





# VIOS 2.2.4 Shared Storage Pools Phase 5 with Tiers and loads more

Power VM



© 2015 IBM Corporation

Nigel Griffiths  
nag@uk.ibm.com





## Marketing: VIOS Shared Storage Pools

1. Enormous reduction in storage man-power
2. Independence from underlying SAN technology & team!
3. Sub-second disk space allocate & connect
  - lu command: create, map, unmap, remove
  - snapshot: create/delete/rollback
4. Autonomic disk mirrors & resilver with zero VM effort
5. Live Partition Mobility ready by default
6. Simple Pool management:
  - pv & failgrp, lssp, alert, VIOS logs
7. DR capability to rebuild a VM quickly on a Remote Pool Copy
8. HMC GUI for fast SSP disk setup across dual VIOS
  - No more: VIOS slot numbers, Cnn or vhosts

IBM  
© 2015 IBM  
2  
Shared Storage Pool 5

## YouTube Videos



  
 © 2015 IBM  
 Shared Storage Pool 5

1.	<a href="#">Shared Storage Pool (SSP) Intro</a>	17 mins
2.	<a href="#">Shared Storage Pools (SSP2) Getting Started</a>	10 mins
3.	<a href="#">Shared Storage Pools (SSP2) Thin Provisioning Alerts</a>	17 mins
4.	<a href="#">Shared Storage Pools (SSP3) New Features</a>	24 mins
5.	<a href="#">Looking Around a Shared Storage Pool SSP3</a>	15 mins
6.	<a href="#">Live Partition Mobility (LPM) with Shared Storage Pool SSP3</a>	7 mins
7.	<a href="#">SSP3 Recover a Crashed Machine's LPAR to Another Machine</a>	25 mins
8.	<a href="#">Migrating to Shared Storage Pool (SSP3) &amp; then LPM</a>	18 mins
9.	<a href="#">Shared Storage Pool 4 (SSP4) Concepts</a>	28 mins
10.	<a href="#">Shared Storage Pools 4 (SSP4) Hands On</a>	19 mins
11.	<a href="#">PowerVC 1.2.1 with Shared Storage Pools</a>	20 mins
12.	<a href="#">Shared Storage Pool in 3 Commands in 3 Minutes</a>	8 mins
13.	<a href="#">Shared Storage Pools Repository is bullet proof</a>	13 mins
14.	<a href="#">Shared Storage Pool Remote Pool Copy Activation for Disaster</a>	22 mins
15.	PowerVM VUG <a href="#">33 VIOS Shared Storage Pools Phase4</a>	100 mins

11,500 views in total so far (Q4 2015)

- YouTube search: Shared Storage Pools Nigel Griffiths

## Recent new information: AIXpert Blog

  
 © 2015 IBM  
 Shared Storage Pool 5

<https://www.ibm.com/developerworks/community/blogs/aixpert/>

- Look for entries with these titles:

SSP4 Best Practice & FAQ

– 35 recommendations + 20 questions

SSP Hands-on Fun with LU by Example

- Rename a LU (offline)
- Backup & restore LU + Backup a snapshot
- Slim down a Thin LU (offline)
- Move a LU between SSPs
- Check if a LU is mapped across whole SSP cluster
- **No Testing in Production (TIP) please!!!**



**New information: AIXpert Blog**

<https://www.ibm.com/developerworks/community/blogs/aixpert/>

- Look for entries with these titles:

- SSP4 a better lu -list command**
  - Script to make the LUs in order + better format
- SSP4 Pool Expansion**
  - Grow the LUNs and the pool grows
- SSP4 Cheat Sheet**
  - Learn the commands by simple examples
- How many SSP in the world?**
  - SSP = no charge option of VIOS so we don't know
  - I guess: SSP=1000's, 10's of TB & loads in production



IBM  
© 2015 IBM  
5  
Shared Storage Pool 5

**Reminder**

IBM  
© 2015 IBM  
6  
Shared Storage Pool 5

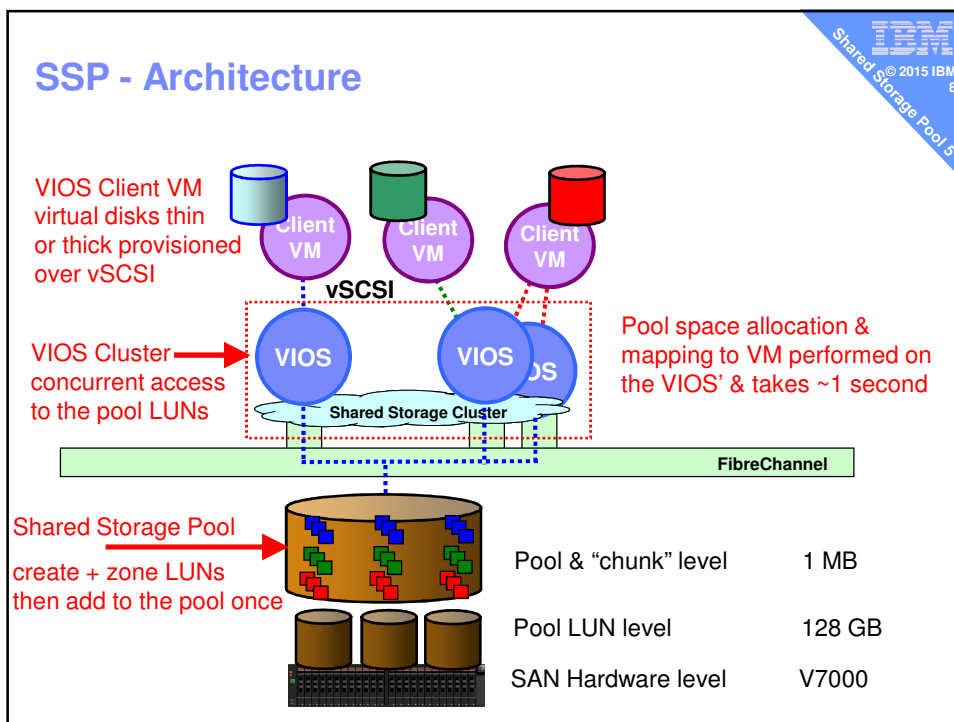
**IBM**  
© 2015 IBM  
Shared Storage Pool 5

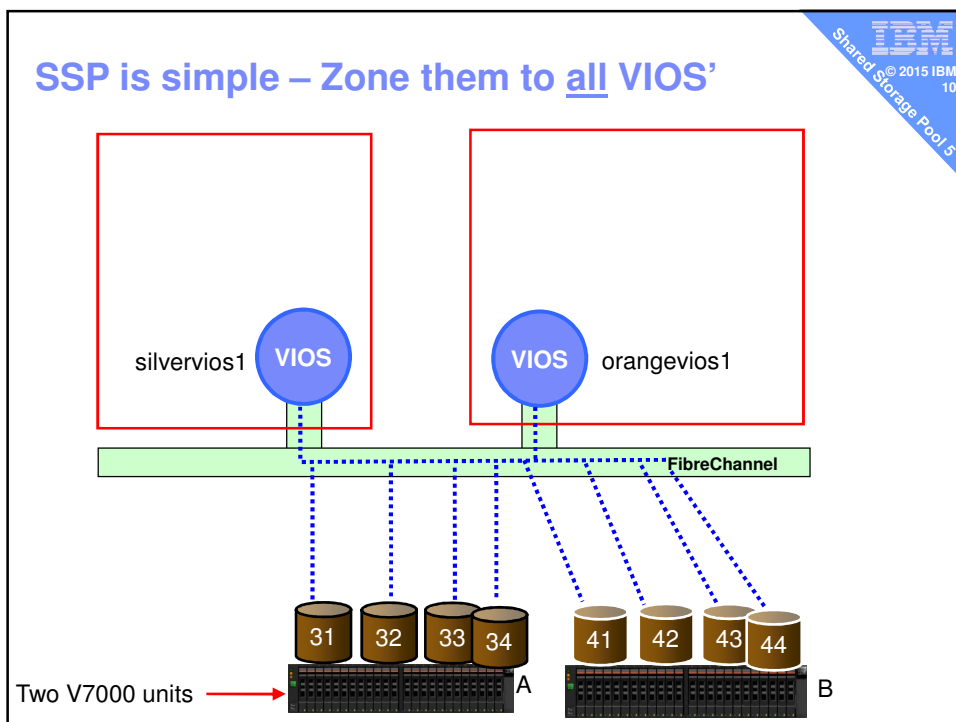
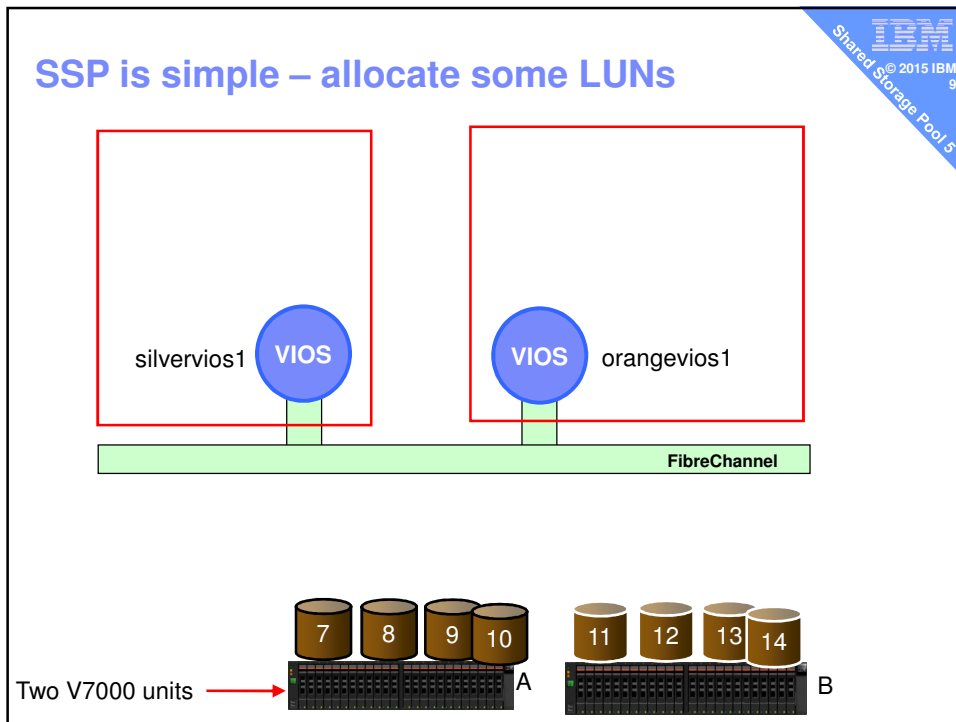
### Readme from FixCentral for each VIOS release

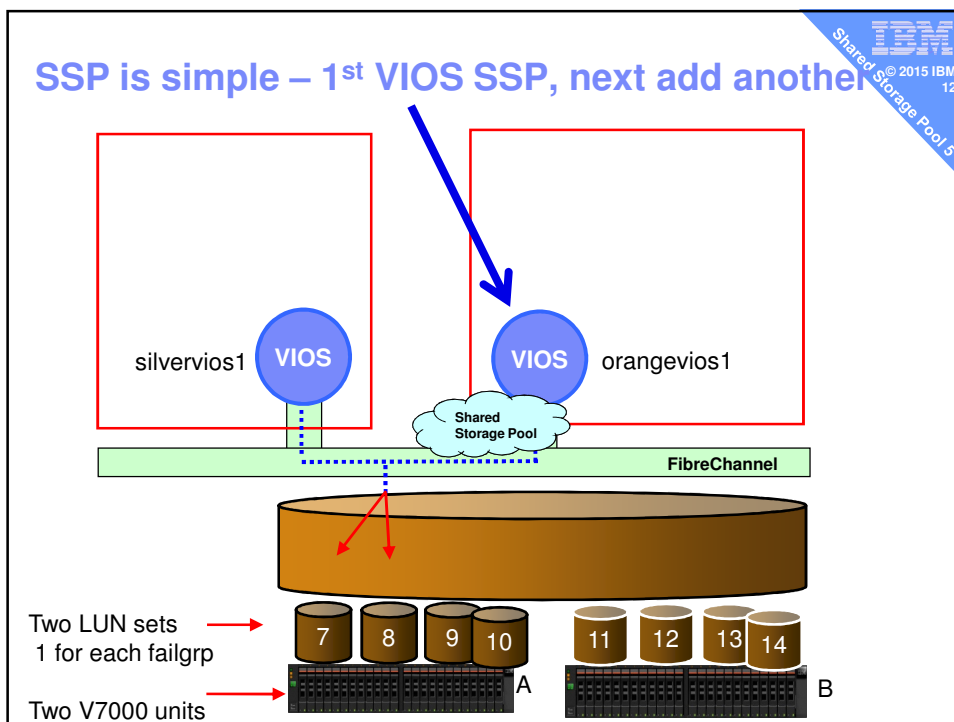
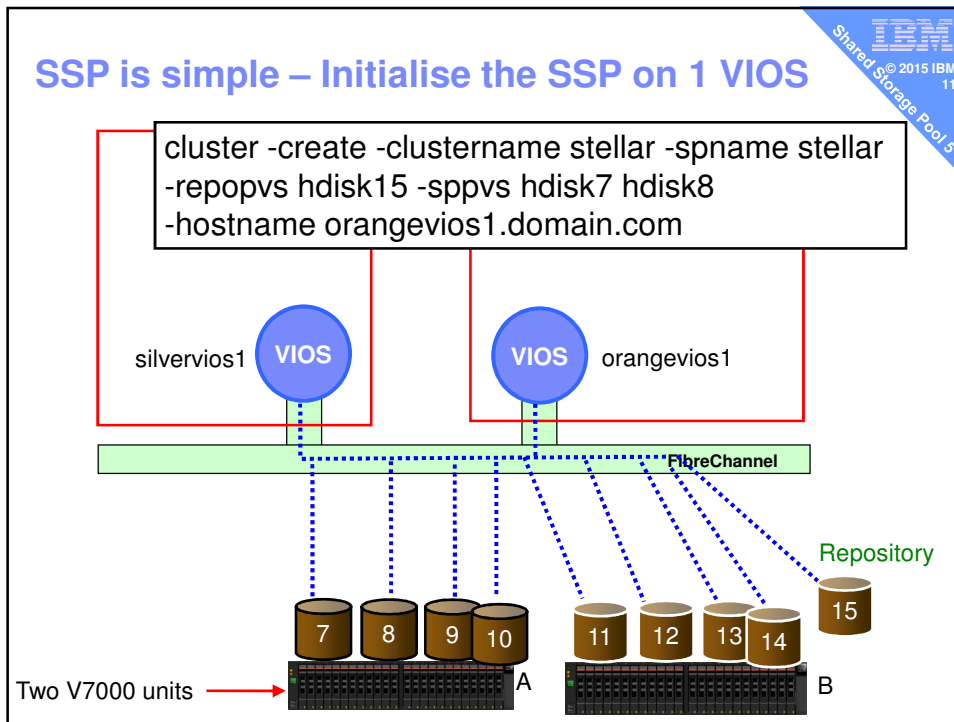
Feature	Min	Max
▪ Number of VIOS Nodes in Cluster	1	<b>16</b>
▪ Number of LUNs in Pool	1	1024
▪ Number of Virtual Disks (LUs) Pool	1	8192
▪ Number of Client LPARs per VIOS node	1	200
▪ Each LUN in Pool size	10GB	16TB
▪ Total Pool size	10GB	<b>512TB</b>
▪ Virtual Disk (LU) size from the Pool	1GB	4TB
▪ Number of Repository Disks	1	<b>1</b>
▪ Capacity of Repository Disk	512MB	1016GB

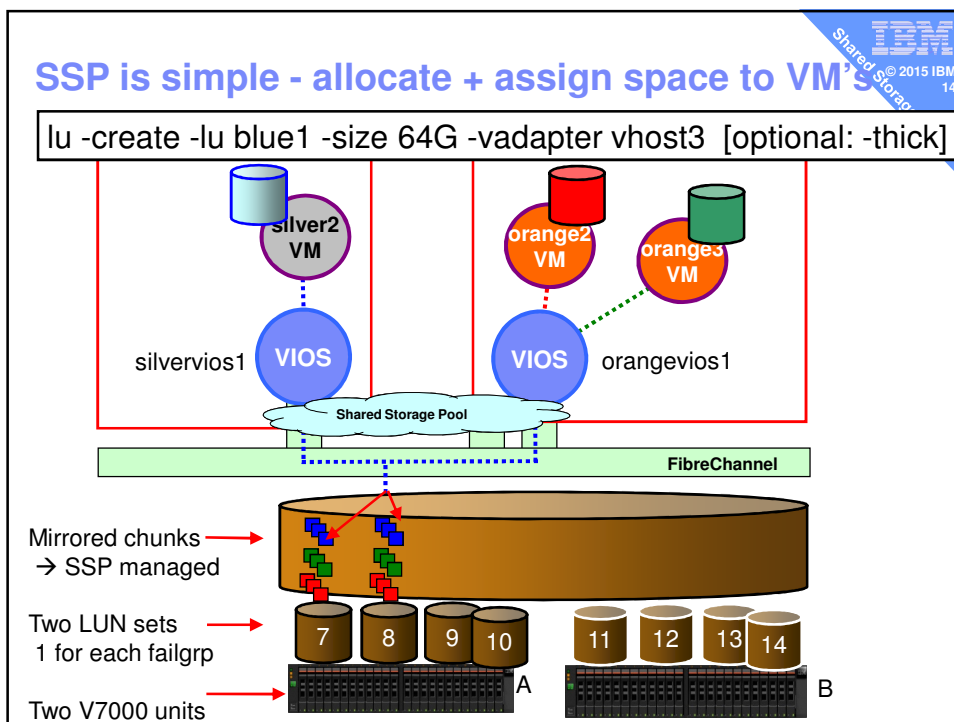
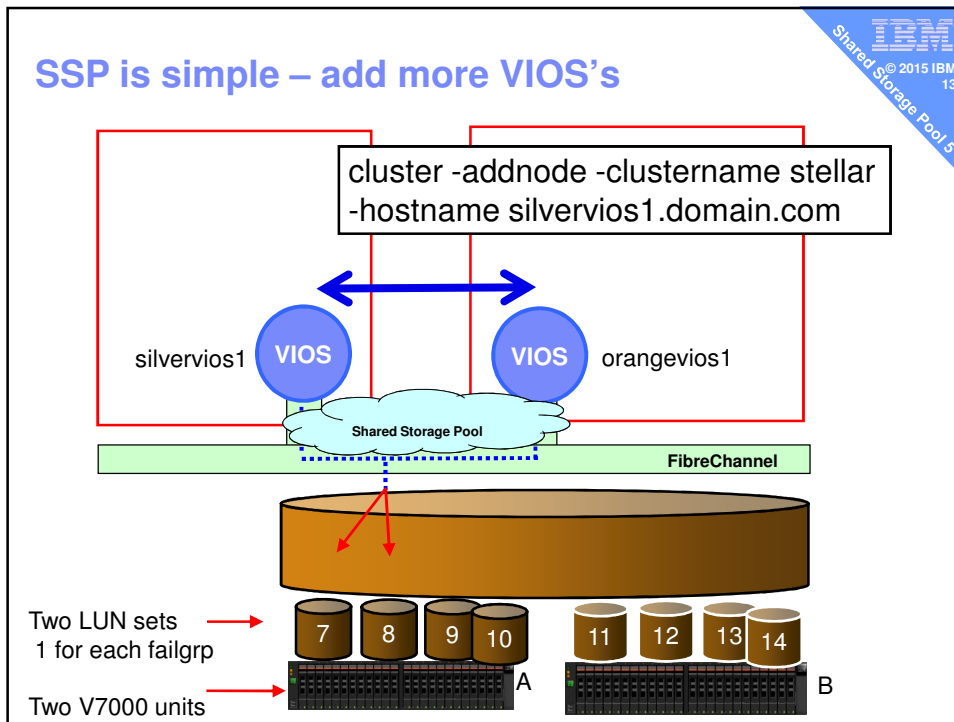
  

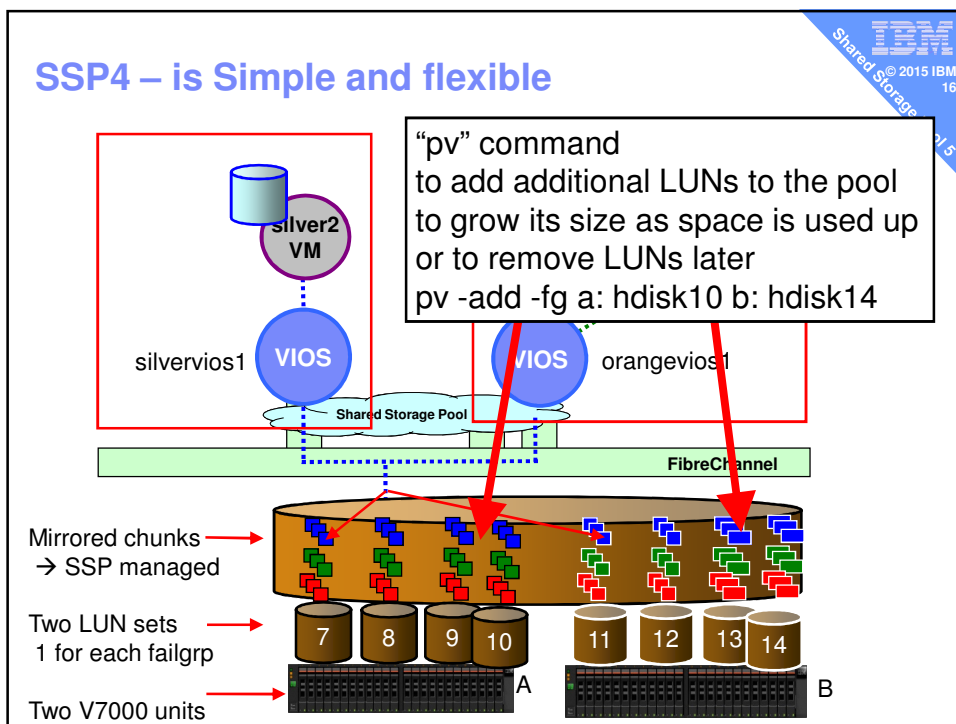
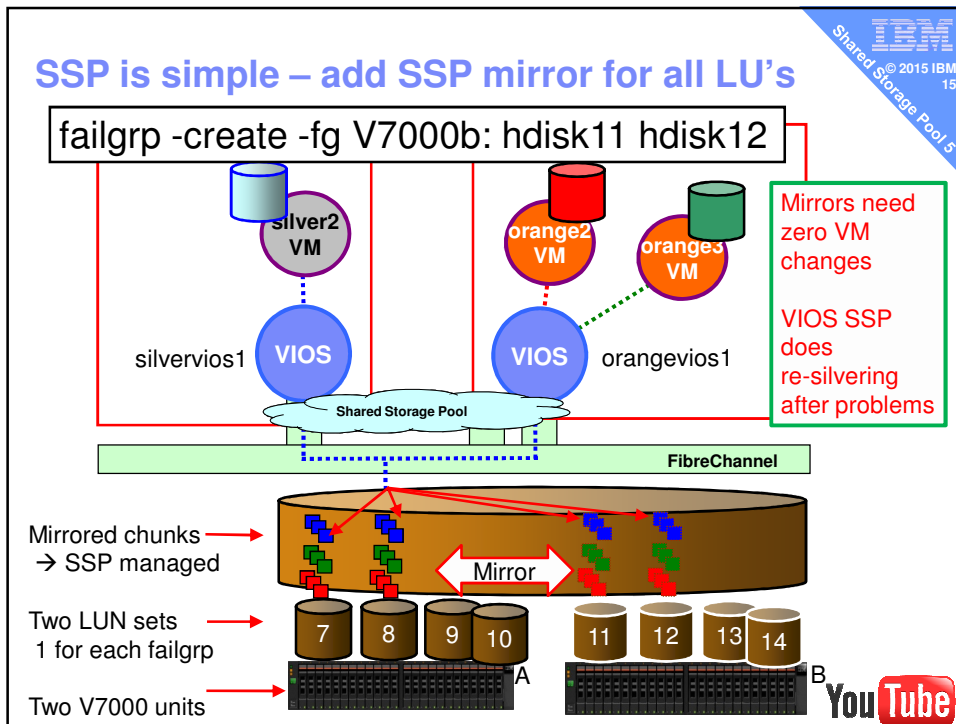
- Nigel's Recommendation
  - 8 LUN minimum
  - LUNs of 128 GB in the pool up to 8 TB
  - Larger LUN for large pools
  - Repository LUN size 1 GB & spare repository LUN of 2 GB



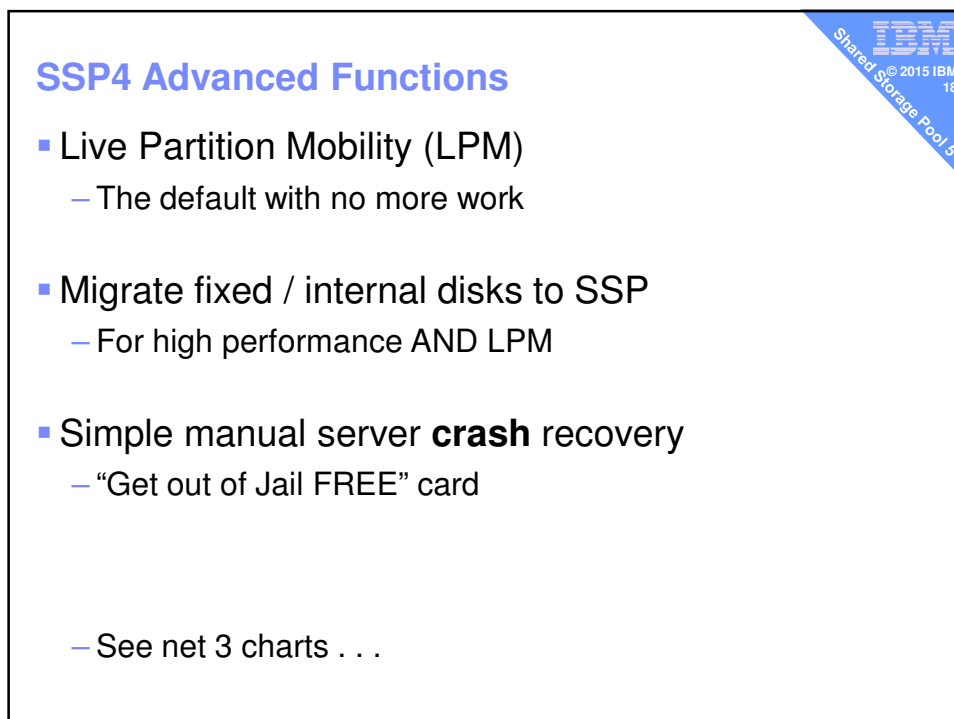
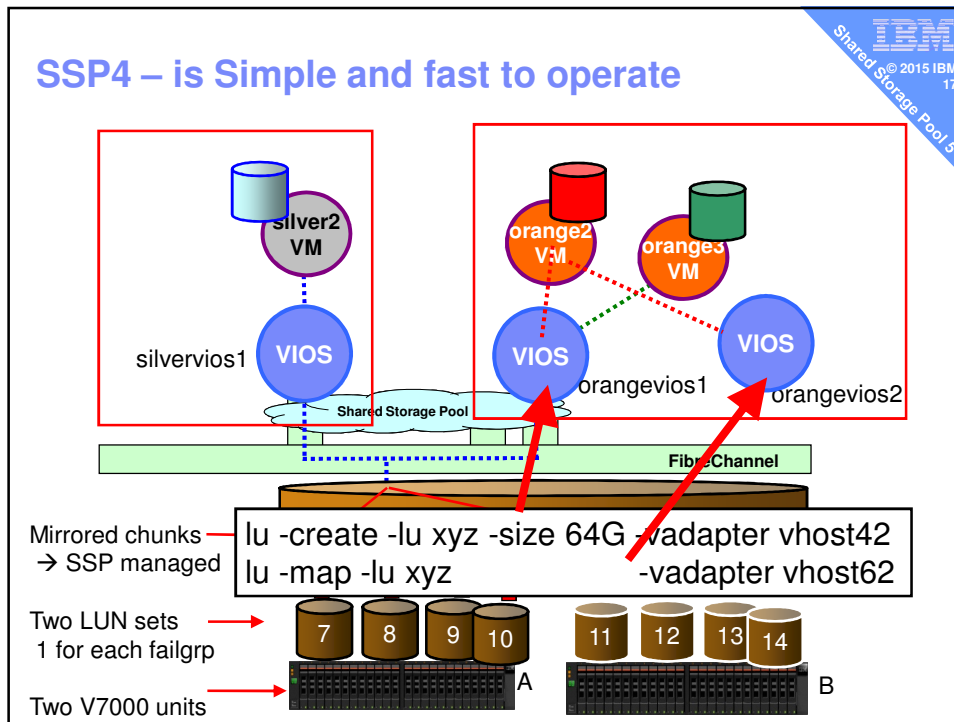












**4**

### Advanced Functions: Live Partition Mobility

IBM © 2015 IBM 19  
Shared Storage Pool 5

- SSP VIOS's already have the LUNs online  
→ no SAN zoning issues
- Provided you have Virtual Ethernet & no physical adapters  
→ you are LPM ready  
→ default is LPM ready  
→ no additional work
- Assuming your machines have PowerVM Enterprise = LPM

The diagram illustrates Live Partition Mobility (LPM). At the top, three Client VMs (purple circles) are shown. Below them are two VIOS instances (blue circles) connected via vSCSI. These VIOS instances are connected to a Shared Storage Pool (cloud shape). The storage pool is connected to a physical storage array (cylinder) with multiple disks. A red arrow points from one Client VM to another, indicating migration. A YouTube logo is in the bottom right corner.

**5**

### Advanced Functions: Migrating to SSP

IBM © 2015 IBM 20  
Shared Storage Pool 5

Got old local disks, VIOS LV or hdisk via vSCSI, NPIV but want to use SSP !!!

Actions:

- Add VIOS to SSP
- Add LU disks
- With AIX
  - migratepv live to SSP
  - bosboot
  - bootlist
  - Remove old disks/adapters
  - and you are LPM ready

The diagram illustrates migrating to a Shared Storage Pool (SSP). It shows three Client VMs (purple circles) connected to three VIOS instances (blue circles) via vSCSI. These VIOS instances are connected to a Shared Storage Pool (cloud shape). The storage pool is connected to a physical storage array (cylinder) with multiple disks. A red 'X' is placed over a disk icon, indicating the removal of old disks. A YouTube logo is in the bottom right corner.

**6**

## Advanced Functions: Box Crash Recovery

- Total box lost!
  - SSP will disks survive

Actions:

- Make a new LPAR
- Map in the SSP LU
- Connect to right network
- Set the bootlist
- Reboot
- and you are running again in, say, 2 minutes!

IBM  
© 2015 IBM  
21  
Shared Storage Pool 5

YouTube

## New Stuff

To install Shared Storage Pool phase 5

- Just upgrade/install VIOS 2.2.4 (or later)
- Available 5<sup>th</sup> December 2015

IBM  
© 2015 IBM  
22  
Shared Storage Pool 5

## VIOS 2.2.4

1. SSP LU Resize (grow = saves admin time)
2. Command: lu -list in alphabetical order
3. Removing pointless -clustername -sname in commands
4. SSP Tiers (multiple pools only better)
  - 10 tiers (think grouping not levels)
  - Fast, medium, slow or IBM, HDS, EMC or prod, in-house, test
5. SSP mirrored now at tier level (was whole SSP)
6. Move a LU between tiers
7. HMC GUI extended for tiers
8. SSP Tier advanced features

© Copyright IBM Corporation 2015

## 1 SSP LU Resize

Why resize, when you can add a new LU virtual Disk?

- Answer: new LU = manual spread data across disks at OS
- New option to the lu command: -resize
- Example: `lu -resize -lu myLU -size 35G` [G = GB or M = MB]
- LU shrink not possible – give you a polite error
- No “add a bit” option (+4GB) → you state the new total size
- Online = Live with the VM using the LU
  - Actually, the VM will not notice more blocks at the end

© Copyright IBM Corporation 2015

## 1 SSP LU Resize on the VIOS

IBM  
© 2015 IBM  
25  
Shared Storage Pool 5

```
$ lu -list -attr lu_name=vm96boot
POOL_NAME: spiral
TIER_NAME: prod
LU_NAME          SIZE (MB)  UNUSED (MB)  UDID
vm96boot         38912      36353        4634dca8b41654ddb39893177d61060e
4634dca8b41654ddb39893177d61060e
```

Original Size

```
$ lu -resize -lu vm96boot -size 40G
Logical unit vm96boot with udid '4634dca8b41654ddb39893177d61060e' has
been successfully changed.
```

```
$ lu -list -attr lu_name=vm96boot
POOL_NAME: spiral
TIER_NAME: prod
LU_NAME          SIZE (MB)  UNUSED (MB)  UDID
vm96boot         40960      38401        4634dca8b41654ddb39893177d61060e
4634dca8b41654ddb39893177d61060e
```

New Size

© Copyright IBM Corporation 2015

## 1 SSP LU Resize on the VM

IBM  
© 2015 IBM  
26  
Shared Storage Pool 5

– OS level details

- AIX: `chvg -g rootvg` → AIX then finds the larger disk space
- AIX client has a VG PP size = minimum you need to grow
  - 32 GB rootvg has default PP size of 64 MB
- IBM i: Does not support LUN/LU resize → actually “dangerous!”
  - For IBM i just give it a new LU – it knows what to do.
- Linux: depends on the volume manager & filesystem in use
  - Good luck!

© Copyright IBM Corporation 2015

## 1 SSP LU Resize on the AIX VM

IBM  
© 2015 IBM  
27  
Shared Storage Pool 5

```
# lsvg rootvg | grep "PP SIZE"
VG STATE:          active          PP SIZE:          64 megabyte(s)
# lsvg rootvg | grep "TOTAL PP"
VG PERMISSION:    read/write    TOTAL PPs:       607 (38848 megabytes)
```

Resize on the VIOS here

Have to grow at least this size & preferably a multiple Or grow by GB's

```
# chvg -g rootvg
```

Pre-resize

```
# lsvg rootvg | grep "TOTAL PP"
VG PERMISSION:    read/write    TOTAL PPs:       639 (40896 megabytes)
#
```

© Copyright IBM Corporation 2015

## 2 lu -list in alphabetical order

IBM  
© 2015 IBM  
28  
Shared Storage Pool 5

- Wow, about time, right!
- I pointed out the random order in the beta testing ☺
- Developers said "No way. Oh cumbs! We will get that fixed."
- And they have

© Copyright IBM Corporation 2015

## 2 lu -list

### Spot the THREE user hostile features? IMHO

```

$ lu -list
POOL_NAME: spiral
TIER_NAME: test
LU_NAME          SIZE (MB)  UNUSED (MB)  UDID
testa            32768      32770        0491ba1cd2bb7a41040f307689636d21
SNAPSHOTS
2015-09-30T11:26:32
testb            40960      40962        5c2ebbf20e7e4edac8d166cb6bec33c1
SNAPSHOTS
before-upgrade
after-upgrade
v234567890123456789012~ 8192      8192         ebffb84803e5ced5401ebf1ed7d6c2fc
vm97boot         38912      36352        a1ccdea9dac4ed18b4ea546de9a69bcc
vm97data         8256       8233         bb033473f8ec78752550ba0fbe940f27

```

© Copyright IBM Corporation 2015


  
 Shared Storage Pool 5
   
 © 2015 IBM
   
 29

## 2 lu -list

```

$ lu -list
POOL_NAME: spiral
TIER_NAME: test
LU_NAME          SIZE (MB)  UNUSED (MB)  UDID
testa            32768      32770        0491ba1cd2bb7a41040f307689636d21
SNAPSHOTS
2015-09-30T11:26:32
testb            40960      40962        5c2ebbf20e7e4edac8d166cb6bec33c1
SNAPSHOTS
before-upgrade
after-upgrade
v234567890123456789012~ 8192      8192         ebffb84803e5ced5401ebf1ed7d6c2fc
vm97boot         38912      36352        a1ccdea9dac4ed18b4ea546de9a69bcc
vm97data         8256       8233         bb033473f8ec78752550ba0fbe940f27

```

(1) On what planet is that useful in the default output?  
32 hexadecimal digits ↓

(2) Mixed in snapshot names = confusing

(3) Truncated LU name at 22 letters

© Copyright IBM Corporation 2015


  
 Shared Storage Pool 5
   
 © 2015 IBM
   
 30

## 2 lu -list

- Trouble is, everyone has there own "perfect layout"
- Fortunately, the SSP designers are very clever
- The lu options allow you any format

```
$ lu -list -field LU_SIZE LU_NAME -fmt :
32768:testa
40960:testb
38912:vm97boot
8256:vm97data
8192:testc
8192:v23456789012345678901234567890
38912:vm96boot
8256:vm96data
```

For a full list of field names use:  
lu -list -verbose

- OK not pretty but nothing awk can't sort out
- In Tier then alphabetical order = actually makes sense

© Copyright IBM Corporation 2015

## 2 lu -list

```
echo "SizeMB UsedMB Used% Type Tier Name"
/usr/ios/cli/ioscli lu -list -fmt : \
-field LU_SIZE LU_USED_SPACE LU_USED_PERCENT \
      LU_PROVISION_TYPE TIER_NAME LU_NAME \
| awk -F: '{ printf "%6d %6d %4d%% %5s %7s %s\n", $1, $2, $3, $4, $5, $6 }'
```

SizeMB	UsedMB	Used%	Type	Tier	Name
32768	0	0%	THIN	SYSTEM	testa
40960	0	0%	THIN	SYSTEM	testb
38912	2562	6%	THIN	SYSTEM	vm97boot
8256	23	0%	THIN	SYSTEM	vm97data
8192	8192	100%	THICK	prod	testc
8192	0	0%	THIN	prod	v23456789012345678901234567890
38912	2561	6%	THIN	prod	vm96boot
8256	26	0%	THIN	prod	vm96data

© Copyright IBM Corporation 2015



## 2 nlu – download AIXpert blog

```

help()
{
    echo $0 "Nigel's lu command with improved layout and column ordering"
    echo $0 "[--sizemb | --usedmb | --used | --type | --tier | --name (default)]"
    exit 0
}

if [[ $(whoami) == "padmin" ]]
then
    command=$(whence $0)
    # echo DEBUG I am padmin so restart $command again as the root user
    echo "$command" $1 | oem_setup_env
else
    # echo DEBUG now I am root
    # lowercase the parameter with tr to avoid input case errors
    case `echo $1 | tr "[A-Z]" "[a-z]"` in
    1 | --sizemb)          COLUMN="-k 1" ;;
    2 | --usedmb)         COLUMN="-k 2" ;;
    3 | --used | --used%) COLUMN="-k 3" ;;
    4 | --type)           COLUMN="-k 4" ;;
    5 | --tier)          COLUMN="-k 5" ;;
    6 | --name)          COLUMN="-k 6" ;;
    ? | -?) help ;;
    *)                   COLUMN="-k 6" ;;
    esac

    echo " SizeMB  UsedMB  Used%  Type  Tier Name"

    /usr/ios/cli/ioscli \
    lu -list -field LU_SIZE LU_USED_SPACE LU_USED_PERCENT LU_PROVISION_TYPE TIER_NAME LU_NAME -fmt : \
    | awk -F: '{ printf "%7d %7d %4d%% %5s %5s %s\n", $1, $2, $3, $4, $5, $6 }' \
    | sort $COLUMN
fi
exit 0

```

© Copyright IBM Corporation 2015

## 2 nlu with sort by field name

```

$ nlu --usedmb

```

SizeMB	UsedMB	Used%	Type	Tier	Name
8192	0	0%	THIN	prod	v23456789012345678901234567890
32768	0	0%	THIN	test	testa
40960	0	0%	THIN	test	testb
8256	23	0%	THIN	test	vm97data
8256	26	0%	THIN	prod	vm96data
38912	2573	6%	THIN	prod	vm97boot
40960	2579	6%	THIN	test	vm96boot
39936	39936	100%	THICK	prod	testc

```

$

```

© Copyright IBM Corporation 2015

### 3 Remove -clustername -spname

- Small item – more of a clean up change
- Only one cluster per VIOS AND one pool per VIOS
  - The VIOS knows the only one possible name
  - QED: don't make the user type it all day
- Removed (optional) from most commands including:
  - lu, failgrp, tier, pv, alert, snapshot
- Nigel's recommendation
  - KISS: make the cluster name & pool name the SAME

© Copyright IBM Corporation 2015

### 3 Remove -clustername -spname

Still needed for

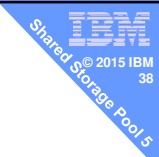
- `cluster -create -cluster XXX -sp YYY`
  - to set the names later seen in command output
- NOTE: it was ~~-spname~~ and now is `-sp`
- Older "list storage pool" command `lssp` has it
  - `lssp -clustername XXX -sp YYY -bd`
- → Same command used for local pool disks (rootvg LV's)
  - Most people use the `lu` command instead via a script
- Note: was always `-sp`

© Copyright IBM Corporation 2015



# 4 Tiers

© Copyright IBM Corporation 2015



## SSP5 – Multiple Pools verses Tiers

- Regular question; I need **multiple pools**, WHEN?
  - With the current version just one pool per VIOS and “hose it all about” policy – I can’t separate workloads
- Different reasons:
  1. Two pools for high speed IBM SAN disks & older slower EMC ones ☺
  2. Different pools for FlashSystem 9000 (prod) & V7000 (test/dev)
  3. Policy to separate production data & “other” workload disks
  4. Have local speedy & remote FC “dark fibre” storage
    - need to set which VM gets which type
- By the way tiers can do all this

## SSP5 – Multiple Pools verses Tiers


- Regular question; I need **multiple pools**, WHEN?
  - With the current version just one pool per VIOS and “hose it all about” policy – I can’t separate workloads
- **But multiple pools has inherit issues:**
  - Live moving a LU between pools is impossible to do safely
  - What if you fill a pool and other is 95% empty
  - Implement in the slow pool but urgently now need fast disks
  - What if VIO Servers have different mixture of pools
    - LPM could be messy & Zone complexity
    - SSP is meant to reduce complexity and save people time

## SSP5 – Multiple Pools verses Tiers

### Conclusion:


- What you wanted multiple pool for, can be done by tiers
- Tiers can do more than multiple pools
- Tiers keep everything simple & fast to operate
- Regular expression: 1,\$s/multiple pools/Tiers/g
  - Over a beer, we could argue they are the same thing!
- Tiers Win! – Hurray!

**SSP5 – Tiers**

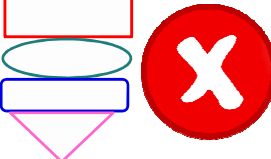
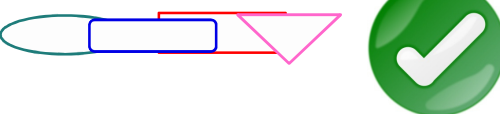


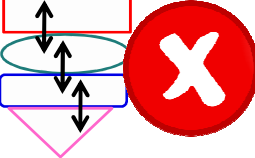
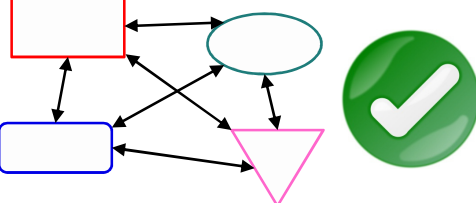
- Ten tiers inside the single shared storage pool
- Tier made up of a set of LUN – giving you data separation
- LUN can be from different disk units
- Use Tiers any way you like:
  1. Speed: FAST, Medium, slow
  2. Vendor: IBM, HDS, EMC
  3. Importance: Critical, Prod, in-house, test, dev
  4. Isolation: customerA and customerB
  5. Isolation: marketing, sales, support and proper techies
  6. Location: Local, computer-room-C, across-campus
  7. Functionality: RDBMS, web-service, archive, video-collection
- Can live / dynamically move an LU between tiers
  - Due to one common set of meta data

**SSP5 – Tiers**




- Misconception 1 Tiers do not imply layers = no top or bottom
 



- Misconception 2 Move from a tier to any other tier
 



- Misconception 3 There is no special tier\*
 

Well there is a tier which has the meta data (SYSTEM) & the Default tier (changeable)

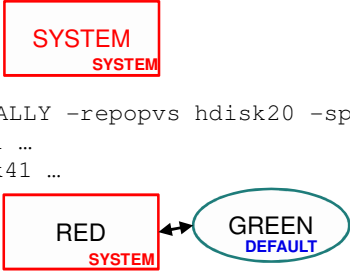

  
 © 2015 IBM
   
 43
   
 Shared Storage Pool 5


### SSP5 – Tiers creation

- A Tier is just a collection of LUNs
- An SSP with one tier works exactly like SSP4
  - The only is called SYSTEM includes the meta data and is the default
- SSP create single first tier
 

```
cluster -create -clustername FRED -repopvs hdisk20 -sp FRED
            -sppvs hdisk31 hdisk32 ...
```
- SSP create with two tier “starter pack”
 

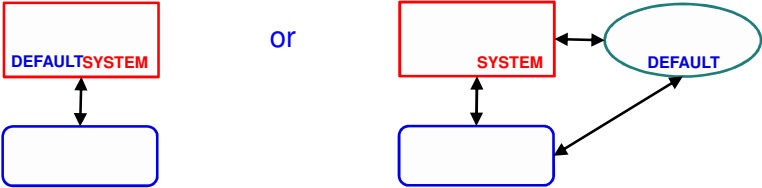
```
cluster -create -clustername SALLY -repopvs hdisk20 -sp SALLY
            -systier RED: hdisk30 hdisk31 ...
            -usrtier GREEN: hdisk40 hdisk41 ...
```
- then you can add more tiers




  
 © 2015 IBM
   
 44
   
 Shared Storage Pool 5

### SSP5 – Later Tier creation

- Add one Tier at a time
- tier -create -tier BLUE: hdisk50 hdisk51



or

- Other Options follow SSP conventions
  - tier -remove -tier NAME
  - tier -modify -tier NAME -attr ATTR=VALUE
  - tier -list [-fmt : | -field | verbose] ...

**SSP5 – Tiers LU allocation**

IBM © 2015 IBM 45  
Shared Storage Pool 5

- One Tier is marked as "Default"
- If you make a LU as before it goes in the Default Tier
  - \$ lu -create -lu vm22boot -size 64G
- Add options to allocate from your choice of Tier
  - \$ lu -create -lu vm23boot \ -size 64G -tier RED
- A LU is only in one Tier


- IMHO: Makes sense for the LU's of one VM to be in the same Tier

**5 SSP5 – Tiers Mirrors**

IBM © 2015 IBM 46  
Shared Storage Pool 5

- Mirrors (a second failgrp) are now at the Tier level
- Same failgrp command but with a -tier parameter
  - failgrp -create -tier SALLY -fg SALLY2: hdisk60 hdisk61 ...
- As before you need a 2<sup>nd</sup> set of disks – on a different disk unit
- So you can have mixed mirrored tiers & unmirrored tiers
  - IMHO that would be unusual
- Have to mirror each tier, of course

Tier for:  
"crash & burn" VMs  
Or  
backup / scratch space



© 2015 IBM  
47  
Shared Storage Pool 5

### SSP5 – Tiers Attributes

```

$ tier -list
POOL_NAME: testsp
TIER_NAME  SIZE (MB)  FREE_SPACE (MB)
MIRROR_STATE
SYSTEM      10112      8000      NOT_MIRRORED
mytier      10110      8000      SYNCED
    
```

**SYSTEM tier no failgrp mirror**  
**mytier has a failgrp mirror & mirrors in sync**

```

$ tier -list -verbose
POOL_NAME: testsp
TIER_NAME: SYSTEM
TIER_TYPE: COMINGLED
TIER_DEFAULT: NO
TIER_SIZE (MB): 10112
FREE_SPACE (MB): 8000
OVERCOMMIT_SIZE (MB): 0
TOTAL_LUS : 5
TOTAL_LU_SIZE: 2112
FG_COUNT : 1
MIRROR_STATE: NOT_MIRRORED
ERASURE_CODE: NONE


POOL_NAME: testsp
TIER_NAME: mytier
TIER_TYPE: USER
TIER_DEFAULT: YES
TIER_SIZE : 10110
FREE_SPACE : 8000
OVERCOMMIT_SIZE: 0
TOTAL_LUS : 3
TOTAL_LU_SIZE: 2110
FG_COUNT : 2
MIRROR_STATE: SYNCED
ERASURE_CODE: MIRROR2
    
```

**COMINGLED – system meta data & user data**  
**Not Default tier for lu –create with no tier option**  
**To change default:**  
**\$ tier -modify -attr TIER\_DEFAULT=YES -tier mytier**  
**'mytier' has been set as default tier successfully.**

**Not mirrored**

**USER data only tier (no meta data)**  
**Default Tier**

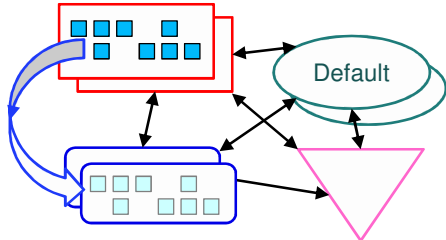
**Mirrors are in sync**  
**Two copies**



© 2015 IBM  
48  
Shared Storage Pool 5

### 6 SSP5 – Tiers: LU Live move

- Live storage migrate a LU's blocks to a different Tier
- With the LU client VM running
- \$ lu -move -lu vm42boot -dsttier bluetier
- dst = destination = IMHO ghastly





## SSP5 – Tiers: LU move

- Effectively Live moving the LU data to a different disk set
- For example:
  - faster SAN disks like Flash
  - more reliable SAN disks
  - similar disks but different location

Some ghastly slow old non-IBM SAN disk array with small disk count

Shiny fast V7000 SAN disk array with many more disks & caching

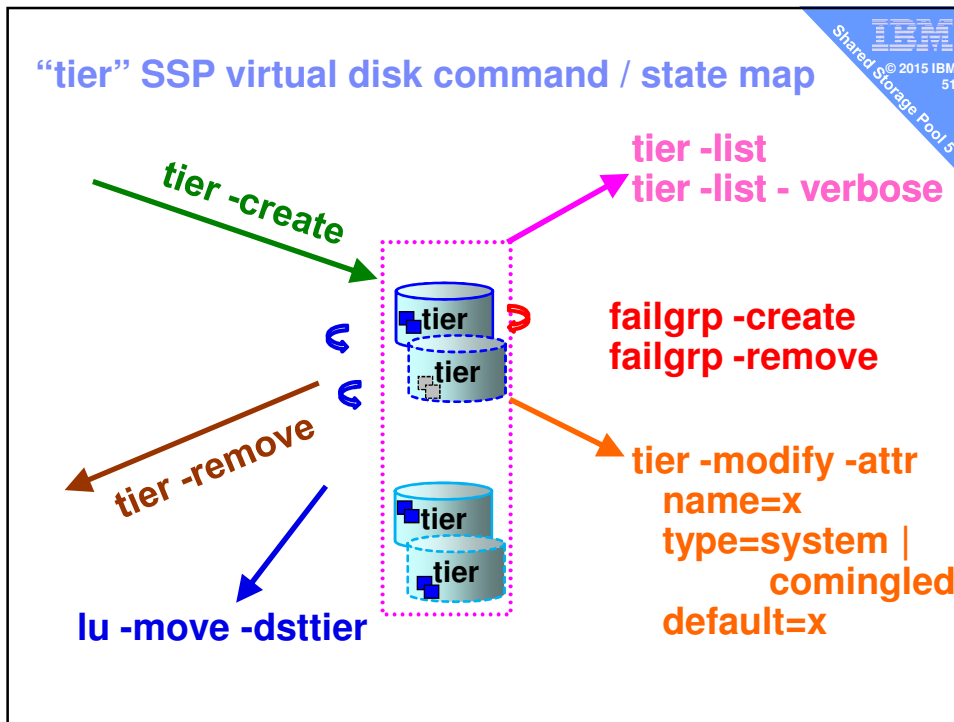
Default

If the workload is disk bound then this is a non-disruptive way to tune the disks or even experiment

## SSP5 – Tiers: LU attributes

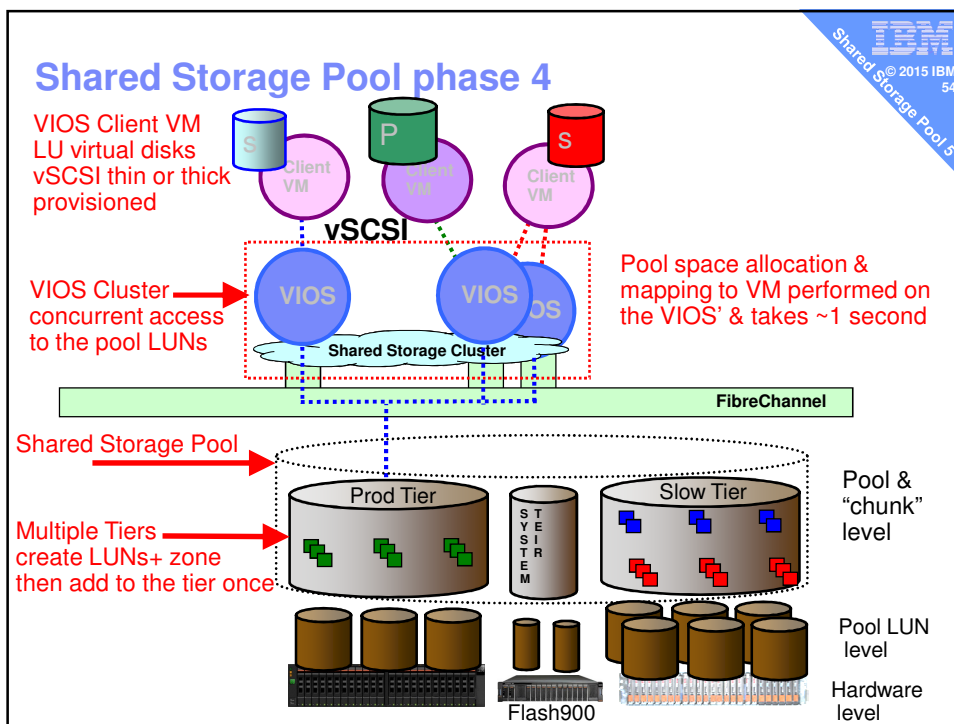
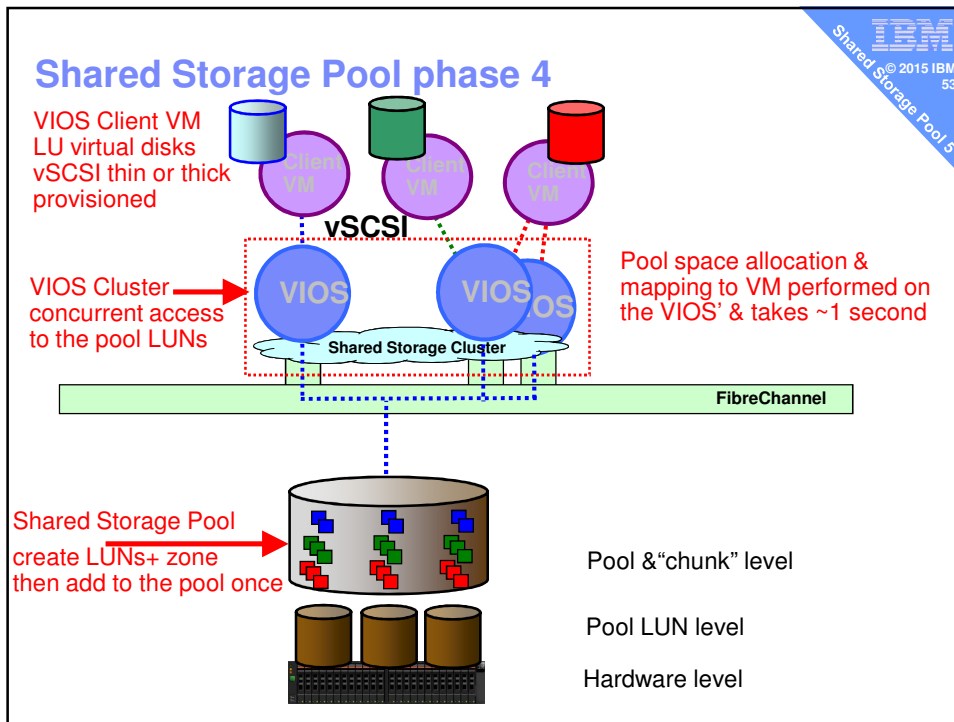
```
$ lu -list -verbose
POOL_NAME:sp1
TIER_NAME:SYSTEM
TIER_RELATION:PRIMARY          or "vacating"
ADDITIONAL_TIERS:N/A         other tier name
LU_NAME:test
LU_UDID:4b9ab8ac36f99fc6d81720528a5dd64b
LU_SIZE (MB) :10
LU_USED_PERCENT:0
LU_USED_SPACE (MB) :0
LU_UNUSED_SPACE (MB) :10
LU_PROVISION_TYPE:THIN
LU_UDID_DERIVED_FROM:N/A
LU_MOVE_STATUS:N/A
LU_SNAPSHOTS:N/A
```

There is no `lu -modify`  
i.e. you can't directly change anything



# SSP5 Worked Example

IBM © 2015 IBM 52  
Shared Storage Pool 5



## SSP4 to SSP5 upgrade?

- SSP4 = VIOS 2.2.3.x
- SSP5 = VIOS 2.2.4.x
  
- Online VIOS upgrade – business as usual
  - Check dual path VIOS
  - clstartstop -stop ...
  - updateios ... plus a VIOS reboot
  - clstartstop -start ...
  - Upgrade last VIOS → Tier functions work
  
- Online:
  - tier -create ...
  - lu -move ...

IBM  
© 2015 IBM  
55  
Shared Storage Pool 5

## 7 SSP5 HMC Classic & Enhanced+

Need to recapture with HMC 840  
as it has much more SSP support including Tiers  
our current HMC 840 has no SSP attached

IBM  
© 2015 IBM  
56  
Shared Storage Pool 5

### SSP5 – HMC Support Classic view

IBM © 2015 IBM 57  
Shared Storage Pool 5

Properties 431.25 32 512

Operations 13 Selected: 1

- Configuration
  - Create Partition
  - System Plans
  - Partition Availability Priority
  - View Workload Management Groups
  - Manage Custom Groups
  - Manage Partition Data
  - Manage System Profiles
- Virtual Resources
  - Shared Processor Pool Management
  - Shared Memory Pool Management
  - Virtual Storage Management
  - Virtual Network Management
  - Reserved Storage Device Pool Management

### SSP5 – HMC Support Classic view

IBM © 2015 IBM 58  
Shared Storage Pool 5

**Virtual Storage Management - P8-ruby-8408-E8E-SN21D494V**

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 ---

Select	Name	Storage Pool	Assigned Partition	Size
<input type="radio"/>	testa	spiral(spiral)	None	32 GB
<input type="radio"/>	testb	spiral(spiral)	None	40 GB
<input type="radio"/>	testc	spiral(spiral)	None	39 GB
<input type="radio"/>	v23456789012345678901234567890	spiral(spiral)	None	8 GB
<input type="radio"/>	vm96boot	spiral(spiral)	vm96(8)	40 GB
<input type="radio"/>	vm96data	spiral(spiral)	vm96(8)	8.06 GB
<input type="radio"/>	vm97boot	spiral(spiral)	vm97(7)	38 GB
<input type="radio"/>	vm97data	spiral(spiral)	vm97(7)	8.06 GB

Show shared storage pool storage

## SSP5 – HMC Support Classic view

  
 © 2015 IBM  
 Shared Storage Pool 5

### Virtual Storage Management - P8-ruby-8408-E8E-SN21D494V

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	Name	Storage Pool	Assigned Partition	Size
<input type="radio"/>	testa	al(spiral)	None	32 GB
<input type="radio"/>	testb	al(spiral)	None	40 GB
<input type="radio"/>	testc	al(spiral)	None	39 GB
<input type="radio"/>	v2345678901234	al(spiral)	None	8 GB
<input type="radio"/>	vm96boot	al(spiral)	vm96(8)	40 GB
<input type="radio"/>	vm96data	al(spiral)	vm96(8)	8.06 GB
<input type="radio"/>	vm97boot	al(spiral)	vm97(7)	38 GB
<input type="radio"/>	vm97data	al(spiral)	vm97(7)	8.06 GB

Create virtual disk...    Modify assignment...    Shared storage pool storage

## SSP5 – HMC Support Classic view

  
 © 2015 IBM  
 Shared Storage Pool 5

### Create Virtual Disk - P8-ruby-8408-E8E-SN21D494V

To create a virtual disk, enter a name and a size for the new disk, and select a storage pool from which to create the new disk. You also can assign 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:

Storage pool name:

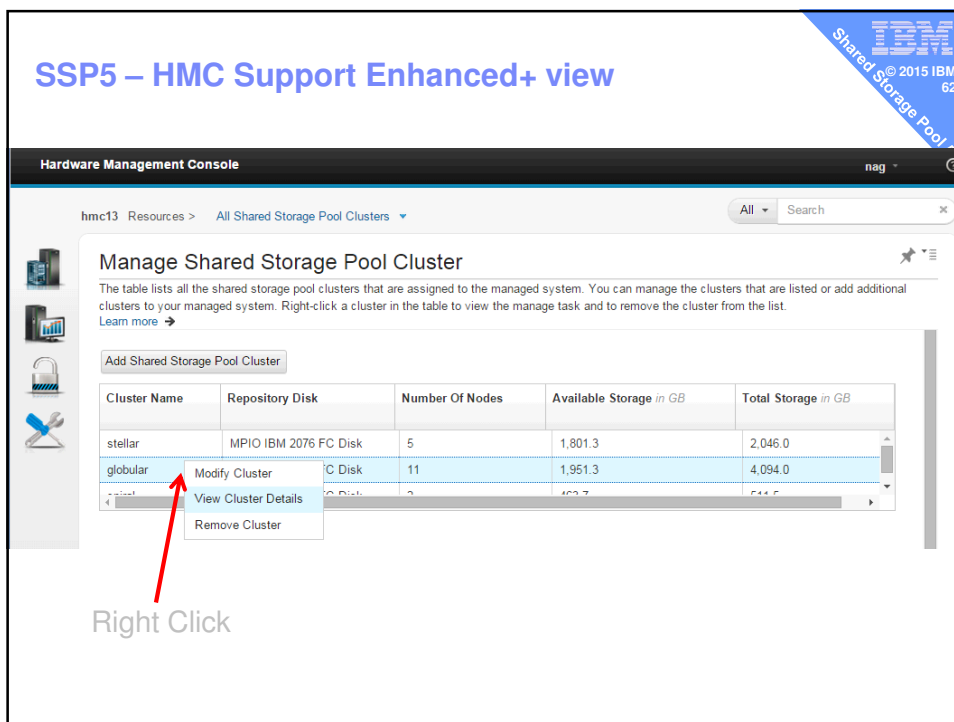
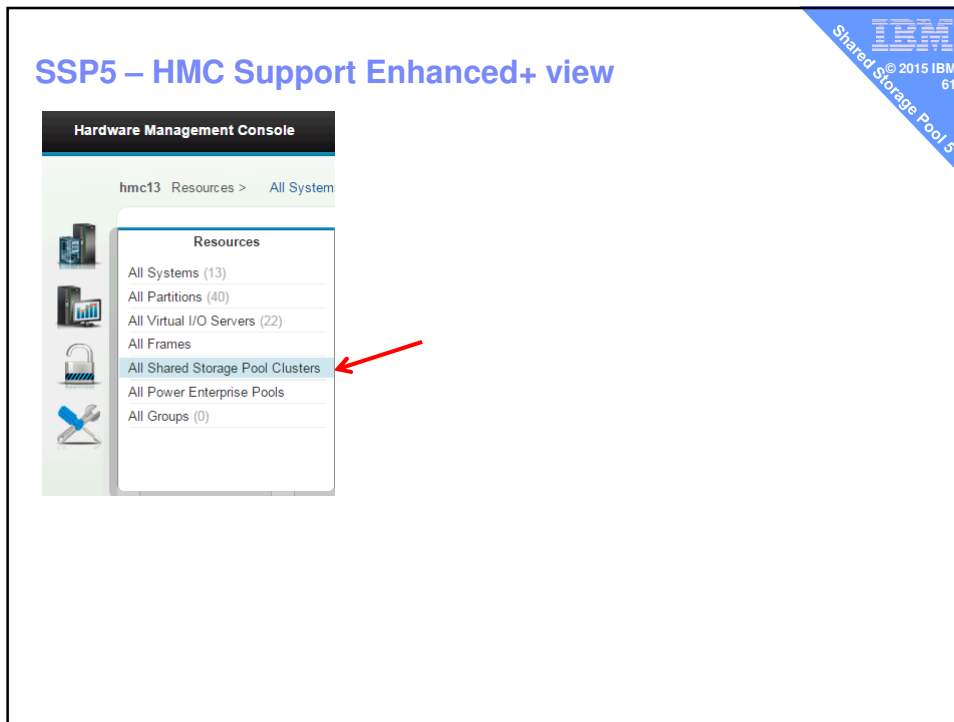
Virtual disk size:  GB

Assigned partition:

Disk type:

Map to VIOS(s):

Select	Virtual IO Server
<input checked="" type="checkbox"/>	rubyvios3
<input checked="" type="checkbox"/>	rubyvios4



### SSP5 – HMC Support Enhanced+ view

Shared Storage Size: 511.5  
 Available Storage: 463.7  
 Threshold Alert Setting: 35.0 %

Cluster Name: spiral

You can view the number of repository disks and nodes that are connected to the selected cluster. The shared storage pool name that is associated with the cluster is also displayed along with their physical volume and shared storage pool volumes details.

Repository Disk (1 required)  
MPIO IBM 2076 FC Disk

Number Of Nodes (max 16)  
2

Shared Storage Pool Name:  
spiral

Physical Volume  
2

SSP Volumes  
8

Add / Remove Nodes Cluster

Cluster Name: spiral

The table lists only the Virtual I/O Servers (VIOS) that can be configured to the shared storage pool volume number (LUN). Select the VIOS from the table to configure them on the cluster. Clear the check-box for any VIOS that you want to remove from the cluster.

	Virtual IO	ID	Cluster Membership	Server
<input checked="" type="checkbox"/>	rubyvios3	3	spiral	P8-ruby-8408-E8E-SN2
<input checked="" type="checkbox"/>	rubyvios4	4	spiral	P8-ruby-8408-E8E-SN2

Ok Cancel

Click these for different information

### HMC 840 – with Tier support

Add Tier Tiers Mirrored? Alert Thresholds

Hardware Management Console

hmc14 Resources > All Shared Storage Pool Clusters > spiral

#### Shared Storage Pool Cluster spiral

You can view and manage different properties for the shared storage pool cluster.

Shared Storage Pool: spiral

Action Add Tier

Tier Name	Available in GB	Total Capacity in GB	Mirroring	Freespace % (Threshold)	Overcommit % (Threshold)
prod (SYSTEM)	210.47	255.75	Yes	10%	50%
test (DEFAULT)	253.21	255.75	Yes	35%	0%

Repository Disk

Replace Disk

Storage Label	Size in GB	UDID	Description
spiral_repo_backup	2.00	01M0ICTTix...	MPIO IBM 2076 FC Disk

Nodes

Add Nodes Remove Node

#### VIOS Nodes

Nodes in Cluster	Node State	ID	System
rubyvios3	Up	3	P8-ruby-8408-E8E-SN21D494V
rubyvios4	Up	4	P8-ruby-8408-E8E-SN21D494V



## 8 SSP5 –Advanced

- Tier modify options
- Separate SYSTEM tier
- lu -move -nonrecursive
- Alert is tier based – warnings on SSP near full or grossly over-committed

IBM  
© 2015 IBM  
65  
Shared Storage Pool 5

### SSP5 – Tier –modify name & Default tier attribute

Tier Name:

Change tier name (like lu & failgrp modify)

`$ tier -modify -tier FRED -attr tier_name=BERT`

Note: 1<sup>st</sup> tier is called “Default” but might not be the Default tier!

#### **Default tier:**

Use `tier -list -verbose`

- And look for TIER\_DEFAULT:YES
- `lu -create` without -tier option go here

Only one tier can be the Default tier - set using

▪ `$ tier -modify -tier GREEN -attr default=yes`

– Note: can use either `default=yes` or `tier_default=yes`

IBM  
© 2015 IBM  
66  
Shared Storage Pool 5

## SSP5 – separate SYSTEM Tier

If your Shared Storage Pool is

- Large (10's of TB)
- High rates of new or changing LUs
- High rates of Thin LU extending – dynamic space allocation
- High numbers of LU tier moves

Then SSP meta-data update I/O rate can be large = effecting performance

### The SSP developers really like the idea of a separate SYSTEM tier disks

1. Separate the meta-data I/O to a different disk set
  - To reduce latency i.e. meta-data I/O not queued behind data I/O
2. Turbo charge SYSTEM tier by using FC Flash storage




## SSP5 – separate SYSTEM Tier

SYSTEM TIER “Rules of thumb”

1. When separate SYSTEM meta-data is recommended
  - If you have some fast LUNs and much slower larger User tier LUNs
  - If you have access to limited Flash LUNs
2. The size of the SYSTEM tier needed
  - 0.3% of the User Tiers





© 2015 IBM  
69  
Shared Storage Pool 5

### SSP5 – SYSTEM Tier

Tier types:

1. **SYSTEM** ← SSP meta-data only
2. **COMINGLED** ← SSP meta-data and LUs
3. **USER** ← All other Tiers for LU data


1<sup>st</sup> tier has to be type SYSTEM or COMINGLED (only one)

Set Tier type “system” or “system comingled”

- \$ tier -modify -tier RED -attr type=system
- \$ tier -modify -tier RED -attr type=comingled

If Type SYSTEM

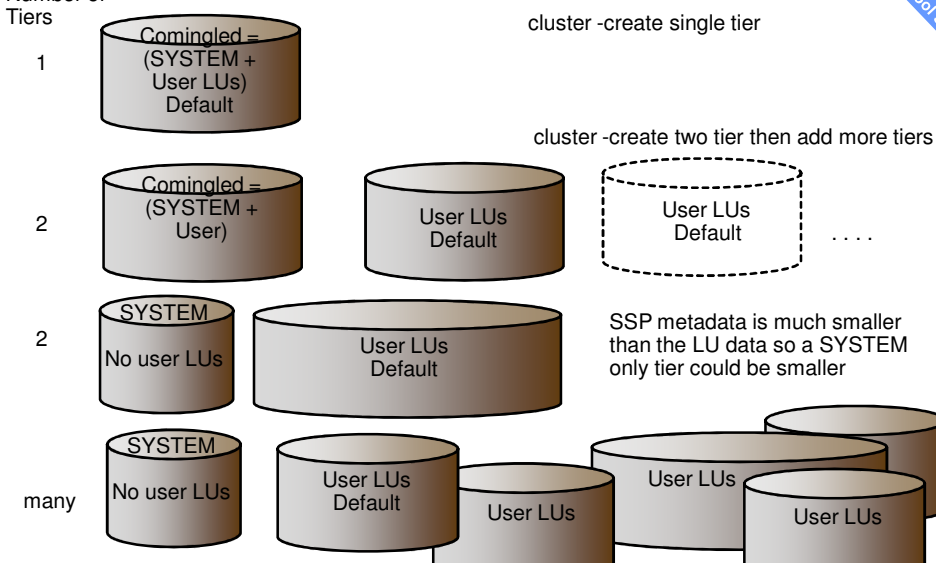
- **Can** lu -move a LU from it ← used to remove LUs from this SYSTEM Tier until no LUs left in it= only meta-data
- **Can't** lu -move a LU to it



© 2015 IBM  
70  
Shared Storage Pool 5

### SSP5 – SYSTEM Tier

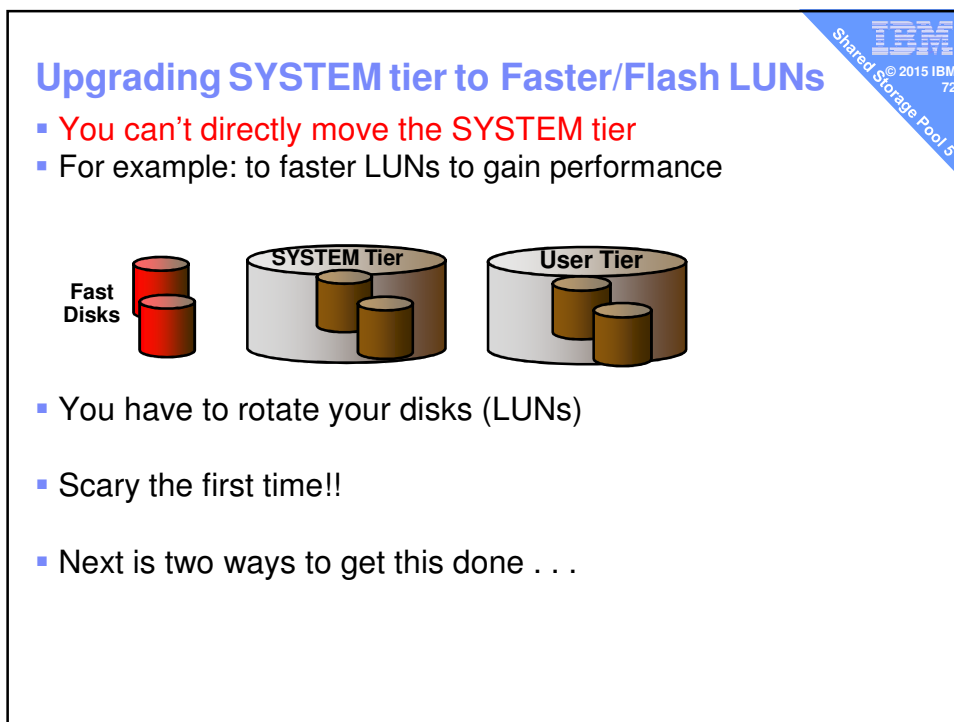
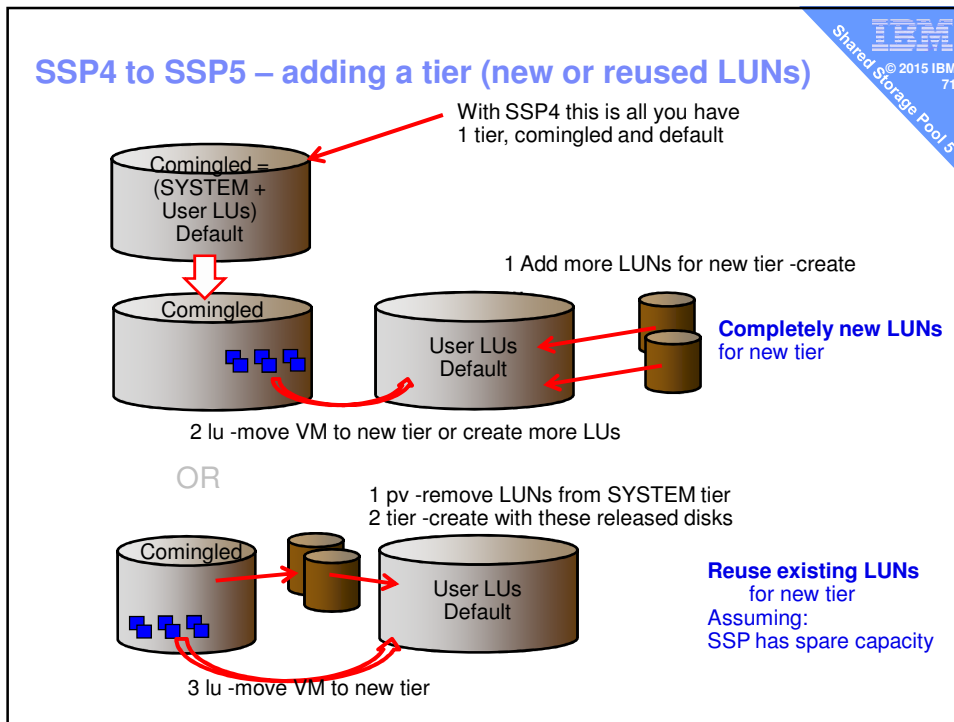
Number of Tiers



cluster -create single tier

cluster -create two tier then add more tiers

SSP metadata is much smaller than the LU data so a SYSTEM only tier could be smaller



### IBM © 2015 IBM 73 Shared Storage Pool 5

## Upgrading SYSTEM tier to Faster/Flash LUNs

- You can move the SYSTEM tier

– 1 Add fast disks: `pv -add -tier sys -fg san1: hdisk98 -fg san2: hdisk99`

– 2 Then remove slow: `pv -remove -pv hdisk42 hdisk52`

– 3 Then add to User tier: `pv -add -tier user -pv hdisk42 hdisk52`

### IBM © 2015 IBM 74 Shared Storage Pool 5

## Upgrading SYSTEM tier to Faster/Flash LUNs

- You can move the SYSTEM tier

– 1 One step Swap out: `pv -replace -oldpv hdisk42 -newpv hdisk89`  
`pv -replace -oldpv hdisk52 -newpv hdisk99`

– 2 Then add to User tier: `pv -add -tier user -pv hdisk42 hdisk52`

## Fast/Flash System Tier in practice

Don't worry – be happy!

- Good to monitor SYSTEM meta-data disk I/O
- ATS has a 128GB min LUN →  $128 / 0.3\% = 39\text{TB}$ 
  - Result just one LUN (pair) for the tier = not good for I/O
  - Flash is expensive → smaller LUN size so lots of LUNs
- VIOS Advisor spots hot disks, small queue depth ...
- SYSTEM I/O rate only an issue on
  - Very large SSP and very busily changing SSP config
  - High numbers of VIOS (future!)

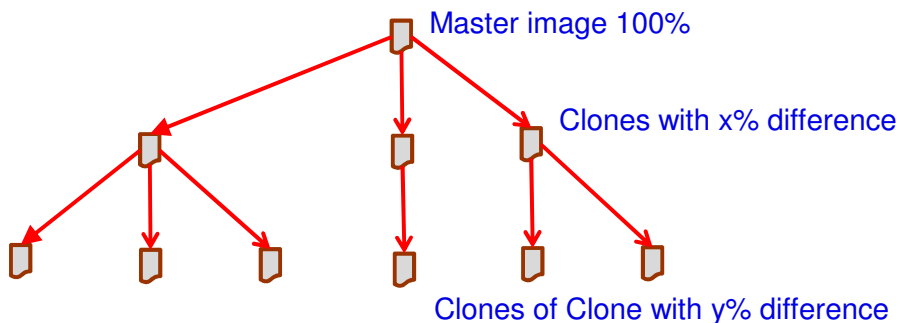
IBM  
© 2015 IBM  
75  
Shared Storage Pool 5

## SSP5 – Tiers: LU Advanced move with Clones

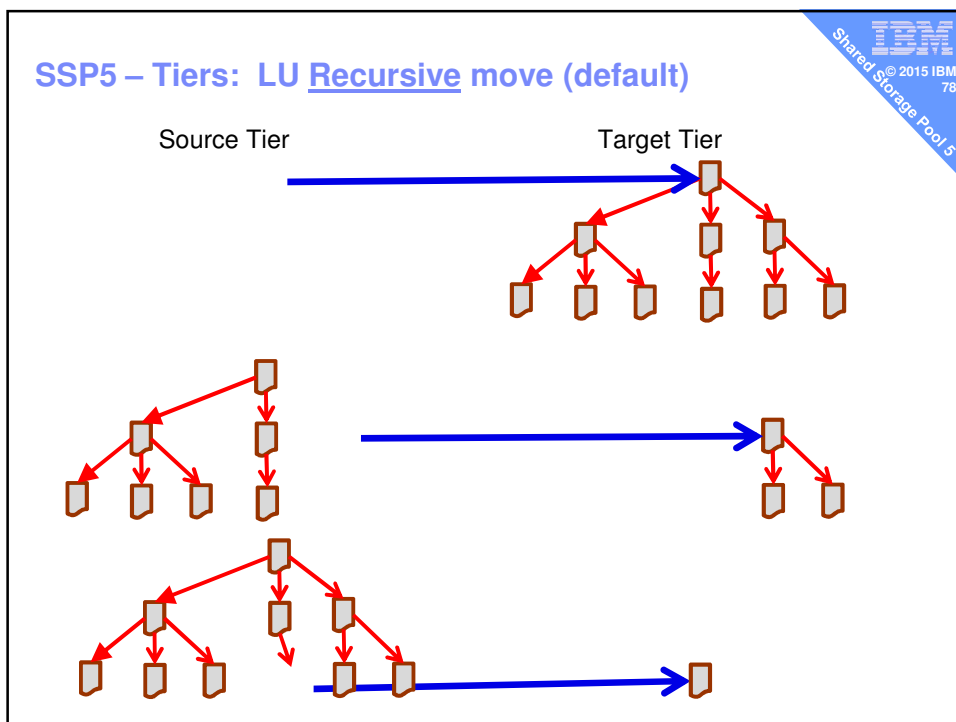
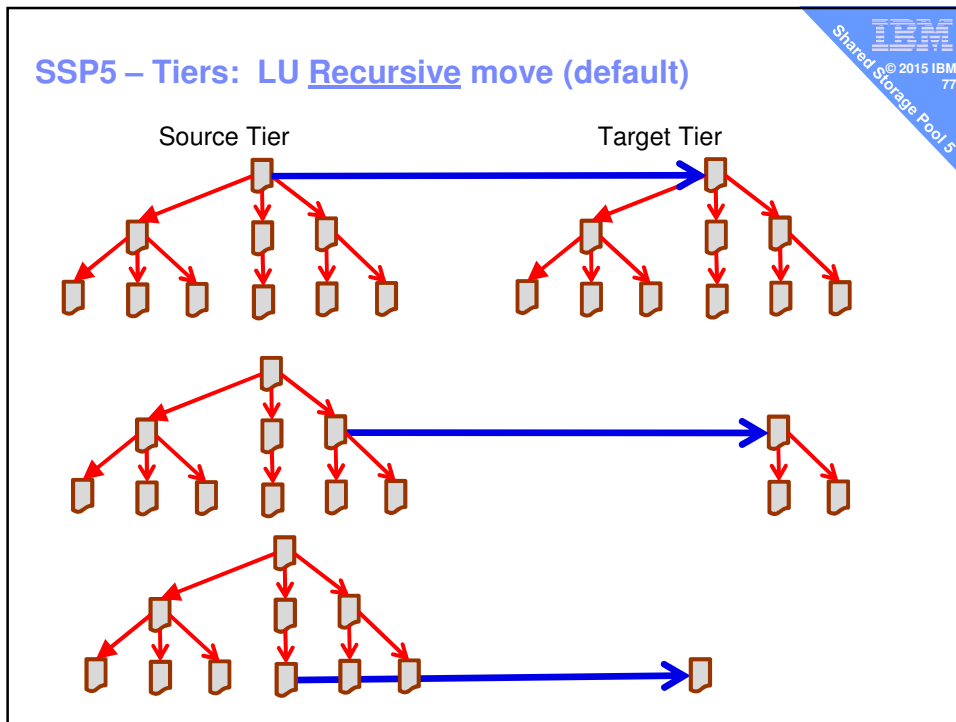
- `lu -move -lu vm42 -dsttier blue -nonrecursive`

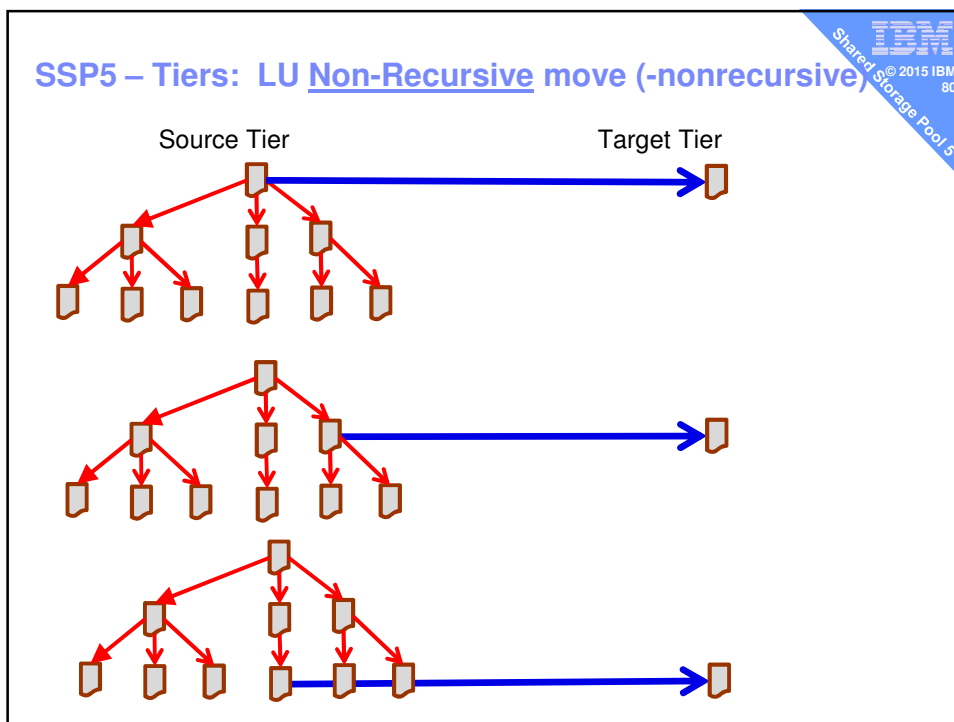
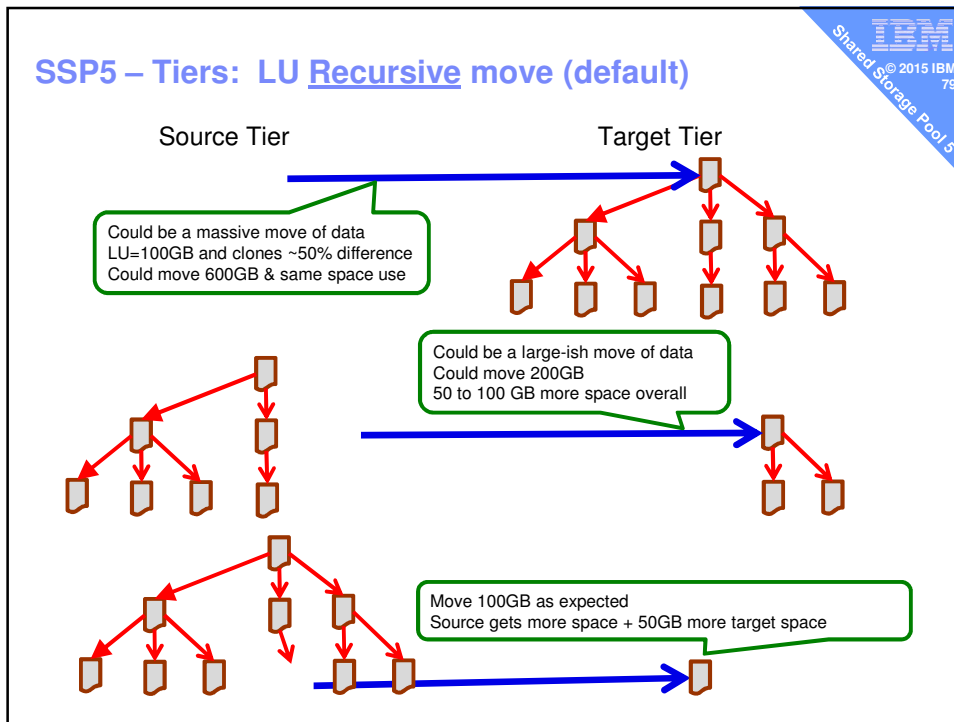
Background

- SSP system admin can't normally create LU clones
- The commands are not documented → Ignoring “under the hood” fiddlers!
- But [System Director](#) and [PowerVC](#) use clones & clones of clones
- Then the clones can have snapshots or further clones

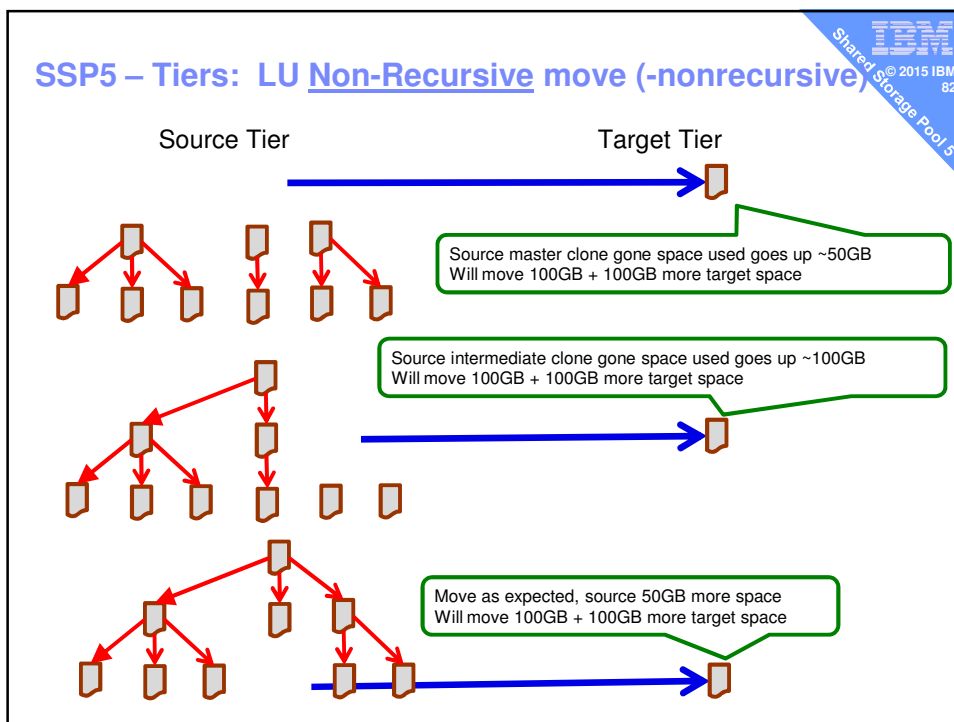
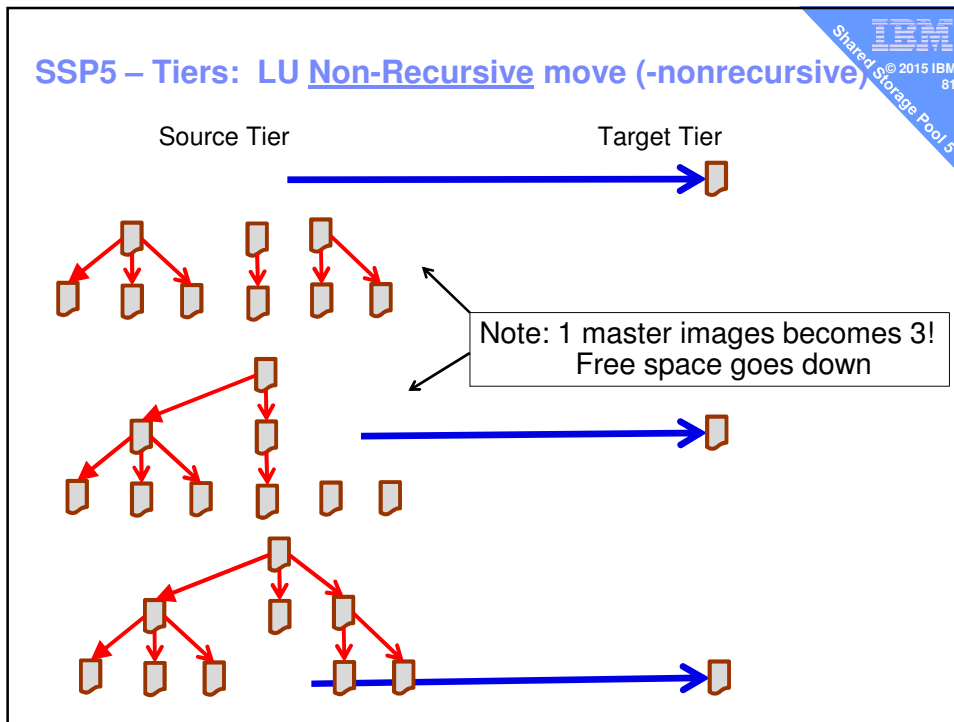


IBM  
© 2015 IBM  
76  
Shared Storage Pool 5









## SSP5 – Summary lu -move with clones



- No clones = no complexity
- With clone hierarchy it can get complex
  - Probably a problem that you which I had never told you about ☺
- It can result in
  - higher than expected large data moves
  - extra space on source **and/or** target
- But it can be used to break a PowerVC clone away
  - Example: clone AIX LU & then install Linux on it = no point being a clone
  - lu -move -lu vm42 -dsttier blue -nonrecursive
  - lu -move -lu vm42 -dsttier red ← Optional move back to original tier

## alert -tier mandatory → Messaged sent to the HMC



- By example - my tier is called “prod”
  - alert -set -tier prod -type threshold -value 10
  - alert -set -tier prod -type overcommit -value 50
  - Threshold is the Pool “free space” getting low  
i.e. 10 means alert when Pool free space crosses to below the 10%
  - Overcommit warns when you go “to far” & risk problems later on.  
50% might be accept on not busy VM's but 500% is sure to bite you!

```
$ alert -list
PoolName:spiral
PoolId:000000009893ED90000000560AA6E3
TierName:prod
ThresholdPercent:10
OverCommitPercent:50

PoolName:spiral
PoolId:000000009893ED90000000560AA6E3
TierName:test
ThresholdPercent:35
OverCommitPercent:N/A
```

*Nigel's Recommendation:  
Set alert values but monitor  
in coming Alerts via HMC  
(which every customer  
should be watching) !!*

← The defaults

## Summary VIOS 2.2.4 with SSP Tiers

1. SSP LU Resize (grow = saves admin time)
2. DIY lu -list
3. SSP Tiers (multiple pools only better)
  - 10 tiers (think grouping not levels)
  - Fast, medium, slow or IBM, HDS, EMC or prod, in-house, test
  - lu -create -lu fred -size 32G -tier prod lu -move -lu fred -dsttier test
4. SSP mirrored now at tier level (was whole pool)
5. SSP LU move between tiers
6. Possible separate SYSTEM tier LUNs
7. HMC GUI Support

© Copyright IBM Corporation 2015

