htm (HP internal)

Example: Private Loop with FC-SCSI Mux

Tachyon HBA --- HUB --- FC-SCSI Mux --- SCSI Disk

The following example explains PDA and LUN addressing in a private loop. A FC-SCSI Mux with 2 FC adapter and 2 SCSI cards has one disk connected.

- 1) The FC-SCSI Mux has a unique **Loop ID** in the range 0 - 125 assigned for each of its FC adapter.
In this example: FC Adapter A = 45, FC Adapter B = 63
- 2) It needs to be clear which of the 4 **SCSI Slots** (0 - 3) of the FCMS is equipped with a SCSI controller card. In this example: SCSI Slot 0, SCSI Slot 1
The connected disk is attached to the SCSI controller in Slot 0
- 3) The **SCSI-ID** of the device needs to be note. For disk arrays you would also need the Lun number.
In this example: SCSI-ID of the SEAGATE disk is 10)



How is the `ioscan` output for the above setup?

```
# ioscan -fn
...
...
ba          1  4          epic          CLAIMED      BUS_NEXUS     PCI Bus Bridge -
epic
fc          1  4/1/0      fcT1          CLAIMED      INTERFACE     HP Fibre Channel
Mass Storage Adapter
lan         2  4/1/0.5      fcT1_cntl    CLAIMED      INTERFACE     HP Fibre Channel
Mass Storage Cntl
/dev/fcms2
fcp         1  4/1/0.8      fcp          CLAIMED      INTERFACE     FCP Protocol
Adapter

(c) ext_bus   6  4/1/0.8.0.45.0    fcpmux       CLAIMED      INTERFACE     HP A3308 FCP-
SCSI MUX Interface
target     4  4/1/0.8.0.45.0.7    tgt          CLAIMED      DEVICE
ctl        5  4/1/0.8.0.45.0.7.0  sctl         CLAIMED      DEVICE        Initiator
/dev/rscsi/c6t7d0
target     5  4/1/0.8.0.45.0.10   tgt          CLAIMED      DEVICE
disk       7  4/1/0.8.0.45.0.10.0 sdisk        CLAIMED      DEVICE        SEAGATE ST32550W
```

				/dev/dsk/c6t10d0	/dev/rdisk/c6t10d0		
(d)	ext_bus	7	4/1/0.8.0.45.1	fcpmux	CLAIMED	INTERFACE	HP A3308 FCP-
	SCSI MUX Interface						
	target	6	4/1/0.8.0.45.1.7	tgt	CLAIMED	DEVICE	
	ctl	6	4/1/0.8.0.45.1.7.0	sctl	CLAIMED	DEVICE	Initiator
				/dev/rscsi/c7t7d0			
(c)*	ext_bus	9	4/1/0.8.0.63.0	fcpmux	CLAIMED	INTERFACE	HP A3308 FCP-
	SCSI MUX Interface						
	target	11	4/1/0.8.0.63.0.7	tgt	CLAIMED	DEVICE	
	ctl	8	4/1/0.8.0.63.0.7.0	sctl	CLAIMED	DEVICE	Initiator
				/dev/rscsi/c9t7d0			
	target	12	4/1/0.8.0.63.0.10	tgt	CLAIMED	DEVICE	
	disk	8	4/1/0.8.0.63.0.10.0	sdisk	CLAIMED	DEVICE	SEAGATE ST32550W
				/dev/dsk/c9t10d0	/dev/rdisk/c9t10d0		
(d)*	ext_bus	10	4/1/0.8.0.63.1	fcpmux	CLAIMED	INTERFACE	HP A3308 FCP-
	SCSI MUX Interface						
	target	13	4/1/0.8.0.63.1.7	tgt	CLAIMED	DEVICE	
	ctl	9	4/1/0.8.0.63.1.7.0	sctl	CLAIMED	DEVICE	Initiator
				/dev/rscsi/c10t7d0			
(a)	ext_bus	8	4/1/0.8.0.255.2	fcpdev	CLAIMED	INTERFACE	FCP Device
	Interface						
	target	7	4/1/0.8.0.255.2.13	tgt	CLAIMED	DEVICE	
	ctl	7	4/1/0.8.0.255.2.13.0	sctl	CLAIMED	DEVICE	HP HPA3308
				/dev/rscsi/c8t13d0			
(b)	ext_bus	11	4/1/0.8.0.255.3	fcpdev	CLAIMED	INTERFACE	FCP Device
	Interface						
	target	14	4/1/0.8.0.255.3.15	tgt	CLAIMED	DEVICE	
	ctl	10	4/1/0.8.0.255.3.15.0	sctl	CLAIMED	DEVICE	HP HPA3308

(* means alternative link to the same device)

Interpretation:

Host interface (HBA):

```
fc 1 4/1/0 fcT1 CLAIMED INTERFACE HP Fibre Channel Mass Storage Adapter
```

FC protocol adapter is 8 for mass storage:

```
fcpl 1 4/1/0.8 fcpl CLAIMED INTERFACE FCP Protocol Adapter
```

The **FC adapter** (a) and (b) of the FC-SCSI Mux are addressed in PDA mode with their Loop-IDs. In ioscan they turn up behind the devices connected to the FC-SCSI Mux:

(a)	ext_bus	8	4/1/0.8.0.255.2	fcpdev	CLAIMED	INTERFACE	FCP Device
	Interface						
	target	7	4/1/0.8.0.255.2.13	tgt	CLAIMED	DEVICE	
	ctl	7	4/1/0.8.0.255.2.13.0	sctl	CLAIMED	DEVICE	HP HPA3308
				/dev/rscsi/c8t13d0			
(b)	ext_bus	11	4/1/0.8.0.255.3	fcpdev	CLAIMED	INTERFACE	FCP Device
	Interface						
	target	14	4/1/0.8.0.255.3.15	tgt	CLAIMED	DEVICE	
	ctl	10	4/1/0.8.0.255.3.15.0	sctl	CLAIMED	DEVICE	HP HPA3308

Bus and target of the FC-adapter are derived from the Loop-ID according to PDA addressing:

for FC-A (Loop-ID 45): Bus = 45 DIV 16 = 2 Target = 45 MOD 16 = 13	for FC-B (Loop ID 63): Bus = 63 DIV 16 = 3 Target = 63 MOD 16 = 15
--	--

The **SCSI controller** of the FC-SCSI mux and the Seagate disks behind are accessed using LUN addressing mode:

```
# ioscan -fd fcpmux
Class      I  H/W Path          Driver S/W State   H/W Type   Description
=====
ext_bus    6  4/1/0.8.0.45.0   fcpmux CLAIMED    INTERFACE  HP A3308 FCP-SCSI MUX
Interface
ext_bus    7  4/1/0.8.0.45.1   fcpmux CLAIMED    INTERFACE  HP A3308 FCP-SCSI MUX
Interface
ext_bus    9  4/1/0.8.0.63.0   fcpmux CLAIMED    INTERFACE  HP A3308 FCP-SCSI MUX
Interface
ext_bus   10  4/1/0.8.0.63.1   fcpmux CLAIMED    INTERFACE  HP A3308 FCP-SCSI MUX
Interface
```

For each FC adapter there is an entry for each SCSI card (here: 0 and 1). This lets you know the number of SCSI cards in the FC-SCSI mux. The SCSI-ID of this cards is shown in the target field and is usually 7:

```
ext_bus      6  4/1/0.8.0.45.0   fcpmux    CLAIMED    INTERFACE    HP A3308
FCP-SCSI MUX Interface
target       4  4/1/0.8.0.45.0.7  tgt        CLAIMED    DEVICE
ctl          5  4/1/0.8.0.45.0.7.0  sctl       CLAIMED    DEVICE      Initiator
                                           /dev/rscsi/c7t7d0
```

The connected **Seagate disk** is shown straight behind this:

```
target       5  4/1/0.8.0.45.0.10  tgt        CLAIMED    DEVICE
disk         7  4/1/0.8.0.45.0.10.0  sdisk      CLAIMED    DEVICE      SEAGATE
ST32550W
                                           /dev/dsk/c6t10d0   /dev/rdisk/c6t10d0
```

The disk devicefile is composed of the Instancenummer of the FC-SCSI Mux SCSI cards (6), the SCSI-ID of the disk (10) and the Lun-ID (0) zusammen:

```
/dev/dsk/c6t10d0
```

In this example the disk can be accessed over a second FC path:

```
target       12  4/1/0.8.0.63.0.10  tgt        CLAIMED    DEVICE
disk         8  4/1/0.8.0.63.0.10.0  sdisk      CLAIMED    DEVICE      SEAGATE
ST32550W
                                           /dev/dsk/c9t10d0   /dev/rdisk/c9t10d0
```

Example: Private Loop versus Public Loop

Following is a simple example of migrating an HP XP256 from a private loop configuration to a fabric topology. A hub has been replaced with a switch resulting in different ioscan outputs.

Tachyon TL/TS/XL2 HBA --- HUB --- XP256

The ioscan output for a **Private Loop** configuration could be as follows:

```
Class      I  H/W Path          Driver   S/W State   H/W Type   Description
-----
```

```

fc          0      0/1/2/0          td          CLAIMED  INTERFACE  HP Tachyon
TL/TS Fibre Channel Mass Storage Adapter
fcp        1      0/1/2/0.8          fcp          CLAIMED  INTERFACE  FCP Protocol
Adapter
ext_bus    4      0/1/2/0.8.0.11.0    fcparray    CLAIMED  INTERFACE  FCP Array
Interface
target     6      0/1/2/0.8.0.11.0.0  tgt          CLAIMED  DEVICE
disk       3      0/1/2/0.8.0.11.0.0  sdisk        CLAIMED  DEVICE      HP OPEN-8
/dev/dsk/c3t0d0 /dev/rdisk/c3t0d0
disk       10     0/1/2/0.8.0.11.0.0.7 sdisk        CLAIMED  DEVICE      HP OPEN-8
/dev/dsk/c3t0d7 /dev/rdisk/c3t0d7
target     7      0/1/2/0.8.0.11.0.1  tgt          CLAIMED  DEVICE
disk       11     0/1/2/0.8.0.11.0.1.0 sdisk        CLAIMED  DEVICE      HP OPEN-9
/dev/dsk/c3t1d0 /dev/rdisk/c3t1d0
disk       18     0/1/2/0.8.0.11.0.1.7 sdisk        CLAIMED  DEVICE      HP OPEN-9
/dev/dsk/c3t1d7 /dev/rdisk/c3t1d7

```

Tachyon TL/TS/XL2 HBA --- SWITCH --- XP256

The ioscan output for a **Public Fabric Loop** configuration could be as follows:

```

Class      I      H/W Path          Driver      S/W State  H/W Type  Description
-----
fc          0      0/1/2/0          td          CLAIMED    INTERFACE  HP Tachyon
TL/TS Fibre Channel Mass Storage Adapter
fcp        1      0/1/2/0.1          fcp          CLAIMED    INTERFACE  FCP Domain
ext_bus    4      0/1/2/0.1.19.0.0  fcparray    CLAIMED    INTERFACE  FCP Array
Interface
target     6      0/1/2/0.1.19.0.0.0  tgt          CLAIMED    DEVICE
disk       3      0/1/2/0.1.19.0.0.0  sdisk        CLAIMED    DEVICE      HP OPEN-8
/dev/dsk/c4t0d0 /dev/rdisk/c4t0d0
disk       10     0/1/2/0.1.19.0.0.0.7 sdisk        CLAIMED    DEVICE      HP OPEN-8
/dev/dsk/c4t0d7 /dev/rdisk/c4t0d7
target     7      0/1/2/0.1.19.0.0.1  tgt          CLAIMED    DEVICE
disk       11     0/1/2/0.1.19.0.0.1.0 sdisk        CLAIMED    DEVICE      HP OPEN-9
/dev/dsk/c4t1d0 /dev/rdisk/c4t1d0
disk       18     0/1/2/0.1.19.0.0.1.7 sdisk        CLAIMED    DEVICE      HP OPEN-9
/dev/dsk/c4t1d7 /dev/rdisk/c4t1d7
fcp        0      0/1/2/0.8          fcp          SCAN       INTERFACE  FCP Protocol
Adapter

```

Looking at the iotree examples, you can see the following:

There has been no change to the adapter path or the associated device file which is used for the fcmsutil diagnostic tool.

- The node 0/1/2/0.8, FCP Protocol Adapter, is in both ioscan outputs. In a private loop configuration, the interface and target devices will reside behind this node. In a Fabric environment, this node may be created as a dummy node if the HBA is scanned when it cannot see the Fabric (for example, no cable attached, switch down, etc.).

In the original private loop implementation of the fibre channel driver, this node of the iotree was used to indicate the fibre channel FC4 “TYPE”. A type of “8” denotes that the FCP protocol is being used to encapsulate the SCSI protocol. With the introduction of fabric, this node contains the “Domain” portion of the N_Port address. To maintain backward compatibility, the domain of 8 is reserved for use with private loop devices.

CAUTION

Do not configure switches with a Domain of 8. This is an unsupported configuration and will not work. The Domain of 8 is reserved for Private Loop devices on HP systems.

- The Fabric configuration now contains an iotree node of `0/1/2/0.1` described as FCP Domain. A node of this type will be built for each Domain the Fabric contains (Domains usually correspond one to one with a switch instance).
- The FCP Array Interface iotree node has changed from `0/1/2/0.8.0.11.0` to `0/1/2/0.1.19.0.0`. The address is still at hardware Path `0/1/2/0`, but the next three elements of the path, which represent the N_Port address, have changed. The old N_Port address of `8.0.11` uses the reserved Domain of 8 and area of 0. In this case, the HPA or Port byte of the N_Port address is 11. In the Fabric iotree, the new N_Port address is `1.19.0`. This corresponds to a Domain id of 1, an area id of 19 and a port id of 0.
Note that for most switches, the Domain will map to a switch instance, an area id will map to a physical connector on the switch, and the port id will only be used (non zero) if there is an Arbitrated Loop configured behind the switch connector. The HPA of the device is then used as the Port portion of the iotree address.
- All targets and disk devices retain their original iotree addresses with the exception that the new Fabric N_Port address has been substituted for the old Arbitrated Loop address.
- New device files have been generated for the new iotree nodes. The old device files will continue to exist until removed with the `rmsf(1M)` command.

Troubleshooting Utilities fcmsutil, tduutil, tdlist, tddiag

fcmsutil

The fcmsutil utility can be found below /opt/fc/bin/ and/or /opt/fcms/bin. It helps you to troubleshoot the FC Loop or SAN. It is invoked using the device file of the FC adapter (see ioscan -fnk).

Tachyon Example:

```
# fcmsutil /dev/fcms2

Local N_Port_ID is = 0x000001
N_Port Node World Wide Name = 0x10000060B03EF669
N_Port Port World Wide Name = 0x10000060B03EF669
Topology = IN_LOOP
Speed = 1062500000 (bps)
HPA of card = 0xFFB4C000
EIM of card = 0xFFFA2009
Driver state = READY
Number of EDB's in use = 0
Number of OIB's in use = 0
Number of Active Outbound Exchanges = 1
Number of Active Login Sessions = 3
```

Tachyon TL/TS/XL2 Example:

```
# fcmsutil /dev/td0

Vendor ID is = 0x00103c
Device ID is = 0x001028
TL Chip Revision No is = 2.3
PCI Sub-system Vendor ID is = 0x00103c
PCI Sub-system ID is = 0x000006
Topology = PRIVATE_LOOP
Local N_Port_id is = 0x000001
Local Loop_id is = 125
N_Port Node World Wide Name = 0x50060b0000010449
N_Port Port World Wide Name = 0x50060b0000010448
Driver state = ONLINE
Hardware Path is = 0/3/0/0
Number of Assisted IOs = 47983
Number of Active Login Sessions = 0
Dino Present on Card = NO
Maximum Frame Size = 960
Driver Version = @(#) libtd.a HP Fibre Channel
Tachyon TL/TS/XL2 Driver B.11.11.09 (AR1201) /ux/kern/ki
su/TL/src/common/wsio/td_glue.c: Oct 11 2001, 11:52:36
```

After pulling the FC cable the Driver state will change to:

```
Driver state = AWAITING_LINK_UP
```

In order to get a summary of the link statistics:

```
# fcmsutil /dev/td0 stat -s
Fri Apr 26 16:05:55 2002
Channel Statistics
```

```

Statistics From Link Status Registers ...
Loss of signal                2      Bad Rx Char                182
Loss of Sync                  40     Link Fail                  4
Received EOFa                  0     Discarded Frame            0
Bad CRC                        0     Protocol Error              0

```

Do not look for high values. Only values that are increasing over time indicate a problem.

Here's how to determine the different topologies using `fcmsutil`:

loop	fabric	topology	fcmsutil output will be
yes	no	Private (Arbitrated) Loop	PRIVATE_LOOP/IN_LOOP
yes	yes	Public Loop	PUBLIC_LOOP/IN_LOOP_FL
no	yes	Switched Point-To-Point	IN_PTTOPT_FABRIC

In the following we see a typical error message:

```
0/4/0/0: Unable to access previously accessed device at nport ID 0xae.
```

Here's how to troubleshoot:

```

# ioscan -fnkH0/4/0/0
fc          0  0/4/0/0          td          CLAIMED     INTERFACE    HP Tachyon TL/TS
Fibre Channel Mass Storage Adapter

fc          0  0/4/0/0.8        fcp          CLAIMED     INTERFACE    FCP Protocol Adapter
ext_bus    4  0/4/0/0.8.0.255.0  fcpdev      CLAIMED     INTERFACE    FCP Device Interface
target     8  0/4/0/0.8.0.255.0.12  tgt         CLAIMED     DEVICE
disk       3  0/4/0/0.8.0.255.0.12.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target     9  0/4/0/0.8.0.255.0.13  tgt         CLAIMED     DEVICE
disk       4  0/4/0/0.8.0.255.0.13.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target    10  0/4/0/0.8.0.255.0.14  tgt         CLAIMED     DEVICE
disk       5  0/4/0/0.8.0.255.0.14.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target    11  0/4/0/0.8.0.255.0.15  tgt         CLAIMED     DEVICE
disk       6  0/4/0/0.8.0.255.0.15.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
ext_bus    5  0/4/0/0.8.0.255.1   fcpdev      CLAIMED     INTERFACE    FCP Device Interface
target    12  0/4/0/0.8.0.255.1.6   tgt         CLAIMED     DEVICE
ctl        4  0/4/0/0.8.0.255.1.6.0  sctl        CLAIMED     DEVICE        HP          A5236A
target    13  0/4/0/0.8.0.255.1.8   tgt         CLAIMED     DEVICE
disk       7  0/4/0/0.8.0.255.1.8.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target    14  0/4/0/0.8.0.255.1.9   tgt         CLAIMED     DEVICE
disk       8  0/4/0/0.8.0.255.1.9.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target    15  0/4/0/0.8.0.255.1.10  tgt         CLAIMED     DEVICE
disk       9  0/4/0/0.8.0.255.1.10.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
target    16  0/4/0/0.8.0.255.1.11  tgt         CLAIMED     DEVICE
disk      10  0/4/0/0.8.0.255.1.11.0  sdisk       CLAIMED     DEVICE        SEAGATE ST39102FC
ext_bus    6  0/4/0/0.8.0.255.2   fcpdev      NO_HW      INTERFACE    FCP Device Interface

```

N_port ID (= AL_PA, because it is a private loop) = 0xae
regarding the conversion table this corresponds to LoopID 34.

```

# fcmsutil /dev/td0 devstat all | grep -e Nport -e Failed
Device Statistics for Nport_id 0x0000ae(Loop_id 34)

```

```

Failed Open of previously opened device          9
Device Statistics for Nport_id 0x0000b9(Loop_id 27)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000ba(Loop_id 26)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000bc(Loop_id 25)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000c3(Loop_id 24)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000c6(Loop_id 22)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000ce(Loop_id 15)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000d1(Loop_id 14)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000d2(Loop_id 13)
Failed Open of previously opened device          0
Device Statistics for Nport_id 0x0000d3(Loop_id 12)
Failed Open of previously opened device          0

```

This is a private loop with PDA Addressing (8.0.255), i.e

```

LoopID = 16*Bus+Target    ==> Bus    = 34 DIV 16 = 2
                        ==> Target = 34 MOD 16 = 2

```

this results in the following HW path:

```

HBA   Domain Area Port Bus Target Lun
0/4/0/0 . 8 . 0 . 255 . 2 . 2 . 0

```

This path is not part of ioscan because the interface above this path is shown as NO_HW, i.e someone disconnected it without rebooting:

```

ext_bus 6 0/4/0/0.8.0.255.2 fcpdev NO_HW INTERFACE FCP Device Interface

```

A5236A is a FC10 JBOD.

Overview of fcmsutil options

fcmsutil Option	Explanation	T/TL
fcmsutil <dev file>	State of the link	T+TL
fcmsutil <dev file> lgninfo_all	Info about devices in the loop	T
fcmsutil <dev file> devstat all	Info about devices in the loop	TL
fcmsutil <dev file> reset	Reset the FC card	T+TL
fcmsutil <dev file> lb tachyon	Internal loopback test, tests the Tachyon chip	T
fcmsutil <dev file> lb plm	External loopback test (needs loopback cable)	T+TL
fcmsutil <dev file> enable	Enable the card (after a HW failure)	T+TL
fcmsutil <dev file> disable	Disable the card	T+TL
fcmsutil <dev file> stat	Obtain statistics maintained by the driver	T+TL
fcmsutil <dev file> stat -s	Obtain statistics summary maintained by the driver	TL
fcmsutil <dev file> clear_stat	Reset the statistics	TL
fcmsutil <dev file> echo <nport_id>	Send packet to other N-port ID (within priv. loop)	T+TL
fcmsutil <dev file> rls <nport_id>	Send packet to another N-port ID (across switch)	TL
fcmsutil <dev file> replace_dsk <nport_id>	Replace disk (disallow authentication)	TL

T means: option available with Tachyon

TL means: option available with Tachyon TL/TS/XL2.

red font means: destructive task, i.e. the communication gets interrupted.

tdutil, tdlis, tddiag

In the directory `/opt/fcms/bin` you can find additional utilities that might be helpful:

- | | |
|---------------------|--|
| <code>tdutil</code> | this is fcmsutil but for td driver only. |
| <code>tdlist</code> | this is a shell script that uses ioscan and tdutil in order to list all devices that are handled by the td driver, i.e. all devices attached to the Tachyon TL/TS/XL2 adapters on the system. The script contains a nice function that translates Loop ID to AL_PA. |
| <code>tddiag</code> | this is a shell script that gathers the following fibre channel related information of the system:
system name, system model, system uptime, memory information, mounted file systems, TachLite version in kernel, system file, patches installed, device special files for TachLite, ioscan output, tdlis output, running processes, infos about each <code>/dev/td#</code> :
device info, device vpd info, device topology sensing mode, chip registers, device statistics, Name Server device statistics, CT Server device statistics, all remote statistics, name server port info (from kernel), name server port info (from switch), device statistics on all targets. |

How to Replace Disks at Tachyon TL/TS/XL2 HBAs

Before a server can talk to a target it has to authenticate at the Tachyon TL/TS/XL2 HBA with its WWN (World Wide Name). This authentication (PLOGI) ensures that the system is talking to the correct device, avoiding data corruption due to user accidentally connecting another device at the same `nport_id`.

The **Adapter** holds a table where the native address (`S_ID` or `AL_PA` and WWN) of every known device of the loop is stored. This table is created upon initialization of the link or when the first communication between host and device happens.

NOTE:

This authentication applies to the devices connected to the TL/TS/XL2 adapter only, Tachyon adapters do not go through the same level of authentication. (`replace_dsk` option is supported by TL/TS/XL2 only).

The `replace_dsk` option of `fcmsutil` is necessary in order to change a device and keep the same `nport_id` but primarily, it should be used when replacing a bad disk. When this option is used, no authentication on that device is performed the next time system communicates with it thus avoiding the following error (syslog):

```
0/4/0/0: 'World-wide name' (unique identifier) for device at loop ID
```

0x5 has changed. If the device has been replaced intentionally, please use the `fcmsutil replace_dsk` command to allow the new device to be used.

In this example the disk with loop_id 5 at Tachyon TL/TS/XL2-Adapter /dev/td0 is going to be replaced:

- Identify the nport_id or loop_id of the disk being moved or replaced

(this step is optional if nport_id or loop_id of the device is known from syslog, dmesg, or other error logs)

Disconnect device (i.e remove it from the loop)

List all devices with which the TL/TS/XL2 card has successfully communicated using the devstat all option:

```
# fcmsutil /dev/td0 devstat all | grep Loop
Device Statistics for Nport_id 0x0000E8 (Loop_id 1)
Device Statistics for Nport_id 0x0000DA (Loop_id 5)
```

Using the echo option, try to login to each device using the loop_id, a message "unable to login" is returned for the disconnected device.

```
# fcmsutil /dev/td0 echo -l 1
Data came back intact
...
# fcmsutil /dev/td0 echo -l 5
Unable to login
```

- Run `replace_dsk` using `nport_id`:

```
# fcmsutil /dev/td0 replace_dsk 0x0000DA
```

in a private loop you can also use the `loop_id`:

```
# fcmsutil /dev/td0 replace_dsk -l 5
```

After running the above command, you will get a message indicating the port will not be authenticated:

```
Disk at nportid 0x0000da (Loop_id 5) will not be authenticated
```

ATTENTION: This step has to be repeated for **any** Tachyon TL/TS/XL2 adapter on the host as well as for other hosts (e.g in a ServiceGuard cluster) that access the same device!

- Replace bad disk with new disk

The new disk (loop_id=5) now can login without getting an authentication failure. The next time the system issues any read/write operations (e.g. ioscan) to the new disk, the new World-Wide Name (WWN) will be recognized and recorded to associate with nport_id 0x0000da. It is important to make sure the correct nport_id or loop_id is used with `replace_dsk` but if you accidentally type in the wrong nport_id or Loop_id, nothing catastrophic will occur. Just keep in mind that there will be no authentication done on that device during the next login. The device will be authenticated on subsequent logins.

Fibre Channel Storage Devices

This is just a short overview about existing devices. If you need details got to HP's storage website <http://www.hp.com/go/storage> or follow one of the links at section [Additional Information](#).

NIKE Model 30 FC Disk Array (legacy)

Product Name: 30 Slot Fibre Channel Disk Array
Codename: „Galaxy“
Product No.: A3661A
PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/7/psp7901.htm (HP internal)

NOTE: This is a legacy product which should not be ordered anymore.
For details go to the [NIKE Disk Array Chapter](#).

FC10 Fibre Channel Disk System (legacy)

Product Name: HP SureStore E Disk System FC10
Codename: „Transformer“
Product No.: A5236A
PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/9/psp9335.htm (HP internal)

The HP SureStore E Disk System FC10 is a 10-slot Fiber Channel JBOD Disk Enclosure that features redundant Link Control Cards (LCC), redundant power supplies and fan modules. The disk modules are native Fibre Channel with 2 FC-ALports per disk. This allows the use of alternate paths for high availability configurations.

FC60 Fibre Channel Disk Array

Product Name: High Availability Fibre Channel Disk Array Model 60/FC
Codename: „Optimus Prime“
ProductNo.: A5277A
PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/10/psp10230.htm (HP internal)

High availability disk array comprised of a controller enclosure and up to six disk enclosures. Each disk enclosure can contain up to ten disks. Disk capacities are 18.2GB, 36.4GB or 73.4GB. RAID Level 0*, 1, 5, and 0/1 are supported.

The HP Array Manager/60 cumulative Patch provides the Manager Software in order to administer the FC60. It contains the commands with manual pages and the startup script (`/sbin/init.d/hparamgr`). The commands can be found in the directory `/opt/hparray/bin/`:

```
amdsp, amcfg, amdload, amfmt, amlog, ammgr, amutil.
```

The communication with the array is done by a daemon (`/usr/sbin/hparray/hparamgrd`) that is invoked by the startup script.

NOTE: The administration of the FC60 is nearly identical to the administration of an AutoRAID.

HP Array Manager/60 Patch:

UX 10.20 [PHCO_23148](#) (or newer)
 UX 11.00 [PHCO_23149](#) (or newer)
 UX 11.11 [PHCO_23150](#) (or newer)

To administer the FC60 through SAM you need at least the following SAM patch:

UX 10.20 [PHCO_19046](#) (or newer)
 UX 11.00 [PHCO_21267](#) (or newer)
 UX 11.11 not necessary

Manual Pages:

amdsp(1M), ammgr(1M), amcfg(1M), amutil(1M), amlog(1M), AM60Srvr(1M).

XP Disk Array Family

Product Name: HP SureStore E Disk Array XP512
 Codename: „Orca”
 ProductNo.: A5950A
 PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/11/psp11235.htm (HP internal)

Product Name: HP SureStore E Disk Array XP48
 Codename: „Cuda”
 ProductNo.: A5920A
 PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/11/psp11559.htm (HP internal)

Information can be found in the [XP Disk Arrays Chapter](#).

HP Virtual Array VA7100/VA7400

The HP Virtual Array VA7100 and VA7400 are utilizing AutoRAID technology. All modules front and rear are hot pluggable.

Product Name: HP Virtual Array VA7100
 Codename: „Cassini”
 ProductNo.: VA7100
 PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/11/psp11761.htm (HP internal)

The chassis accomodates:

- 3 EIA rack mountable chassis.
- 15 low profile 3.5 inch fibre channel hard disk drives loaded from the front.
- 1 or 2 fibre channel controllers mounted in the rear. Each controller has one 1Gb fibre channel interface to the host.
- 2 power supply and fan modules mounted in the rear. The power supply and fan components are combined into a single module.

Product Name: Virtual Array VA7400
 Codename: „Cronus”
 ProductNo.: VA7400
 PSP: http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/12/psp12062.htm (HP internal)

Compared to the VA7100 this array has the following additional features:

- Two types of chassis enclosures: The main that holds the AutoRAID controllers and

up to six add on enclosures that allow more drives to be added behind the AutoRAID controllers.

- Optional 2Gb configurable fibre channel interface to the host
- Add on chassis has 2 fibre channel LCC (Link Controller Card) that allows a daisy chain type of connection to the main. Each LCC has two 1Gb fibre channel interface.

You can manage the VA either by the front panel or by installing a is a graphical + command line user interface - [HP Command View SDM](#). The administration of the VA is similar to the administration of the FC60 or the AutoRAID. The commands start with `arm*` instead of `am*` or `array*` at the FC60 or AutoRAID respectively. The table at <http://ren.nsr.hp.com/howto/array.html> (HP internal) gives you a comparison between the admin commands of the three disk arrays.

You can download SDM from an external website. Just search for “sdm ux” at <http://www.hp.com>. It can also be found (together with latest Firmware) at the HP intranet: http://tce-web.boi.hp.com/prod_port/ (HP internal).

HP SureStore Disk System 2405 (ds2405)

Product Name:	HP Surestore Disk System 2405
Codename:	„Apex“
ProductNo.:	A6250A
PSP:	http://wwwpsp.atl.hp.com/lmx_mount/supplan/psp/12/psp12704.htm (HP internal)

The HP Surestore Disk System 2405 (ds2405) supports the latest 2 Gb Fibre Channel technology. With its modular design, the ds2405 is a highly scalable and flexible storage system. In a compact 3U form factor, each enclosure holds up to 15 disk drives or over 1 TB of capacity. The ds2405 can be upgraded from a standalone disk system to add-on storage for the HP Virtual Array.

Glossary

AL_PA	ArbitratedLoop_PhysicalAddress; The address of a fibre channel node in an <i>arbitrated loop</i> . (1byte; 0x00..0x7E)
Arbitrated Loop	An interconnection scheme which supports up to 126 ports on a loop in a shared medium topology. This is HP's standard implementation of a loop Same as <i>FC-AL</i> .
Cascading	An interconnection of individual switches used to create larger Fabric configuration.
DFA	Direct Fabric Attach - the connection between an N-port and an F-port.
E-Port	Switch port to cascades switches.
EPL	Enhanced Private Loop; HP's FC-AL implementation with limitations in order to use legacy devices (FC-devices in <i>PLDA</i> environment; e.g. Galaxy Disk Array) in a SAN but access them like <i>PLDA</i> .
F-Port	Switch Port that operates according to the <i>P2P</i> protocol.
Fabric	A fibre channel interconnection method which allows multiple hosts and/or storage devices connected with a multi-port hub allowing multiple, simultaneous and concurrent data transfers. The existence of a FC switch is fundamental.
FC-AL	Abbreviation for Fibre Channel Arbitrated Loop; see <i>Arbitrated Loop</i> .
FC-AL-2	Whereas FC-AL allows private loop only, FC-AL-2 allows private and public loop. In order to run public loop the device needs to be able handle 8 bit addresses within the private loop (FC-AL) and 24bit addresses when communicating over the switch.
FC-FLA	FC Fabric Loop Attach; ANSI standard that describes the communication between <i>NL-</i> und <i>FL-Ports</i> .
Fibre	Thin filament of glass. An optical waveguide consisting of a core and a cladding which is capable of carrying information in the form of light. Fibre is also a general term used to cover all physical media types supported by Fibre Channel, such as optical fiber, twisted pair, and coaxial cable.
Fibre Channel	A high-speed, serial, bi-directional, topology independent, multi-protocol, highly scalable interconnection between computers, peripherals, and networks.
FL-Port	Switch port that operates in Loop protocol.
Frame	Smallest unit of the transmission protocol between 2 N-Ports. It consists of start sequence, header, payload (SCSI data), CRC and end sequence.
ISL	Inter Switch Link. A cable to cascade FC switches.
Loop	All nodes are attached in a ring topology (similar to FDDI or Token Ring).
Loop ID	Counterpart to AL_PA. Address of the node in <i>FC-AL</i> (126 to 0).

N-Port	Port of a storage device or host that operates according to <i>P2P</i> protocol.
NL-Port	A storage device or host that operates according to <i>FC-AL</i> protocol.
Node	Device that is connected over fibre channel (host, disk array, bridge, ...).
P2P	Point-To-Point; direct connection between two nodes.
PLDA	Private Loop Direct Attach (synonym: private loop).
Port	FC-interconnection of a node. Ususally means node.
Private Loop	A private loop is an <i>Arbitrated Loop</i> which has no attached switch port (NL-port to NL-port). In theory, private loop devices can only communicate with other devices on the same loop, but see <i>Translative Mode</i> .
Public Loop	Arbitrated Loop with a switch port (FL-Port) that allows the nodes to communicate with other nodes outside the loop.
QL	Quick Loop - a technique by which a Brocade switch can combine several fibre channel links into a single virtual private arbitrated loop.
TranslativeMode	A technique by which a switch allows private targets which only speak the private loop protocol to communicate with public initiators which are not actually on the loop. The switch creates the illusion that these private targets physically connected via private loop ports are on the fabric, by creating "phantom" fabric addresses which can be the source or destination of messages. The switch sees and captures these messages and reroutes them to the correct address.

Additional Information

The technical development on the fibre channel market is fast and rapidly growing. Answers to questions regarding supported configurations and newest technology can be found on the following websites:

HP's current state of FC support

http://techcom.cup.hp.com/dir_fcms/supinfo_index.htm

<http://turbo.rose.hp.com/spock> (HP internal)

EMEA Central SAN Competency Center

http://hprtnt06.grc.hp.com/central_san_cc (HP internal)

this site contains very good training material! It links to the good old *Tachyon Times* website.

WW Storage Support Portal

<http://invent.atl.hp.com/WWStoragePortal/> (HP internal)

Manuals and Release Notes

<http://docs.hp.com/hpux/netcom/#Fibre%20Channel>

Fibre Channel Association

<http://www.fibrechannel.com> (non HP)