



VIOS Shared Storage Pools Phase 3 (also called SSP3)

→ Q4 2012

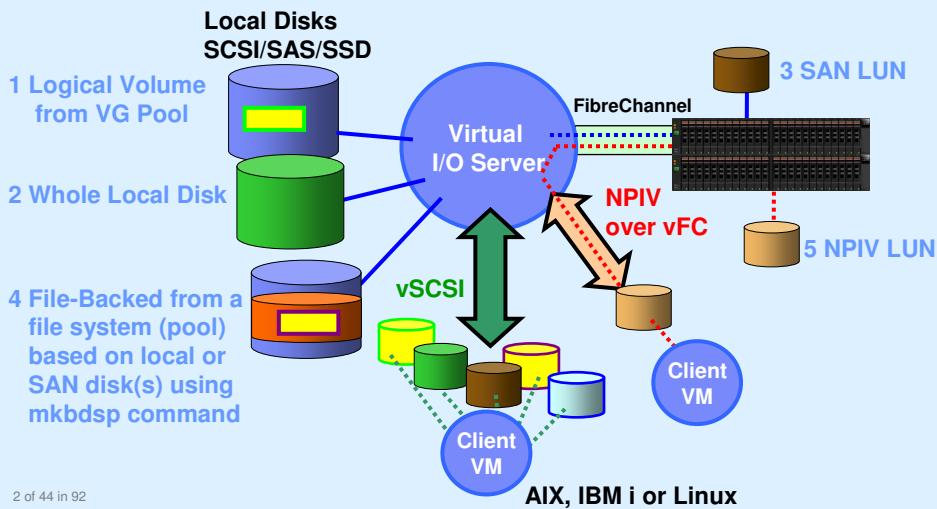


Nigel Griffiths
IBM Power Systems
Advanced Technology Support, Europe

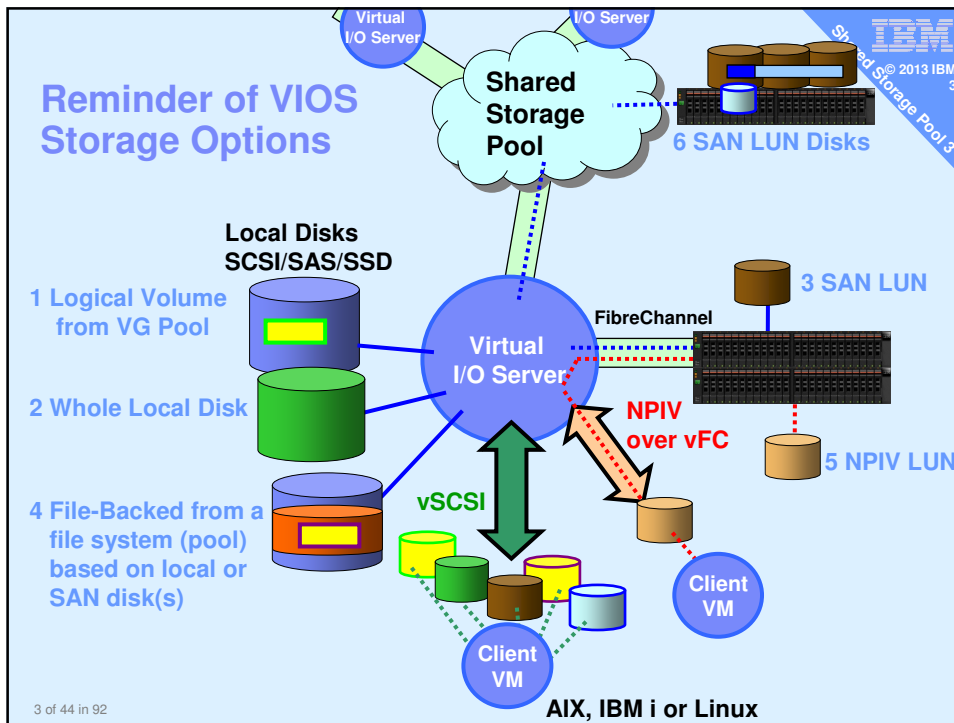
© 2013 IBM Corporation

Reminder of VIOS Storage Options

IBM
© 2013 IBM
Shared Storage Pool 3



2 of 44 in 92



VIO Shared Storage Pool phase 3

- New function delivered with VIOS 2.2.2.1 now 2.2.2.2 fix
- Please read the Readme notes:
 - <http://www.ibm.com/support/docview.wss?uid=hpc1vios117f5701>
 - <http://www.ibm.com/support/docview.wss?uid=hpc1vios610c6192>

VIOS 2.2.1.3 FixPack 25 SP01 Virtual I/O Server 2.2.1.0 VIOS 2.2.1.3 FixPack 25 SP 01 Readme

Readme file for: VIOS 2.2.1.3 FixPack 25 SP01
 Product Component Release: 2.2.1.0
 Update Name: VIOS 2.2.1.3 FixPack 25 SP 01
 Fix ID: VIOS_2.2.1.3_FP25-SP01
 Publication Date: 14 Dec 2011
 Last modified date: 14 Dec 2011

Contents

Download location
Known issues

Installation information

Prior to installation
Installing
Performing the necessary tasks after installation

Rate this page:
Average rating 0 user(s)

Add comments

Document information

Virtual I/O Server

Reference #: 00000676
Modified date: 2011-12-14

- Then read VIOS 2.2.2 User Guide
 - <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hb1/p7hb1.pdf>

4 of 44 in 92

Shared Storage Pool phase 3 Requirements

1 of 3




Read the Readme Notes

- Platforms: **POWER6 & POWER7** only (includes **Blades**)
- VIOS Storage Pool (minimums):
 - Direct fibre-channel attached LUNs:
 - **1 for repository ~10 GB (NEW INFO: 1GB is enough) &**
 - **1 or more for data, 10 GB → in practice lots more [like 1TB+]**
- Pool Storage Redundancy: Repository & pool storage must be **RAIDed**
- VIOS **name resolution** to resolve hostnames
- Virtual I/O Server(s):
 - **Minimum CPU: Entitlement=1+, VP=1+** (shared, uncapped is OK)
 - **Minimum Memory: 4+ GB** (no skinny VIOS)


5 of 44 in 92

Shared Storage Pool phase 3 Limits


2 of 3



- Max VIOS nodes: **16**



(was 4)



- Max physical disks (LUNs) in a pool:
- Max virtual disks (LUs) in a cluster:
- Number of Client LPARs per VIOS (or pair)
- Capacity of Physical Disks in Pool (each)
- Storage Capacity of Storage Pool (total)
- Capacity of each Virtual Disk (LU) in Pool
- Number of Repository Disks

	Phase 3	Phase 2
	1024	(was 256)
	8192	(was 1024)
	200	(was 40)
	16TB	(was 4TB)
	512TB	(was 128TB)
	1GB to 4TB	(same)
	1 (CAA limit)	(same)
	→ new recovery options	

6 of 44 in 92

IBM
© 2013 IBM
7
Shared Storage Pool 3

Shared Storage Pool phase 3 Limits

2 of 3

- Max VIOS nodes:

16

7 of 44 in 92

IBM
© 2013 IBM
8
Shared Storage Pool 3

Shared Storage Pool phase 3 Limits

2 of 3

- Max VIOS nodes: **16** (was 4)

200 VIOS 200 VIOS 200 VIOS ... 16

200 VIOS VIOS 200 VIOS VIOS ... 8

<ul style="list-style-type: none"> Max physical disks (LUNs) in a pool: Max virtual disks (LUs) in a cluster: Number of Client LPARs per VIOS (or pair) 	<table border="0"> <tr> <td style="width: 50%;">Phase 3</td> <td style="width: 50%;">Phase 2</td> </tr> <tr> <td style="color: red;">1024</td> <td>(was 256)</td> </tr> <tr> <td style="color: red;">8192</td> <td>(was 1024)</td> </tr> <tr> <td style="color: red;">200</td> <td>(was 40)</td> </tr> </table>	Phase 3	Phase 2	1024	(was 256)	8192	(was 1024)	200	(was 40)
Phase 3	Phase 2								
1024	(was 256)								
8192	(was 1024)								
200	(was 40)								
<ul style="list-style-type: none"> Capacity of Physical Disks in Pool (each) Storage Capacity of Storage Pool (total) 	<table border="0"> <tr> <td style="width: 50%;">16TB</td> <td style="width: 50%;">(was 4TB)</td> </tr> <tr> <td style="color: red;">512TB</td> <td>(was 128TB)</td> </tr> </table>	16TB	(was 4TB)	512TB	(was 128TB)				
16TB	(was 4TB)								
512TB	(was 128TB)								
<ul style="list-style-type: none"> Capacity of each Virtual Disk (LU) in Pool 	<table border="0"> <tr> <td style="width: 50%;">1GB to 4TB</td> <td style="width: 50%;">(same)</td> </tr> </table>	1GB to 4TB	(same)						
1GB to 4TB	(same)								
<ul style="list-style-type: none"> Number of Repository Disks 	<table border="0"> <tr> <td style="width: 50%;">1 (CAA limit)</td> <td style="width: 50%;">(same)</td> </tr> <tr> <td style="color: red;">→ new recovery options</td> <td></td> </tr> </table>	1 (CAA limit)	(same)	→ new recovery options					
1 (CAA limit)	(same)								
→ new recovery options									

8 of 44 in 92

Shared Storage Pool phase 3 Restrictions

3 of 3



Network

- Reliable & not congested
- DNS should use local /etc/hosts first
- Forward & reverse lookup must work
- Recommended to synchronise clocks
- SEA must use default threaded mode

Storage

- Can't resize a LUN
- SSP may take more CPU
- No SCSI reservations (Reserve/Release)
- HA SAN solutions used to mitigate outages
- SANCOM not supported
- Don't use vSCSI adapter "Any client partition can connect"
- AMS or Suspend/Resume can't use SSP for Paging Space

Lots of restrictions were dropped for this release

9 of 44 in 92

Cluster create on 1st node

Create cluster on one VIOS (here called bluevios1)

```
$ cluster -create -clustername galaxy \  
-repopvs hdisk2 \  
-spname atlantic -sppvs hdisk3 hdisk5 \  
-hostname bluevios1.ibm.com
```

...

Cluster galaxy has been created successfully.

It will take a minute or two, then output Cluster created
You will find a bunch of new daemons running.

If it complains the disks are "in use" check.
If certain they are correct, wipe the disk content with:

```
# cleandisk -r hdiskX  
# cleandisk -s hdiskX
```

It may ask you to confirm y/n ?

10 of 44 in 92

Reminder

Then on that 1st node - add other nodes

On the first VIOS running the cluster

```
$ cluster -addnode -clustername galaxy \  
        -hostname redvios1.ibm.com
```

```
Partition redvios1.aixncc.uk.ibm.com has been added to the galaxy cluster  
$
```

Add other node(s) as necessary.

11 of 44 in 92

IBM
© 2013 IBM
11
Shared Storage Pool 3
Reminder

Allocate disk space & assign to client VM

```
$ mkbdsp -clustername galaxy \  
        -sp atlantic 16G -bd vdisk_diamond6a \  
        -vadapter vhost2
```

```
Logical Unit vdisk_diamond6a has been created with udid:  
615af85de5acad39a8827e9cd01d6b36.  
Assigning file "vdisk_diamond6a" as a backing device.  
Vtscsi3 Available.  
$
```

Notes:

- 16 GB is not actually allocated until written too
- vdisk_diamond6a is just a name = reminder of the VM using it
- vhost2 is the virtual SCSI adapter for client VM diamond6

12 of 44 in 92

IBM
© 2013 IBM
12
Shared Storage Pool 3
Reminder

Skipping SSP phase 2 ...

- Cluster command for details of the VIOSs
- Setting up Dual VIOS multiple pathing to SSP
- Add a new LUN to grow the pool
- Replace a LUN also allows migration

- Thin Provisioning of disk space at 64 MB chunk
 - Optional regular provisioning with -thick ☺
- Alerting on Pool Space getting LOW !!
 - There is a new Alert for extreme Over-Commit level
- Snapshots
 - Learnt a rollback to earlier snapshot removes later ones

13 of 44 in 92

Reminder

IBM
© 2013 IBM
13
Shared Storage Pool 3

Monitoring Disk use with lssp – SSP3

```
$ lssp -clustername stars -sp atlantic -bd
Lu Name      Size(mb) ProvisionType %Used Unused(mb) Lu Udid
orange7a     32768    THIN           9% 29615      7d58538152 ...
orange7b     32768    THIN           0% 32770      76136907aa ...

$ lssp -clustername stars
POOL_NAME:      atlantic
POOL_SIZE:      130944
FREE_SPACE:     125514
TOTAL_LU_SIZE:  65536
OVERCOMMIT_SIZE: 0
TOTAL_LUS:      2
POOL_TYPE:      CLPOOL
POOL_ID:        000000009893E510000000050740962
```

See who is using most disk
& who might run out

14 of 44 in 92

IBM
© 2013 IBM
14
Shared Storage Pool 3

Monitoring Disk use with lssp – SSP3

```

$ lssp -clustername stars -sp atlantic -bd
Lu Name      Size (mb) ProvisionType %Used Unused (mb) Lu Udid
orange7a    32768    THIN           9% 29615      7d58538152 ...
orange7b    32768    THIN           0% 32770      76136907aa ...

$ lssp -clustername stars
POOL_NAME:    atlantic
POOL_SIZE:    130944
FREE_SPACE:   125514
TOTAL_LU_SIZE: 65536
OVERCOMMIT_SIZE: 0
TOTAL_LUS:    2
POOL_TYPE:    CLPOOL
POOL_ID:      0000000009893E510000000050740962
    
```

See who is using most disk & who might run out

Over-commit
Good to know the "worst case"

15 of 44 in 92

Four User Interface

- VIOS Command Line (CLI)
 - Already shown in this presentation by example
 - Some feature make sense via CLI → like: cluster –create/-addnode
- VIOS cfgassist menu
 - This is the VIOS version of smitty
 - Menu driven interface for CLI
 - Fully covers SSP functions
- HMC Graphical User Interface (HMC 7.7.4+)
 - SSP virtual disk list, create and connect to your VMs
 - Note: don't have any VIOS vSCSI slots in "Any node" = bad practice anyway
 - See screen shots ...
- System Director – **New in SSP3**
 - Already has Storage Pools concept and features
 - SSP is just another storage pool type
 - Then may adds new unique items – like VMControl appliance deploy to SSP disk space or Linked-Cloning


Reminder

New

16 of 44 in 92

Hardware Management Console

- HMC from October 2011
 - For SSP support & LPM of SSP LPARs
 - Addition feature to Virtual Storage Management



- Shipped with HMC upgrade V7 R7.4 SP0+
 - Not part of the VIOS package

HMC

Virtual Disk	Storage Pool	Assigned Partition	Size	Disk Type
vdisk_diamond5a	atlantic	diamond5-AIX7-SSP2(5)	16 GB	Thin
vdisk_diamond6a	atlantic	diamond6-AIX7-SSP2(3)	16 GB	Thin

Virtual Storage Management - diamond-8233-E8B-SN100271P

Use virtual storage management tasks to manage your Shared Storage Pool (SSP) Devices. Select a query.

VIOS/SSP: [galaxy] Query

Storage Details

Virtual disks are logical entities on the VIOS partitions. To perform management tasks for virtual disks, select the task to perform. You also can create a virtual disk.

Create Virtual Disk - diamond-8233-E8B-SN100271P

To create a virtual disk, enter a name and a size for the new disk, select a storage pool from which to create the new disk. You also select the new disk to a logical partition. This task can take several minutes to complete if you are creating a virtual disk in a file-based storage pool.

Virtual disk name: disk_diamond8a
 Storage pool name: atlantic(galaxy) (46.6 GB free, 51.62 GB total)
 Virtual disk size: 16 GB
 Assigned partition: diamond8-AIX7 TL1 beta(8)
 Disk type: Thick
 Map to VIOS(s): Select Virtual IO Server
 diamondvios1-SSP

OK Cancel Help

SSP Pool

HMC

Virtual Storage Management - diamond-8233-E8B-SN100271P

Use virtual storage management tasks to manage virtual storage for your VIOS virtual servers and your Shared Storage Pool (SSP) Devices. Select a VIOS virtual server or an SSP Device to query.

VIOS/SSP:

Storage Details

Virtual Disks | Storage Pools | Physical Volumes | Optical Devices | Virtual Fibre Channel

Virtual Storage Management - diamond-8233-E8B-SN100271P

Use virtual storage management tasks to manage virtual storage for your VIOS virtual servers and your Shared Storage Pool (SSP) Devices. Select a VIOS virtual server or an SSP Device to query.

VIOS/SSP:

Storage Details

Virtual Disks | Storage Pools | Physical Volumes | Optical Devices | Virtual Fibre Channel

Virtual disks are logical entities on the VIOS partition that provide storage for client partitions. To perform management tasks for existing virtual disks, select a virtual disk then select the task to perform. You also can create a new virtual disk.

--- Select Action ---

Storage Pool	Assigned Partition	Size
clientvg	None	25 GB
clientvg	None	256 MB
clientvg	None	16 GB
clientvg	None	16 GB
clientvg	None	16 GB
clientvg	None	16 GB
clientvg	diamond8-AIX7 TL1 beta(8)	16 GB

Show shared storage pool storage

--- Select Action ---

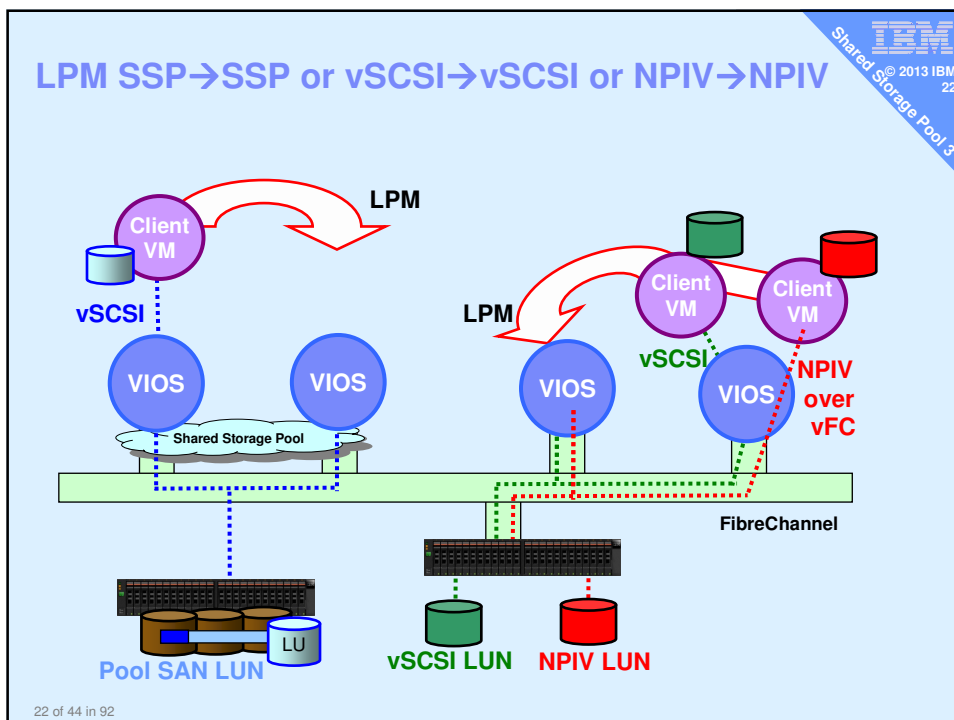
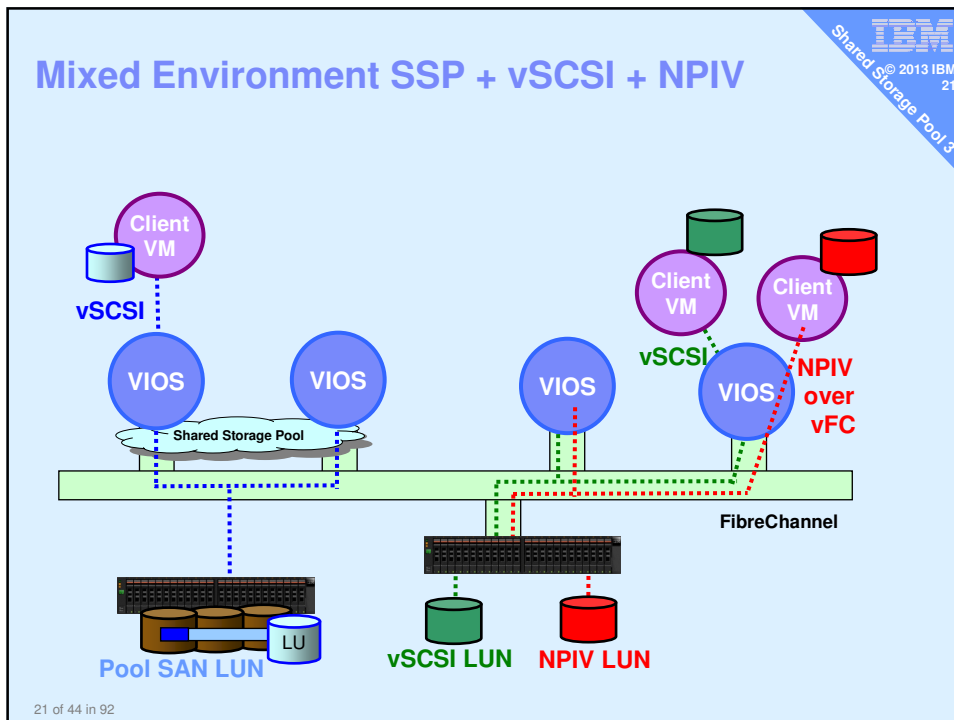
Select	Name	Storage Pool	Assigned Partition	Size
<input type="radio"/>	fslv00	clientvg	None	25 GB
<input type="radio"/>	loglv00	clientvg	None	256 MB
<input type="radio"/>	vdisk_diamond5a	atlantic(galaxy)	diamond6-AIX616-SSP(3)	16 GB
<input type="radio"/>	vdisk_diamond6a	atlantic(galaxy)	diamond6-AIX616-SSP2(3)	16 GB
<input type="radio"/>	vdisk_diamond8a	atlantic(galaxy)	diamond8-AIX7 TL1 beta(8)	16 GB
<input type="radio"/>	xdiamond4_1lv	clientvg	None	16 GB
<input type="radio"/>	xdiamond7_2lv	clientvg	None	16 GB
<input type="radio"/>	xdiamond7_4lv	clientvg	None	16 GB
<input type="radio"/>	xdiamond8	clientvg	diamond8-AIX7 TL1 beta(8)	16 GB

Show shared storage pool storage


SSP3 in practice

IBM
© 2013 IBM
20
Shared Storage Pool 3

20 of 44 in 92




Can mix on the same VIOS



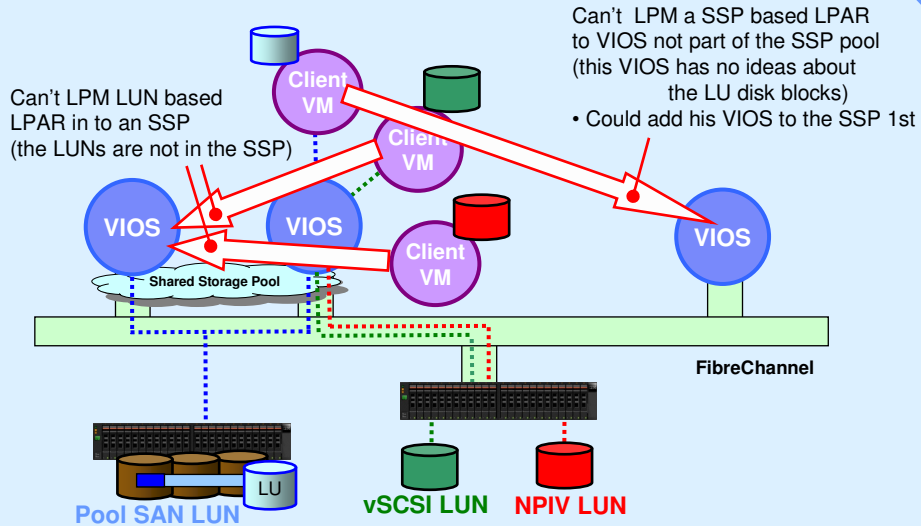
23 of 44 in 92

Can LPM



24 of 44 in 92


Can't LPM → if it means a different disk layout



25 of 44 in 92

Adopt an AIX LPAR

Want to move your client VM to use the SSP!

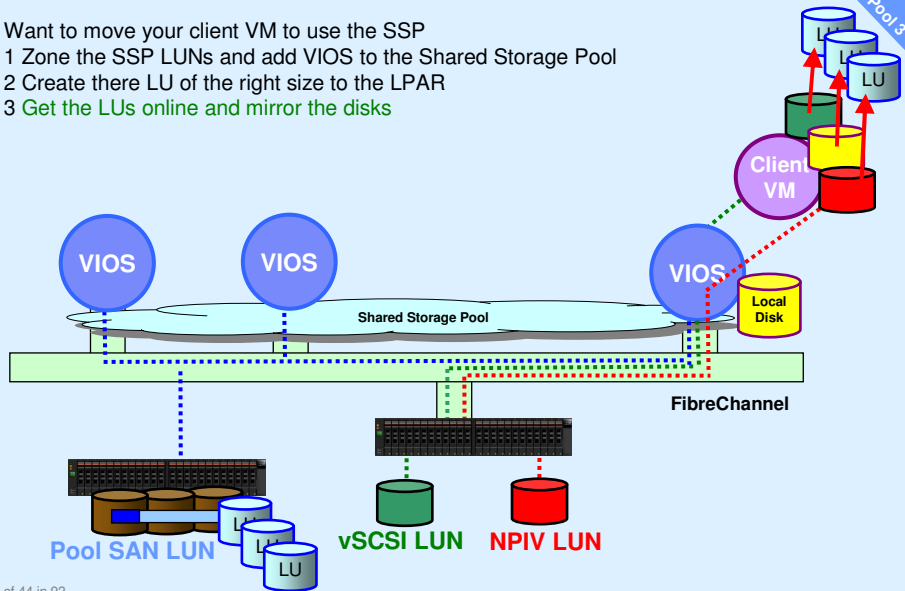


26 of 44 in 92

Adopt an AIX LPAR

Want to move your client VM to use the SSP

- 1 Zone the SSP LUNs and add VIOS to the Shared Storage Pool
- 2 Create there LU of the right size to the LPAR
- 3 Get the LUs online and mirror the disks

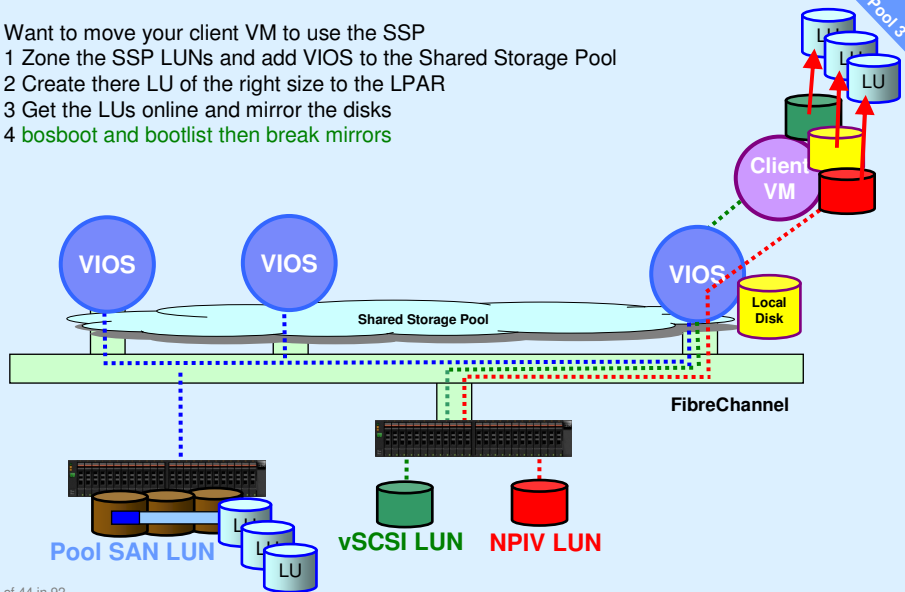


29 of 44 in 92

Adopt an AIX LPAR

Want to move your client VM to use the SSP

- 1 Zone the SSP LUNs and add VIOS to the Shared Storage Pool
- 2 Create there LU of the right size to the LPAR
- 3 Get the LUs online and mirror the disks
- 4 bosboot and bootlist then break mirrors




30 of 44 in 92

Adopt an AIX LPAR

Want to move your client VM to use the SSP

- 1 Zone the SSP LUNs and add VIOS to the Shared Storage Pool
- 2 Create there LU of the right size to the LPAR
- 3 Get the LUs online and mirror the disks
- 4 bosboot and bootlist then break mirrors
- 5 Remove original disks

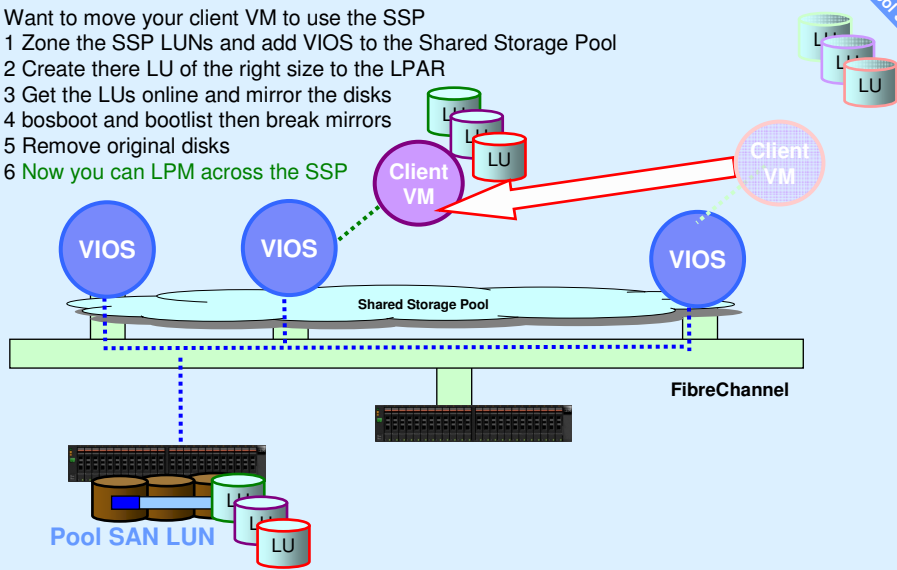


31 of 44 in 92

Adopt an AIX LPAR

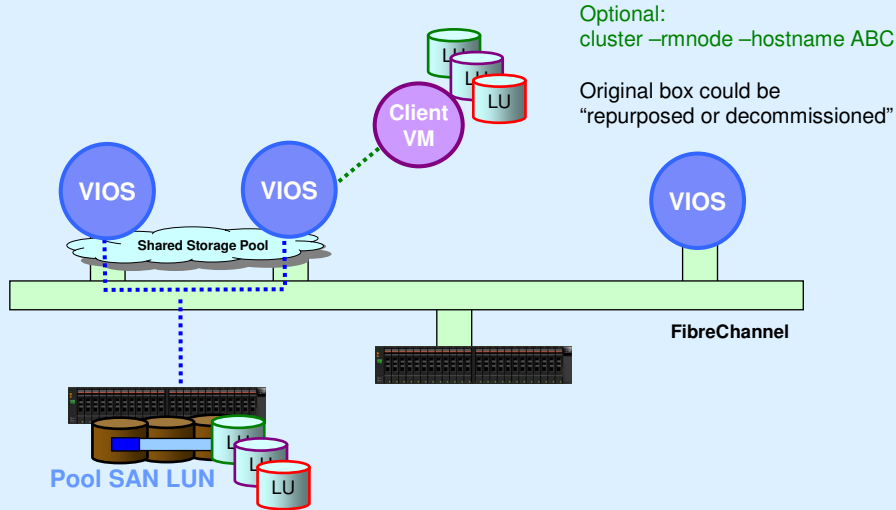
Want to move your client VM to use the SSP

- 1 Zone the SSP LUNs and add VIOS to the Shared Storage Pool
- 2 Create there LU of the right size to the LPAR
- 3 Get the LUs online and mirror the disks
- 4 bosboot and bootlist then break mirrors
- 5 Remove original disks
- 6 Now you can LPM across the SSP




32 of 44 in 92

Adopt an AIX LPAR = "You will be assimilated!" The Borg.

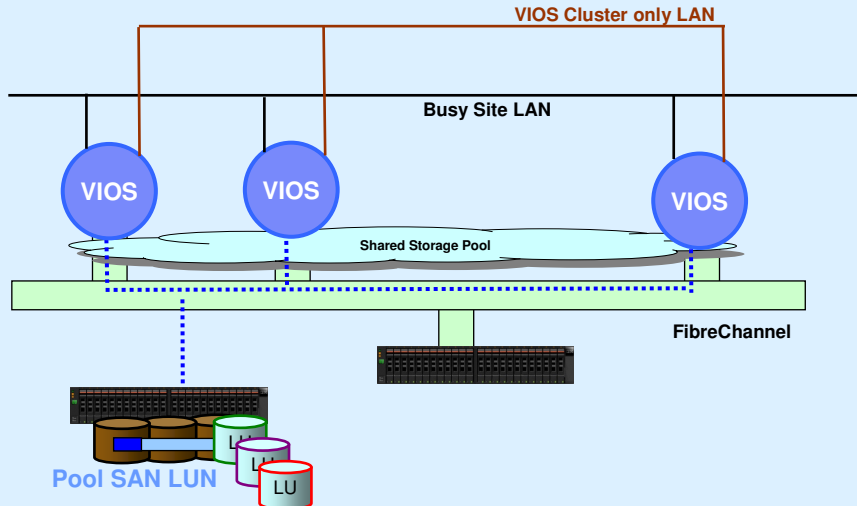


Dedicated VIOS Cluster network



Dedicated VIOS Cluster network


Setup network route to the other VIOS to use this new network interface
Not really a VIOS SSP feature but now tested & supported



35 of 44 in 92

Production Separation

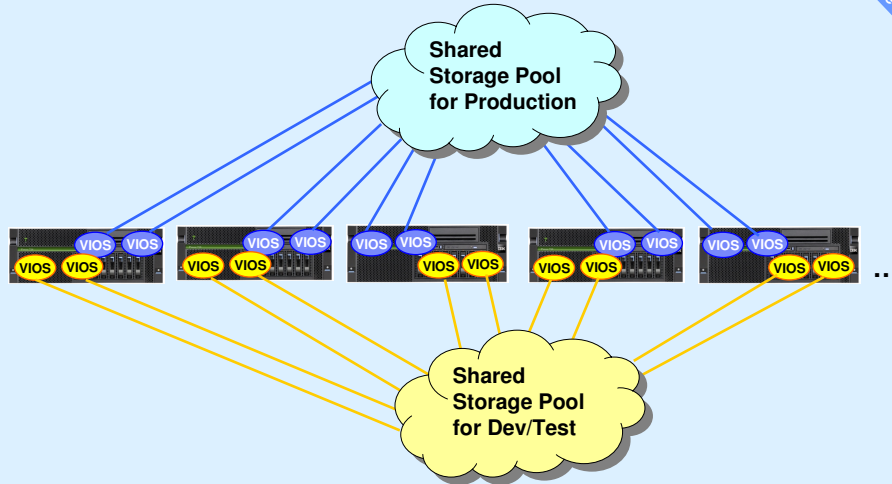
Popular configuration is:
dual VIOS for Production
dual VIOS for Other workloads



How can we add Shared Storage Pool and keep them separate ?

36 of 44 in 92

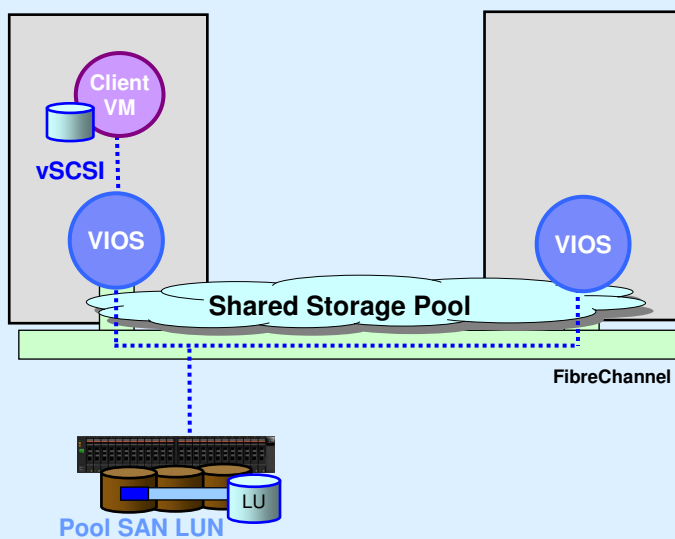
One machine on multiple pools



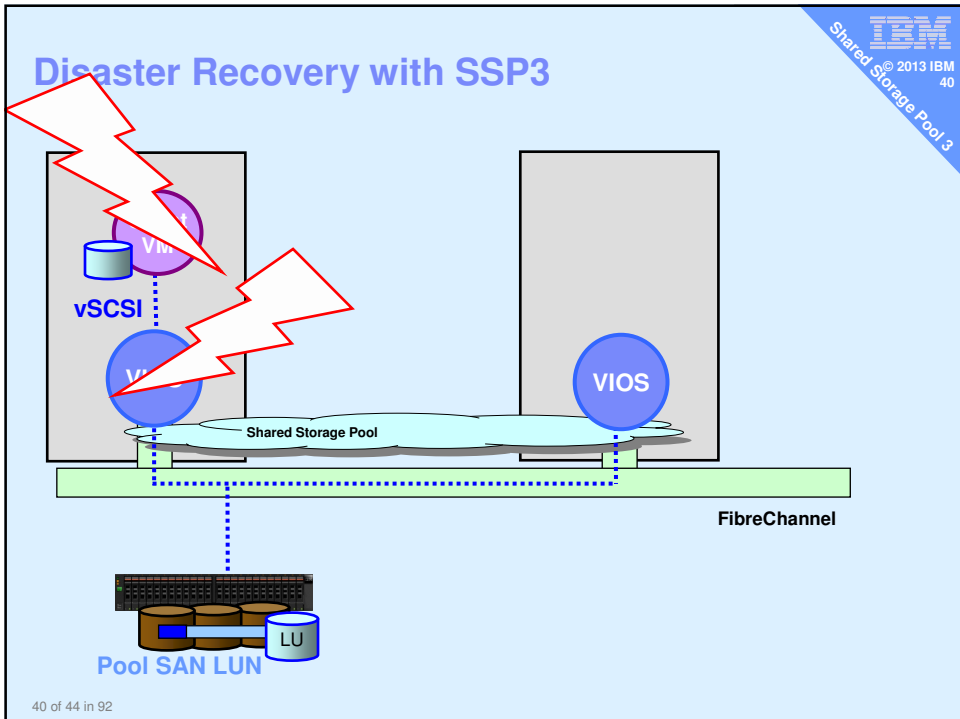

Note: Any one VIOS can only be part of only one SSP cluster but different VIOS's can be on difference SSP clusters

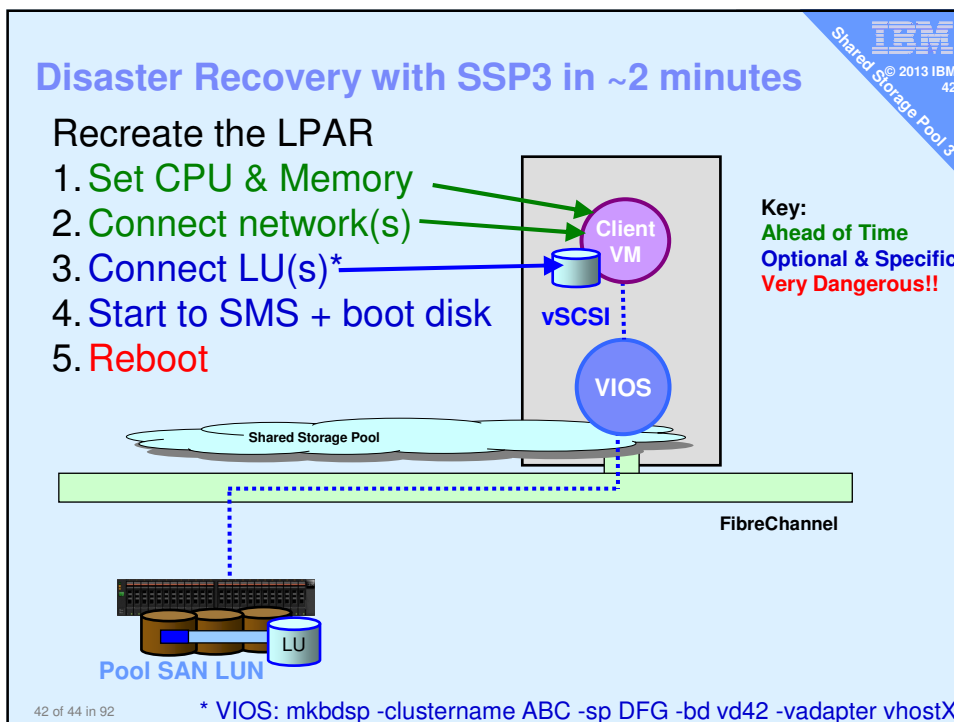
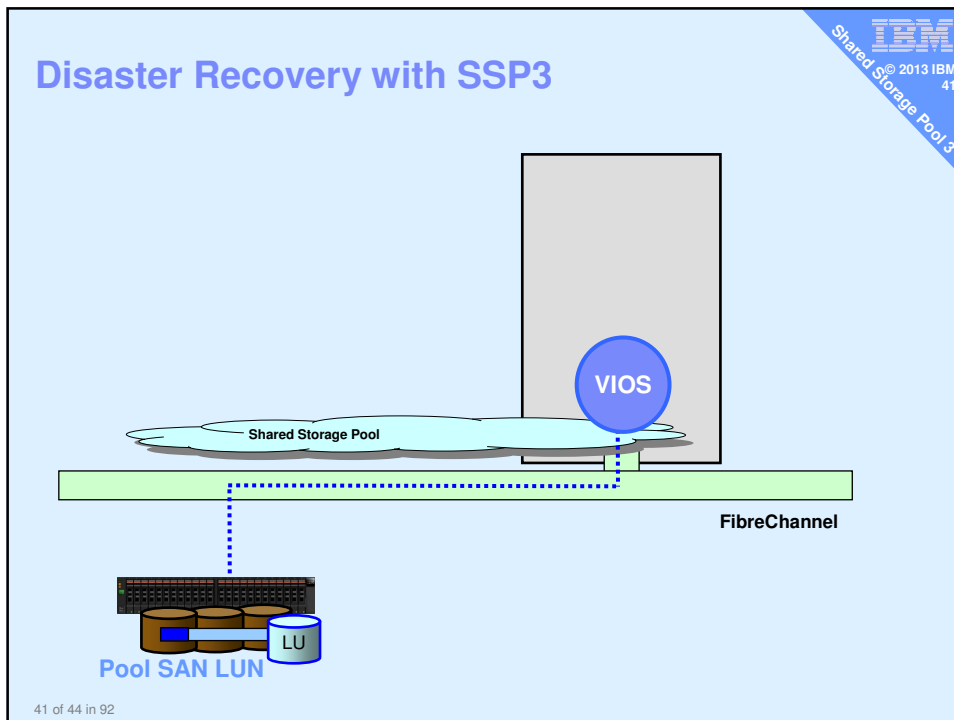
37 of 44 in 92

Disaster Recovery with SSP3



38 of 44 in 92





Want more?

Shared Storage Pools - five videos to Get You Started Today

- https://www.ibm.com/developerworks/community/blogs/aixpert/entry/shared_storage_pools_five_videos_to_get_you_started_today
- [Shared Storage Pool phase 3 \(SSP3\) New Features](#) - presentation on whats new.
- [Migrating to Shared Storage Pool \(SSP3\) & then LPM](#) - get a LPAR on local disk to SSP in 5 minutes and then LPM to a new machine in 1 minute
- [SSP3 Recover a Crashed Machine's LPAR to Another Machine](#) - not everything need HACMP but it would be nice to get those LPAR running quickly
- [Live Partition Mobility \(LPM\) with Shared Storage Pool SSP3](#) - load balance across the computer room, evacuate a machine for maintenance and use a new machine on day 1
- [Looking around a Shared Storage pool via commands and HMC](#) - see how easy it is to understand and operate SSP

45 of 44 in 92

BACKUP – 50 more slides


Topics

- Clusters
- Snapshot
- Working with Dual VIOS
- Thin Provisioning
- Alerts when the pool is nearly empty
- Storage mobility and repository recovery
- SSP commands reminder sheet

- Demo details – now 4 major movies available

46 of 44 in 92


Terms Shared Storage Pool 3




1 SSP2 cluster = set of co-operating Virtual I/O Servers
The VIOS can only be in one cluster

Here we show two clusters

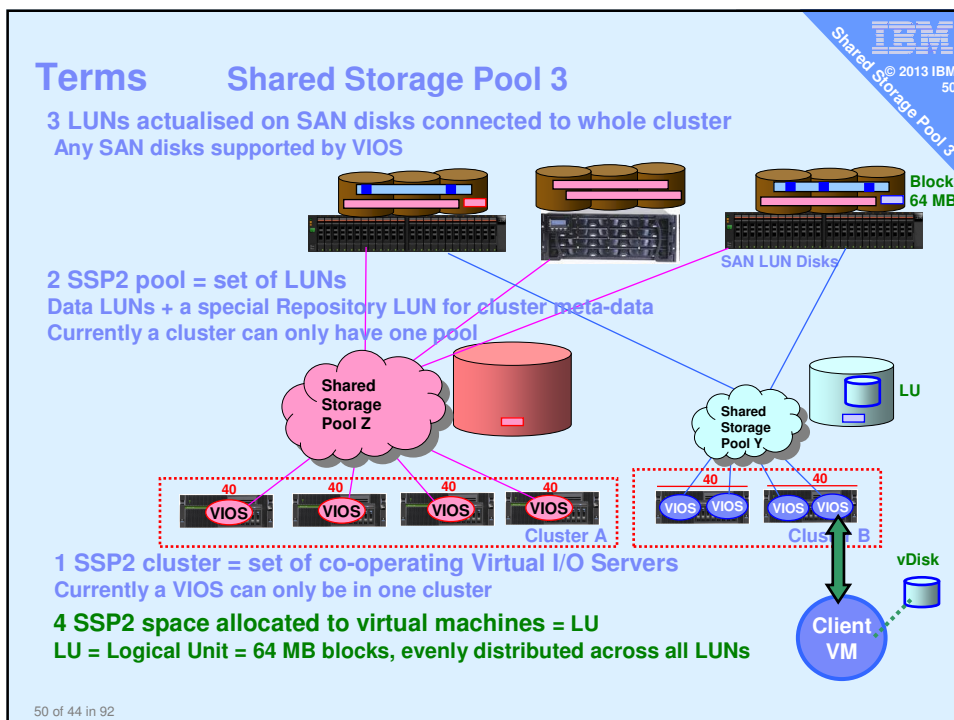
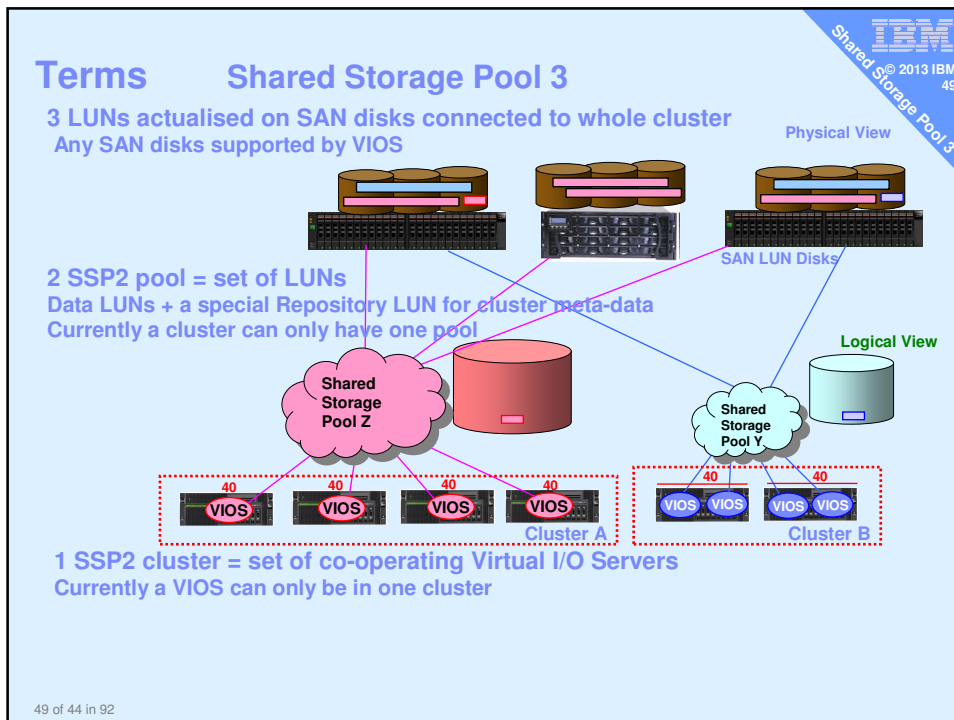
Terms Shared Storage Pool 3

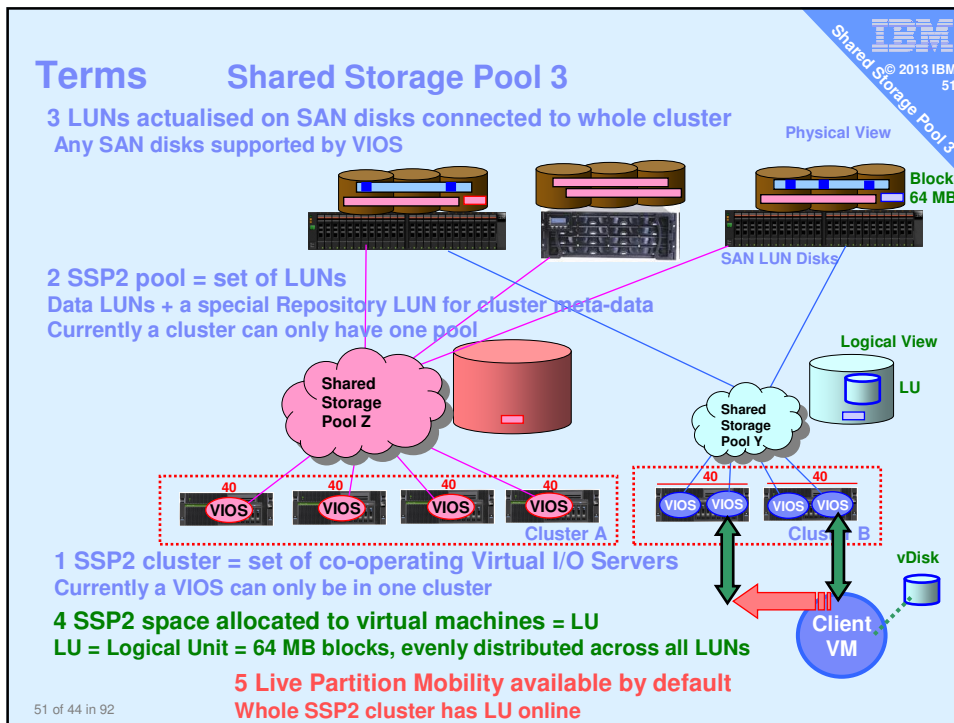


2 SSP2 pool = set of LUNs
Data LUNs + a special Repository LUN = cluster meta-data
Currently a cluster can only have one pool



1 SSP2 cluster = set of co-operating Virtual I/O Servers
Currently a VIOS can only be in one cluster





- Preparation**
- All the Cluster VIOSs need the LUNs online
Make sure they are available = Zoned
 - **BEFORE** you start the cluster for all LUNs
`chdev -dev <device name> -attr reserve_policy=no_reserve`
 - Don't forget this for extra disks that you add later
 - Forgetting this = a real mess to make the low level disk attribute change
-
- IBM © 2013 IBM 52
Shared Storage Pool 3
- 52 of 44 in 92

List cluster & cluster nodes

```
$ cluster -list
CLUSTER_NAME:      galaxy
CLUSTER_ID:        64517962b01c11e1ac6aba367e934e03
$
$ cluster -status -clustername galaxy
Cluster Name      State
galaxy            OK

Node Name          MTM                Partition Num  State  Pool
                  8233-E8B02100271P  2             OK     OK
diamondvios1
diamondvios2      8233-E8B02100271P  1             OK     OK
redvios1          8203-E4A0310E0A41  1             OK     OK
$
```



Example of a 3 node cluster

House keeping

You can remove a node from the cluster

- LPM any important client Virtual machines elsewhere
- Stop remaining VMs
- Remove client VMs
- Remove virtual disks
- then

```
$ cluster -rmnode -clustername galaxy \  
-hostname redvios1.ibm.com
```

You can also remove the cluster completely

- Once all disk space unassigned & nodes removed

```
$ cluster -delete -clustername galaxy
```

Pool Disk Space Management

Content

- Allocate pool disk space and give to a VM
 - Ditto as two commands
- Removing the disk space
- Monitoring the pool

Allocate disk space & assign to client VM

```
$ mkbdsp -clustername galaxy \  
-sp atlantic 16G -bd vdisk_diamond6a \  
-vadapter vhost2
```

```
Logical Unit vdisk_diamond6a has been created with udid:  
615af85de5acad39a8827e9cd01d6b36.  
Assigning file "vdisk_diamond6a" as a backing device.  
Vtscsi3 Available.  
$
```

Notes:

- 16 GB is not actually allocated until written too
- vdisk_diamond6a is just a name = reminder of the VM using it
- vhost2 is the virtual SCSI adapter for client VM diamond6

Same but 2 steps (create LU then assign it)

1) Create Logical Unit Note: no -vadapter option
`$ mkbdsp -clustername galaxy -sp atlantic 10G -bd LU42`
 Lu Name:LU42
 Lu Udid:374a609cb072e4015d558ff290b9f0bd

List the pool contents

```
$ lssp -clustername galaxy -sp atlantic -bd
Lu Name      Size(mb) ProvisionType  Lu Udid
LU42         10240      THIN              374a609cb072e4015d558ff290b9f0bd
...
```

2) Example of two ways using "-bd LU42" or "-luudid <hexidecimal>"

- -bd only works if LU42 is unique
- Note: below **no Size argument** (or it creates another LU with same name!)

```
$ mkbdsp -clustername galaxy -sp atlantic -bd LU42 -vadapter vhost2
Assigning file "LU42" as a backing device.
VTD:vtscsi1
```

- or -

```
$ mkbdsp -clustername galaxy -sp atlantic \
-luudid 374a609cb072e4015d558ff290b9f0bd -vadapter vhost2
Assigning file "374a609cb072e4015d558ff290b9f0bd" as a backing device.
VTD:vtscsi1
```

Dual path via Two VIOSs

1 Setup virtual SCSI adapter pairs as normal

- client VM virtual SCSI adapter A ↔ VIOS C
- client VM virtual SCSI adapter B ↔ VIOS D

2 on VIOS C: use "lsmap -all" to map slot to vhostN

- `mkbdsp -clustername galaxy -sp atlantic 16G -bd vdisk_red6a -vadapter vhostN`


3 on VIOS D: use lsmap -all to map slot to vhostM

- `mkbdsp -clustername galaxy -sp atlantic -bd vdisk_red6a -vadapter vhostM`

4 On the client VM

```
$ lspath
Enabled hdisk0 vscsi0
Enabled hdisk0 vscsi1
```

Note: No size (16G) 2nd time



5 LPM still available – dual VIOS to dual VIOS

Removing an LU (Logical Unit)

Assuming it is NOT used !!

On the VIOS remove disk space
rmbdsp = remove backing device from storage pool

By name

```
$ rmbdsp -clustername galaxy -sp atlantic  
-bd vdisk_diamond6a
```

By LU hexadecimal id

```
$ rmbdsp -clustername galaxy -sp atlantic  
-luudid 858152297879adfe0d75b05f586d36ee
```

59 of 44 in 92

IBM
© 2013 IBM
59
Shared Storage Pool 3
Reference Only

House keeping

Add more physical LUNs to the Pool

```
$ chsp -add -clustername galaxy -sp atlantic  
hdisk8
```

Remove a physical LUN from the Pool

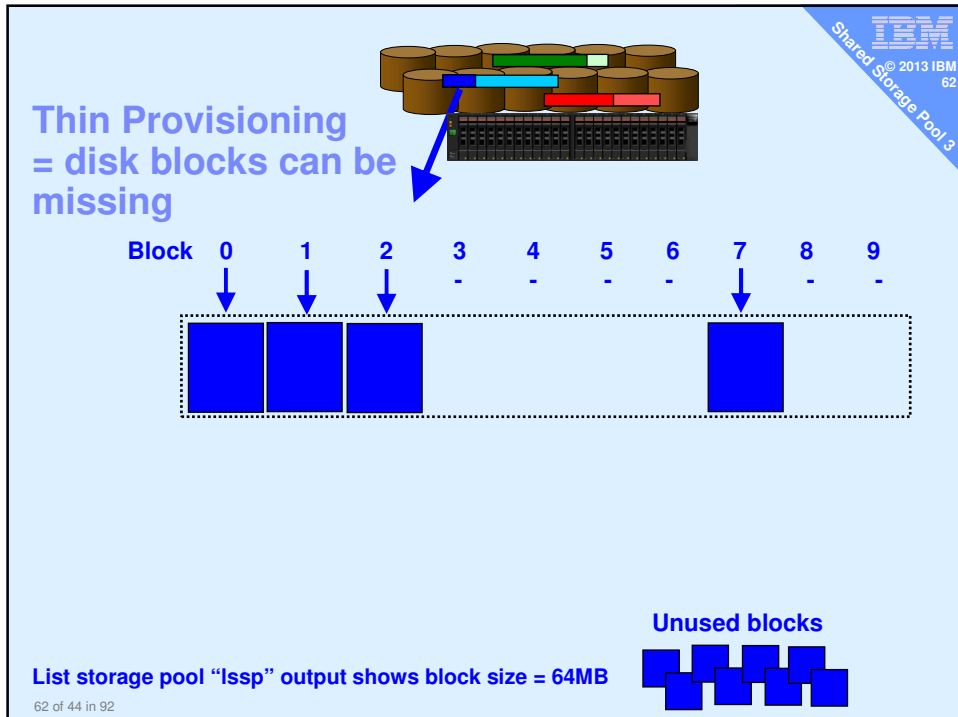
- You can not with this release
- We can replace a disk but not remove one
 - Replacement disk - equal or larger size

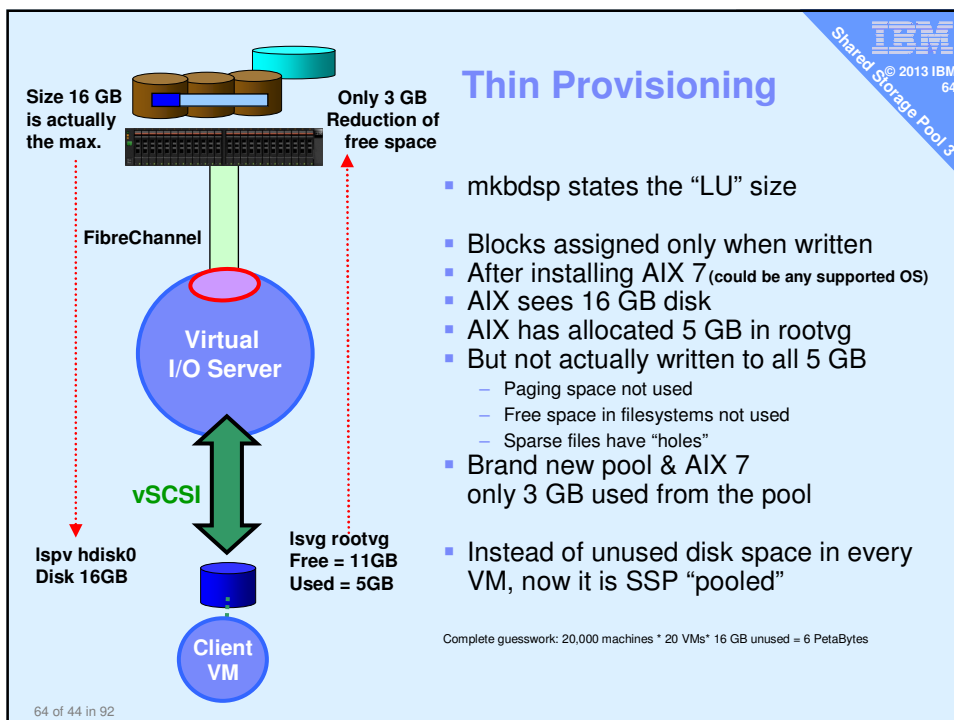
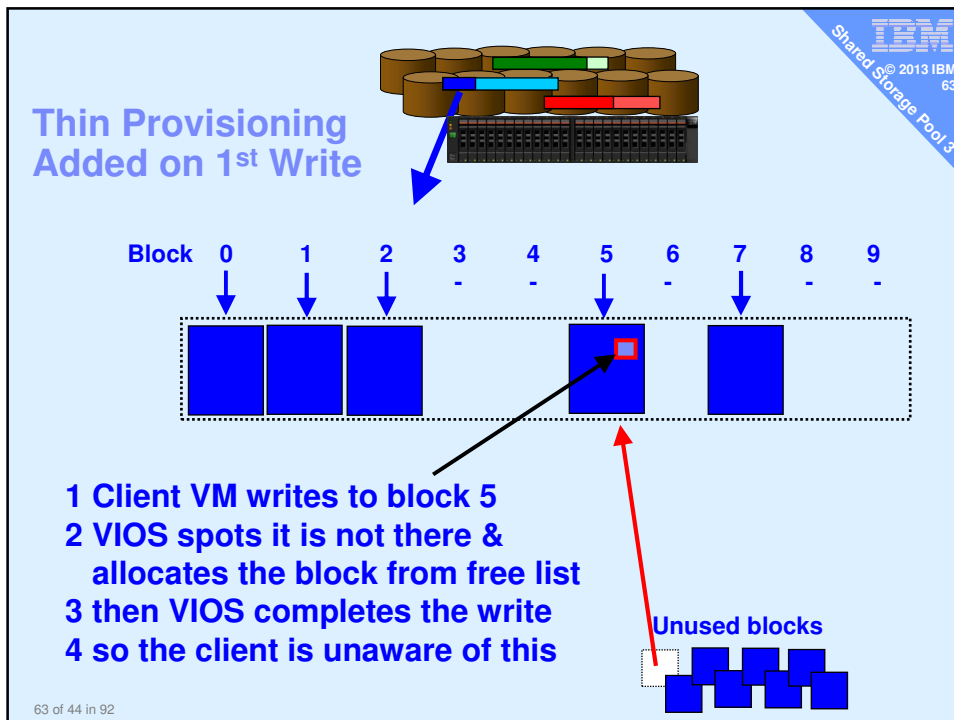
60 of 44 in 92

IBM
© 2013 IBM
60
Shared Storage Pool 3

Experiments in Thin provisioning

= Allocating disk blocks only when they are used i.e. written





Thick Provisioning

- Doh! A no-brainer!
- Like Thin but actually allocate all the disk space
- New option: `mkbdsp ... -thick`

The point is

- No problem, if the free list empties
- Good for more important work/production or you prefer not to dynamically add blocks

Monitoring: topas on VIOS then "D"

```

Topas Monitor for host:      diamondvios1Interval: 2      Fri Jan 14 14:46:00 2011
=====
Disk  Busy%  KBPS    TPS    KB-R    ART    MRT    KB-W    AWT    MWT    AOW    AOD
cldisk2  41.0  17.6K  493.0    0.0    0.0  174.6   17.6K   1.1  14.6    0.0    0.0
cldisk3  34.0  20.0K  160.0    0.0    0.0  186.4   20.0K   2.9  13.1    0.0    0.0
cldisk1   3.0   24.0    6.0    0.0    0.0  112.0    24.0    0.6 158.8    0.0    0.0
hdisk0   0.0    8.0    2.0    0.0    0.0   10.2    8.0    4.1  64.2    0.0    0.0
caa_priva 0.0  17.0    5.0    9.0    0.1    2.1    8.0    0.5   6.9    0.0    0.0
hdisk1   0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0   7.2    0.0    0.0
cd0      0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0   0.0    0.0    0.0
  
```

One client VM running: yes `>/tmp/x`

Disk I/O spread across disks

Allocation unit is 64MB (was in the lssp output)

Monitoring Disk use with lssp – SSP2

```
$ lssp -clustername galaxy -sp atlantic -bd
```

Lu(Disk) Name	Size(MB)	ProvisionType	Lu	Udid
vdisk_diamond6a	16384	THIN	615af	. . .
vdisk_diamond8a	16384	THIN	917c0	. . .
vdisk_diamond5a	8192	THICK	f1442	. . .
vdisk_diamond5b	8192	THICK	ebedc	. . .
vdisk_diamond3a	10240	THIN	afcec	. . .

```
$ lssp -clustername galaxy
```

```
POOL_NAME: atlantic
POOL_SIZE: 47552
FREE_SPACE: 17945
TOTAL_LU_SIZE: 59392
TOTAL_LUS: 5
POOL_TYPE: CLPOOL
POOL_ID: 000000009893EDD000000004F174D22
```

47522 Pool Physical Size
17945 Pool Physical Free
29607 Pool Physical Used
Pool use $29607/47522 \times 100 = 62\%$

59392 Allocated
Pool Over commit $59392/47522 = 1.25$
allocated 25% more than I have!
= Thin provisioning

Monitoring Disk use with lssp – SSP3

```
$ lssp -clustername stars -sp atlantic -bd
```

Lu Name	Size(mb)	ProvisionType	%Used	Unused(mb)	Lu	Udid
orange7a	32768	THIN	9%	29615	7d58538152	...
orange7b	32768	THIN	0%	32770	76136907aa	...

```
$ lssp -clustername stars
```

```
POOL_NAME: atlantic
POOL_SIZE: 130944
FREE_SPACE: 125514
TOTAL_LU_SIZE: 65536
OVERCOMMIT_SIZE: 0
TOTAL_LUS: 2
POOL_TYPE: CLPOOL
POOL_ID: 000000009893E510000000050740962
```

See who is using most disk
& who might run out

Over-commit
Good to know the "worst case"




IBM
© 2013 IBM
69
Shared Storage Pool 3

Thin provisioning risks running out of space → Ek!
Next write needing a new SSP block, gets a disk error!
Just don't go there – you need to be warned!

69 of 44 in 92

Thin Provisioning Alerts



IBM
© 2013 IBM
70
Shared Storage Pool 3

- To list the alert threshold:
 - **alert -list -clustername *galaxy* -spname *atlantic***
 - \$ **alert -list -clustername *galaxy* -spname *atlantic***
 - PoolName: atlantic
 - PoolID: 0000000009893EDD000000004F174D22
 - ThresholdPercent: 35 ←these are the defaults
 - OverCommitPercent: n/a
- Set alerts to warn on free pool space getting too low %
 - **alert -set -clustername *galaxy* -spname *atlantic***
 - type threshold -value 10**
- Set alerts to warn on overcommit getting too large %
 - **alert -set -clustername *galaxy* -spname *atlantic***
 - type overcommit -value 30**
- To remove the alert:
 - **alert -unset -clustername *galaxy* -spname *atlantic* -type threshold**
 - Threshold is set to 0 (zero) – it will not happen!!

70 of 44 in 92

House keeping – Alert Reporting

- Reported on **any one of the VIOS cluster**

- padmin user: errlog

- Like AIX errpt
- \$ errlog | more

```
IDENTIFIER  TIMESTAMP  T C  RESOURCE_NAME  DESCRIPTION
0FD4CF1A   0215112612 I O  VIOD_POOL      Informational Message
```

2

- \$ errlog -ls | more

See example on the next page

- Can also be reported to high levels SM like Systems Director etc.



```
$ errlog -ls ...
LABEL:      VIO_ALERT_EVENT 1
IDENTIFIER: 0FD4CF1A
```

```
Date/Time:   Wed Feb 15 11:26:32 CST 2012
Sequence Number: 86
Machine Id:  00F602714C00
Node Id:     diamondvios2
Class:       O
Type:        INFO
WPAR:        Global
Resource Name: VIOD_POOL 2
```

```
Description
Informational Message 3
```

```
Probable Causes
Asynchronous Event Occurred
```

```
Failure Causes
PROCESSOR
```

```
Recommended Actions
Check Detail Data
```

```
Detail Data
Alert Event Message 4
```

```
25b8001
A Storage Pool Threshold alert event occurred on pool D_E_F_A_U_L_T_061310 pool id 92d2fd5f2ec45382 in cluster galaxy cluster id 00841e2a422711e194cbf60271715fc2 The alert event received is: Threshold Exceeded.
```

```
Diagnostic Analysis
Diagnostic Log sequence number: 250
Resource tested: sysplanar0
Menu Number:      25B8001
```

```
Description:
A Storage Pool Threshold alert event occurred on pool D_E_F_A_U_L_T_061310 pool id 92d2fd5f2ec45382 in cluster galaxy cluster id 00841e2a422711e194cbf60271715fc2 The alert event received is: Threshold Exceeded.
```

5



House keeping - Thin Provisioning Alerts

- It is vital that you get these warning messages
- Suggest on ALL VIOS
 1. Email the Pool stats every night to the admin guys (cron as root)

```
. /home/padmin/.profile
lssp -clustername galaxy | /usr/bin/mailx -s "SSP stats" ops@acme.com
```

2. Script to check and if free space is low then email or send phone TEXT message or escalate
- Possible reactions are:
 - Add a new LUN to the pool,
 - Delete allocated space = unused LU or entire VM & space
 - Drop a Snapshot
 - Look and check the larger VM really need the space.

House keeping - Thin Provisioning Alerts

IBM Systems Director captures these events


- Discover, access and Inventory all VIOS.

The screenshot shows the IBM Systems Director web interface. The 'Problems' tab is active, displaying a table of active problems for the system 'goldvios1'. A tooltip is shown over the first row, which is a 'Warning' event.

Sel...	Name	Severity	System	Component	Category	Time Re...	Dx
<input type="checkbox"/>	A Storage Pool Thres...	Warning	goldvios1	goldvios1	Hardware Status	15 Feb 201...	A

Tooltip text: A Storage Pool Threshold alert event occurred on pool D_E_F_A_U_J_I_T_061310 pool id 92d2f5f2ec45362 in cluster galaxy cluster id 00841e2a422711e194cbf60271715fc2 The alert event received is: Threshold Exceeded.

Snapshot



Snapshots on VM disks and Cloning

Snapshot available using

- Advanced SAN disks or SAN Volume Controller (SVC)

but now VIOS admin can do this too!

Currently: no way
to save the
snapshot off-line

Snapshot + Drop

- Very quick
- Allows point in time backup
- Later delete the original to reclaim the space

Examples:

- Backup VM stopped, quiesce, live

Snapshot + Roll-back

- Very quick
- Useful for lots of reasons →
- Stop the client VM
- Restart on original copy
- Discard newer copy

Examples:

- Practice OS or App update
- Training & reset
- Benchmark & reset
- Failure & avoid recovery from tape
- Save points for batch runs

Supports single disk or a consistent set of disks

Snapshot – create, list, delete or rollback

Snapshot Usage:

```
snapshot -create <filename> -clustername galaxy -spname atlantic -lu LUs
```

```
snapshot -delete <filename> -clustername galaxy -spname atlantic -lu LUs
```

```
snapshot -rollback <filename> -clustername galaxy -spname atlantic -lu LUs
```

```
snapshot -list -clustername galaxy -spname atlantic
```

Notes:

- Alternatively swap “-lu LU_name(s)” for “-luudid Hexadecimal”
- LUs means a space separated list disk names

Snapshot – create and list

Create

```
$ snapshot -create diamond5s.snap -clustername galaxy  
-spname atlantic -lu vdisk_diamond5a
```

List

```
$ snapshot -list -clustername galaxy -spname atlantic  
Lu Name      Size(mb)    ProvisionType  Lu Udid  
vdisk_diamond5a 16384      THIN           b3f3a . . .  
Snapshot  
diamond5s.snap
```

Also snap shots appear in the lssp output

```
$ lssp -clustername galaxy -sp atlantic -bd  
Lu Name      Size(mb)    ProvisionType  Lu Udid  
vdisk_diamond5a 16384      THIN           b3f3a . . .  
Snapshot  
diamond5s.snap  
  
vdisk_diamond6a 16384      THIN           4c9e9 . . .
```

Snapshot – delete or rollback

When sure you never want to rollback
Delete original & continue on the current blocks

```
$ snapshot -clustername galaxy -delete diamond5t.snap
  -spname atlantic -lu vdisk_diamond5a
```

Rollback to a snapshot
Stop the virtual machine/LPAR then


```
$ snapshot -clustername galaxy -rollback diamond5t.snap
  -spname atlantic -lu vdisk_diamond5a
```

Warning:
You loose any updates you made since that snapshot
Any snapshots since that snapshot are removed

Snapshot Model




Original Set



Snapshot Model



Original Set becomes the Snapshot




New working set

Creating a snapshot only involved copying the meta data
i.e. list of the blocks within the LU (not the block themselves)

Snapshot + Update





Original Set becomes the Snapshot



New working set

- 1 Client VM update to block 7
- 2 VIOS allocates a new block
- 3 Copies original 64MB
- 4 VIOS completes the write





Storage Management

- Reminder currently,
 - One pool of large LUNs
 - Pool can be on a mix of brands or generations of disk sub-systems
 - 64 MB chunks are spread as evenly as possible across LUNs

84 of 44 in 92

IBM © 2013 IBM 84 Shared Storage Pool 3

Live Storage Mobility

Server Admin would like to :

- A. Replace a faulty LUN
- B. Move all blocks off one disk subsystem (retiring a disk subsystem)
- C. Recover from repository failure
- D. Select which disk subsystems a particular VM uses
- E. Ensure mirrors are on different subsystems (even different sites)

- A and B → via replace physical disk

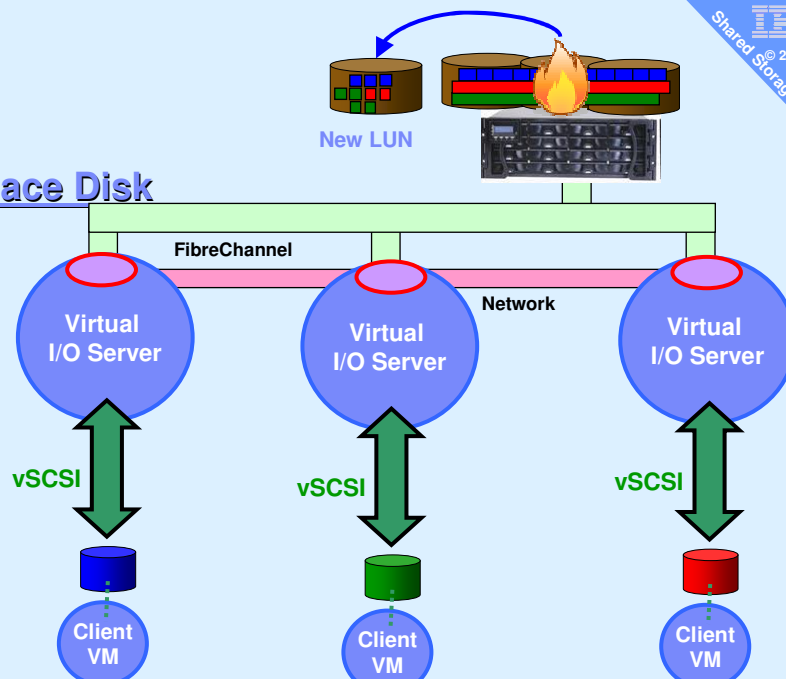
- `chsp -replace -clustername galaxy -sp atlantic -oldpv hdisk4 -newpv hdisk24`

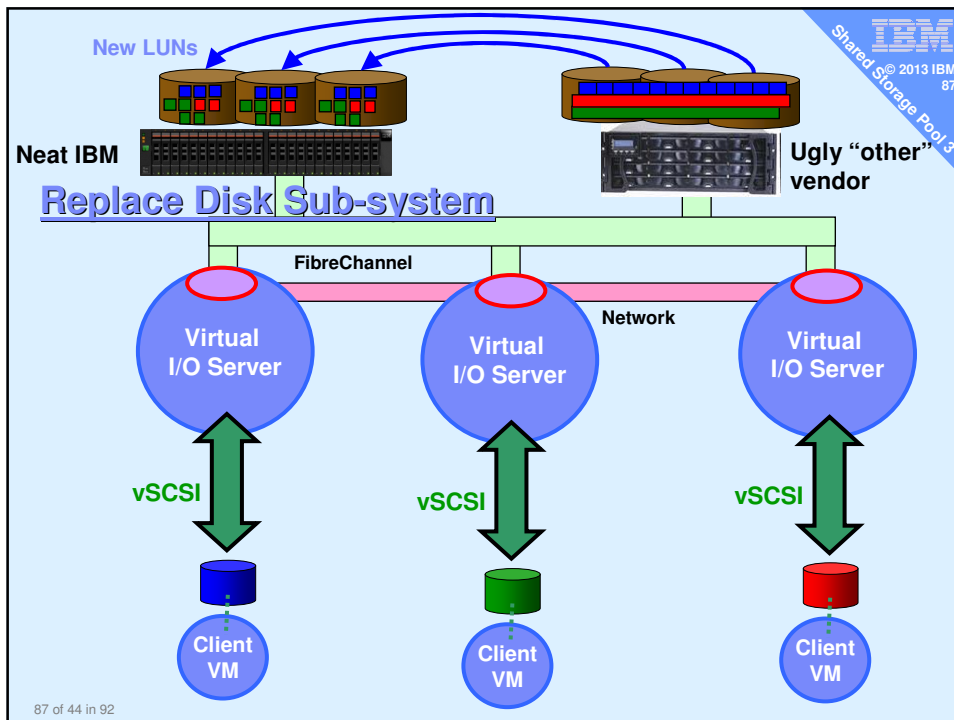
- C → see `viosbr` command (later)

- D and E → in a later SSP release

- Could use `SVC` now for lower level mirror (E)

Replace Disk





What if you loose the VIOS?

- Updated **viosbr** supports backup / restore of SSP config
 - **Warning: this saves the config but not the data**
- Backup – will perform regular backups for you


```
viosbr -backup -clustername Name -file File \
[-frequency daily|weekly|monthly [-numfiles fileCount]]
```
- View


```
viosbr -view -file File -clustername Name [-type devType][-detail | -mapping]
```
- Restore


```
viosbr -restore -clustername N -file F -subfile NodeFile [-validate | -inter | -force][-type devType]
viosbr -restore -clustername N -file F -repopvs disks [-validate | -inter | -force][-type devType][-currentdb]
viosbr -recoverdb -clustername N [ -file F ]
viosbr -migrate -file F
```
- Can recover from
 1. Repository Disk is corrupted (see -repopvs)
 2. One SSP VIOS is reinstalled
 3. SSP Database is corrupted
 4. Restore to old configuration on the VIOS node
 - Changes done to SSP mappings on the node after a backup

88 of 44 in 92

SSP2 command cheat sheet **Reference Only**

```
1. chdev -dev <device name> -attr reserve_policy=no_reserve
2. cluster -create -clustername galaxy -repopvs hdisk2
   -spname atlantic -sppvs hdisk3 hdisk5 -hostname bluevios1.ibm.com
3. cluster -list
4. cluster -status -clustername galaxy
5. cluster -addnode -clustername galaxy -hostname redvios1.ibm.com
6. cluster -rmnode [-f] -clustername galaxy -hostname redvios1.ibm.com
7. cluster -delete -clustername galaxy
8. lscluster -s or -d or -c or -i = CAA commands
9. chsp -add -clustername galaxy -sp atlantic hdisk8 hdisk9
10. chsp -replace -clustername galaxy -sp atlantic -oldpv hdisk4 -newpv hdisk24
11. mkbdsp -clustername galaxy -sp atlantic 16G
   -bd vdisk_red6a -vadapter vhost2 [-thick]
12. rmbdsp -clustername galaxy -sp atlantic -bd vdisk_red6a
13. lssp -clustername galaxy -sp atlantic -bd
14. lssp -clustername galaxy
15. alert -set -clustername galaxy -spname atlantic -type threshold -value 10
16. alert -list -clustername galaxy -spname atlantic
17. errlog -ls
18. snapshot -create name -clustername galaxy -spname atlantic -lu LU42
19. snapshot -delete name -clustername galaxy -spname atlantic -lu LU42
20. snapshot -rollback name -clustername galaxy -spname atlantic -lu LU42
21. snapshot -list -clustername galaxy -spname atlantic
22. viosbr -backup -clustername galaxy -file Daily -frequency daily -numfiles 10
23. viosbr -view -file File -clustername Name ...
24. viosbr -restore -clustername Name ...
```

89 of 44 in 92