



# Exploiting Virtualisation on IBM Power Systems with PowerVM



**Nigel Griffiths**  
IBM Power Systems  
Advanced Technology Support  
EMEA



© 2011 IBM  
2

## Processor Virtualisation (CPU Sharing)



Pre-2000

## Processor Virtualisation (CPU Sharing)




<b>Old Style</b> Separate Systems	<b>LPAR</b> Server Consolidation LPAR size via start time boundaries
Pre-2000	~2001

## Processor Virtualisation (CPU Sharing)




<b>Old Style</b> Separate Systems	<b>LPAR</b> Server Consolidation LPAR size via start time boundaries	<b>DLPAR</b> Dynamic live boundary changes Manual or scripts
Pre-2000	~2001	~2002

## Processor Virtualisation (CPU Sharing)



Old Style	LPAR	DLPAR	SPLPAR
Separate Systems	Server Consolidation LPAR size via start time boundaries	Dynamic live boundary changes Manual or scripts	Shared Processor automatic adjusts at millisecond level by Hypervisor
Pre-2000	~2001	~2002	~2005

## Processor Virtualisation (CPU Sharing)



**SPLPAR**  
Shared Processor automatic adjusts at millisecond level by Hypervisor

~2005

## Processor Virtualisation (CPU Sharing)



### SPLPAR

Shared Processor  
automatic adjusts  
at millisecond level  
by Hypervisor

~2005

## Processor Virtualisation (CPU Sharing)



### SPLPAR

Shared Processor  
automatic adjusts  
at millisecond level  
by Hypervisor

~2005

### Harvesting

“Spare” capacity  
ready for adding  
more workloads at  
zero hardware cost

~2006

## Processor Virtualisation (CPU Sharing)



### SPLPAR

Shared Processor  
automatic adjusts  
at millisecond level  
by Hypervisor

~2005

### Harvesting

“Spare” capacity  
ready for adding  
more workloads at  
zero hardware cost

~2006


### Partition Mobility

Make a cluster of  
your machines &  
flow your workload  
between them



~2008

## Virtualisation - Value Proposition

## Virtualisation - Value Proposition




## Virtualisation - Value Proposition




## Virtualisation - Value Proposition

"Your new system will be ready in ..."  
"20 Minutes" or "20 Days"




Quick Deployment




Reduced Infrastructure

## Virtualisation - Value Proposition


"Your new system will be ready in ..."  
"20 Minutes" or "20 Days"



Quick Deployment



Reduced Infrastructure



Server Consolidation


## Virtual Adapters

The I/O centric view of the world:


- CPU used to “modify & feed data” to the networks & disks



## Virtual I/O Server (Adapter Sharing)




Pure Direct

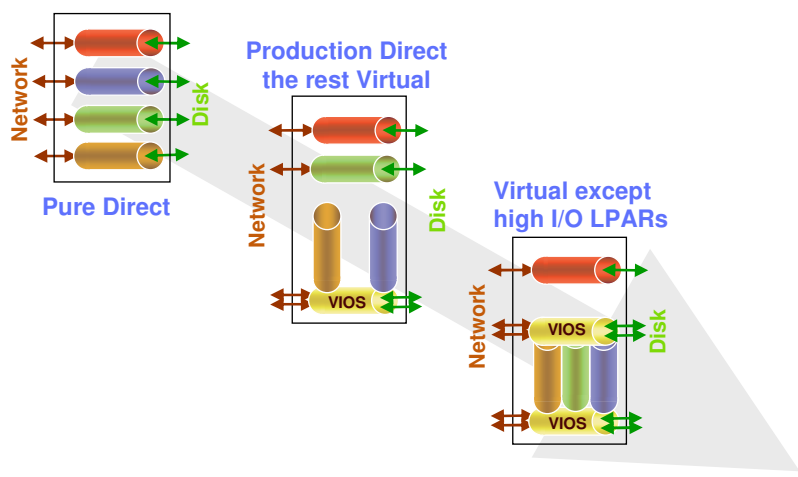




### Virtual I/O Server (Adapter Sharing)



### Virtual I/O Server (Adapter Sharing)





**Where are You?**

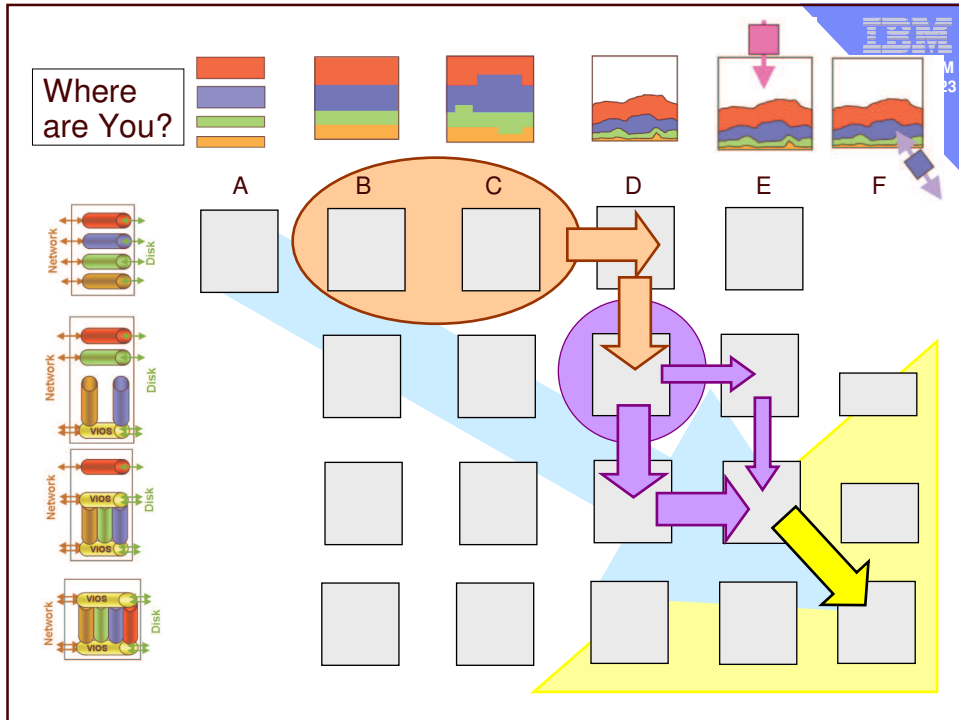
M 21

	A	B	C	D	E	F
Network						
Disk						
Network						
Disk						
Network						
Disk						
Network						
Disk						


**Where are You?**

M 22


	A	B	C	D	E	F
Network						
Disk						
Network						
Disk						
Network						
Disk						
Network						
Disk						



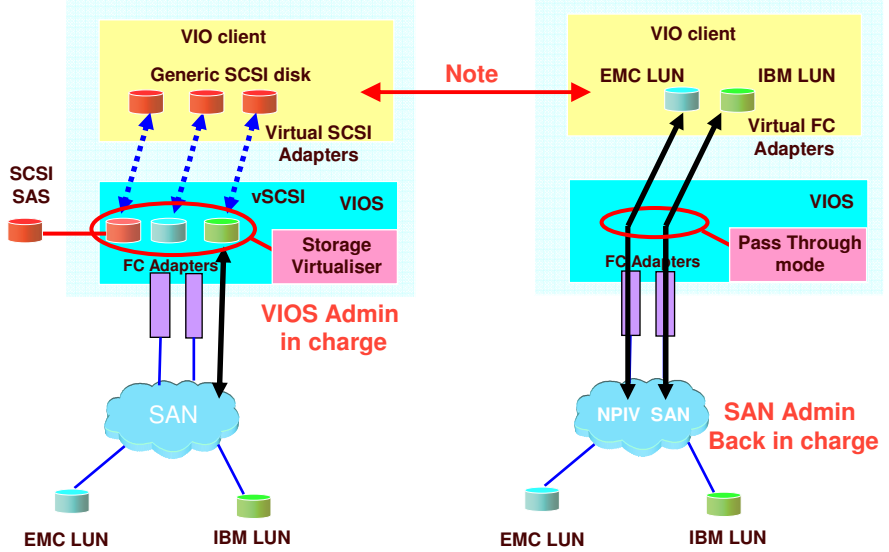
### Storage Virtualisation since 2005



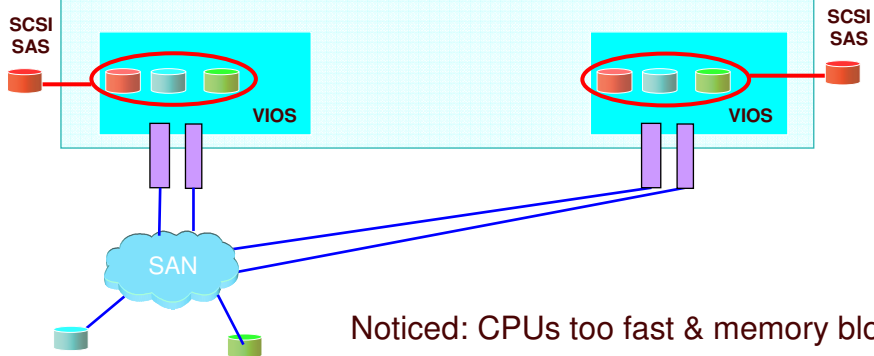
### Virtual CD/DVD & Virtual Tape



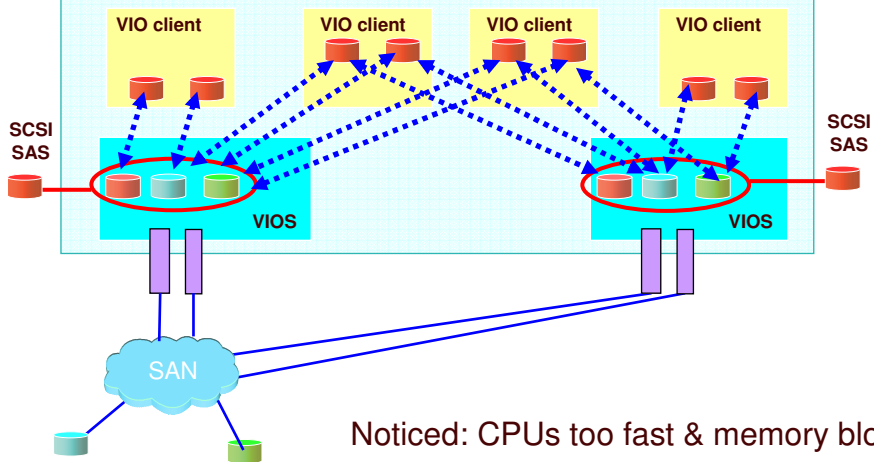
### Storage Virtualisation with N-Port ID Virtualisation



Dual VIOS since 2001 allows concurrent updates  
→ supporting 2 to 200+ Logical Partitions



Dual VIOS since 2001 allows concurrent updates  
→ supporting 2 to 200+ Logical Partitions



## What about Virtual Memory!

Virtual Memory available in UNIX to 30+ years

Modern Software need extreme amounts of memory

Power has two mechanisms to increase memory use

- Active Memory Expansion (AIX only)
- Active Memory Sharing
  
- AIX Workload Partition can also save memory

## AME Conceptual Model

Active Memory Expansion

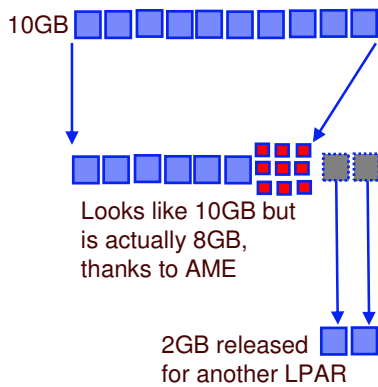
Memory Pages




Dynamically adjusted depending on compression ratio & target

# AME - What is your Plan?


**Memory Shrinking**  
to release RAM for other uses



**Memory Growing**  
for RAM optimisation & performance




## Classic Virtual Memory (LPAR)







## Active Shared Virtual Memory (LPAR)



## Active Memory Sharing (Virtual Virtual Memory!)

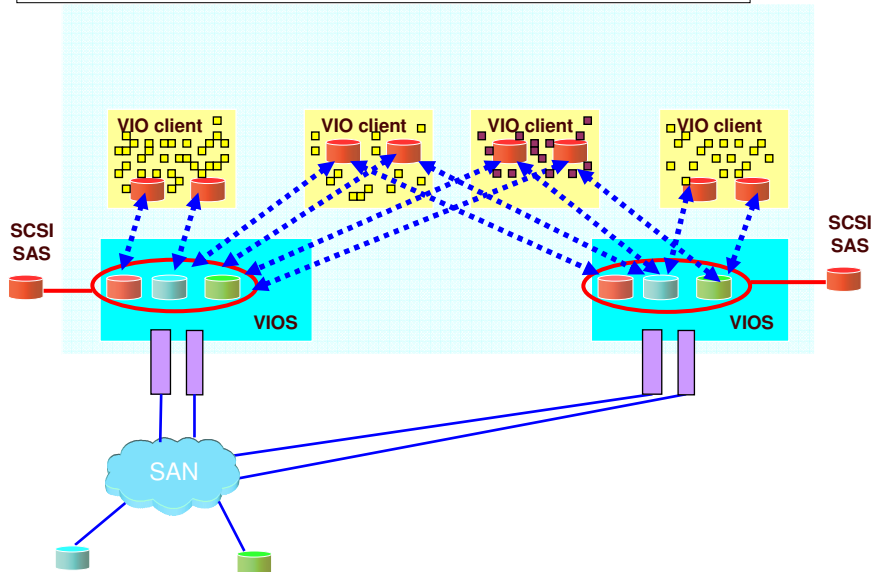


### Active Memory Sharing (Virtual Virtual Memory!)




Loaning Mode

### Active Memory Sharing (Virtual Virtual Memory!)



Active Memory Sharing (Virtual Virtual Memory!)



AME & AMS Comparisons

Active Memory Expansion

- Jan 2010
- AIX6 TL4+ on POWER7
  - Not Linux nor IBM I
- Machine Activation (LPP)
  - 60 day trial
- Pure Virtual LPAR
- Internal to single LPARs
- Assume "spare" CPU cycles for compression
- Simple to setup in LPAR
- Use amepat to predetermine the compression factor
- Use topas/nmon to monitor

Active Memory Sharing

- May 2009
- POWER6
- AIX6 TL3+, Linux & IBM i 6.1
- PowerVM Enterprise
- Pure Virtual LPAR
- Cooperating group of LPARs
- Assumes loanable RAM
- Pages flow between LPARs at a few MB/s
- More complex to setup on VIOS & LPARs
- Use topas -C to monitor

## Summary so far ....

### Dedicated World: still available

Dedicated CPU  
Dedicated Virtual Memory  
Dedicated Disk adapters - SCSI/SAS/SAN  
Dedicated Network adapters  
NFS or NIM !  
CPU load balancing in minutes/hours  
Memory load balancing in minutes/hours

### Virtual World

Virtual CPUs  
Shared Virtual Memory (AMS) + AME  
Virtual Disks (vSCSI & NPIV)  
Virtual Networks  
Virtual CD/DVD & Tape & NFS & NIM  
CPU load balancing in milli-seconds  
Active Memory Sharing between LPARs in seconds

### Live Partition Mobility

## Summary so far ....

### Dedicated World: still available

Dedicated CPU  
Dedicated Virtual Memory  
Dedicated Disk adapters - SCSI/SAS/SAN  
Dedicated Network adapters  
NFS or NIM !  
CPU load balancing in minutes/hours  
Memory load balancing in minutes/hours

### Virtual World

Virtual CPUs  
Shared Virtual Memory (AMS) + AME  
Virtual Disks (vSCSI & NPIV)  
Virtual Networks  
Virtual CD/DVD & Tape & NFS & NIM  
CPU load balancing in milli-seconds  
Active Memory Sharing between LPARs in seconds

### Live Partition Mobility

“Expensive = Old school”  
“Cheap & Flexible”

## Live Partition Mobility (LPM)

High levels of efficient Virtualisation means

- Hypervisor support (system firmware)
- Virtual I/O Server (VIOS)
- HMC for management
- Operating System support

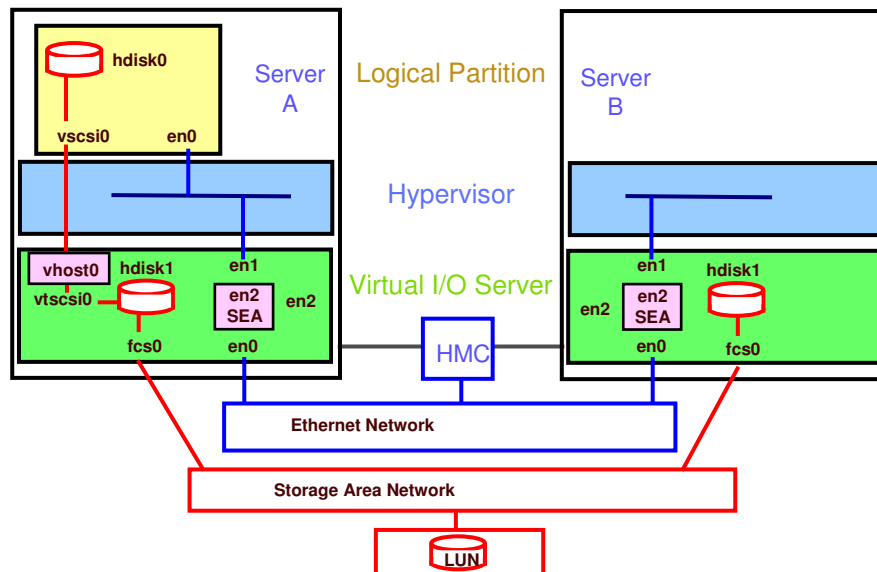
For new features & fixes ... updates are required

LPM allows for zero downtime updates

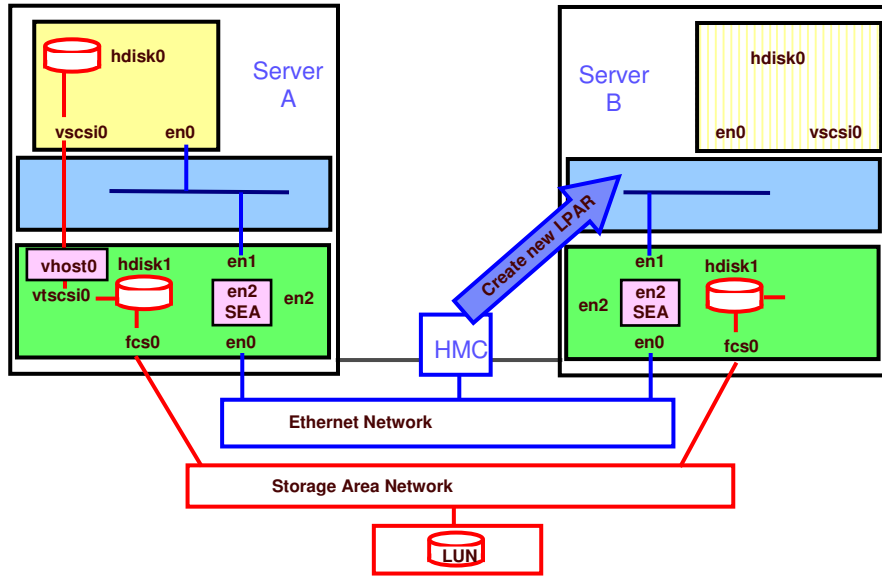
Also allows

1. Workload balancing across machine room
2. New machine use at day 1
3. Repair actions & upgrades

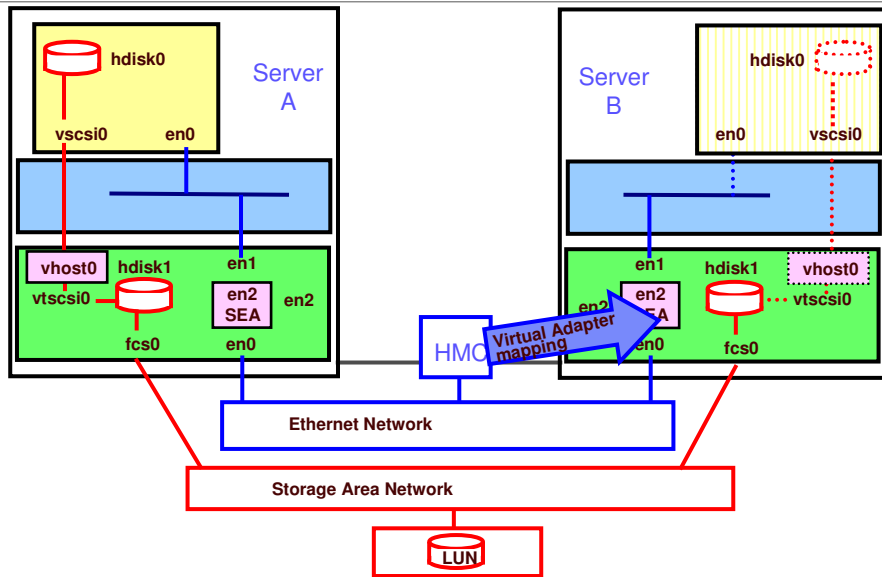
## Live Partition Mobility (1 of 6)



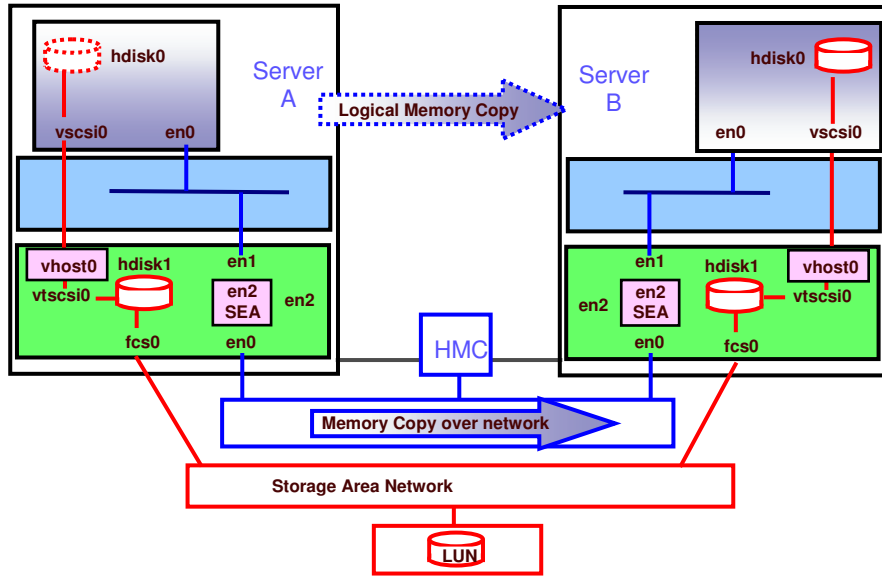
## Live Partition Mobility (2 of 6)



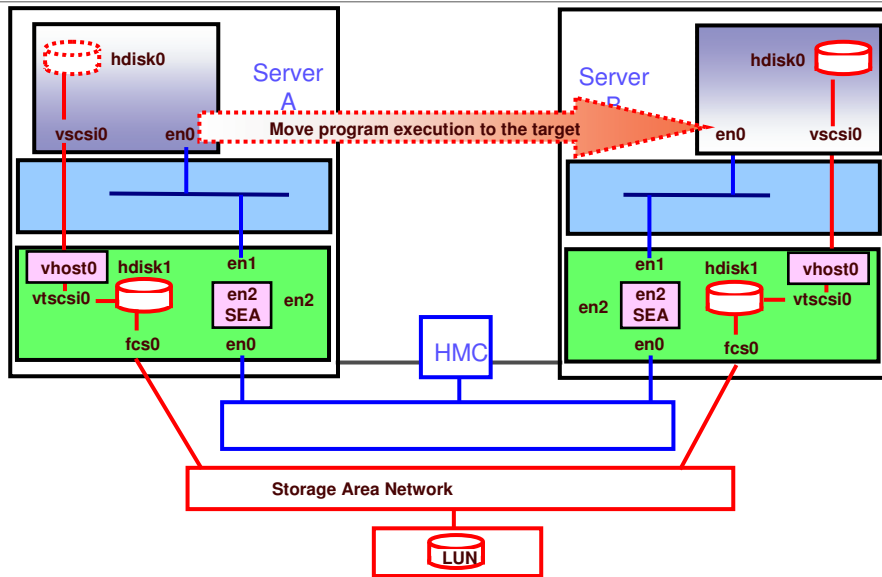
## Live Partition Mobility (3 of 6)



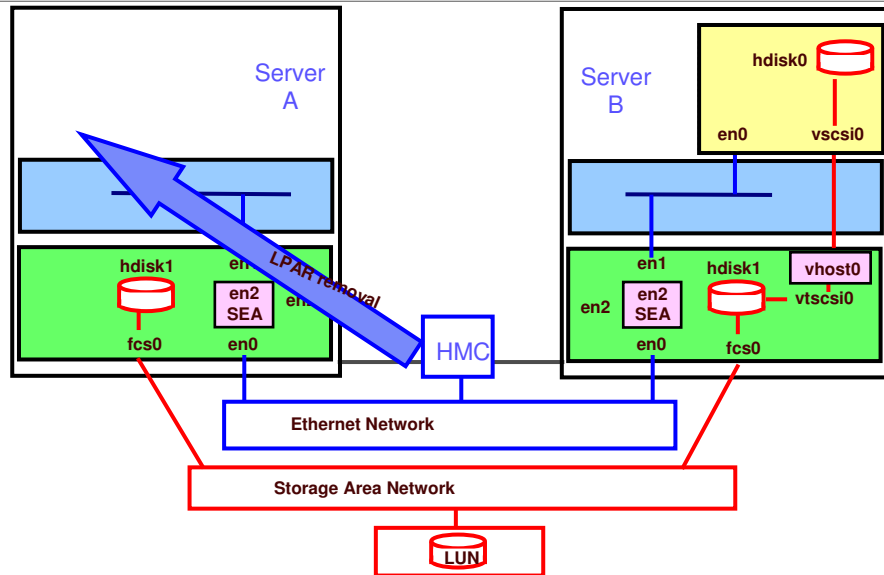
## Live Partition Mobility (4 of 6)



## Live Partition Mobility (5 of 6)



## Live Partition Mobility (6 of 6)



## Suspend and Resume

### Resource balancing

- suspend low-priority for more urgent processes

### Simplified maintenance

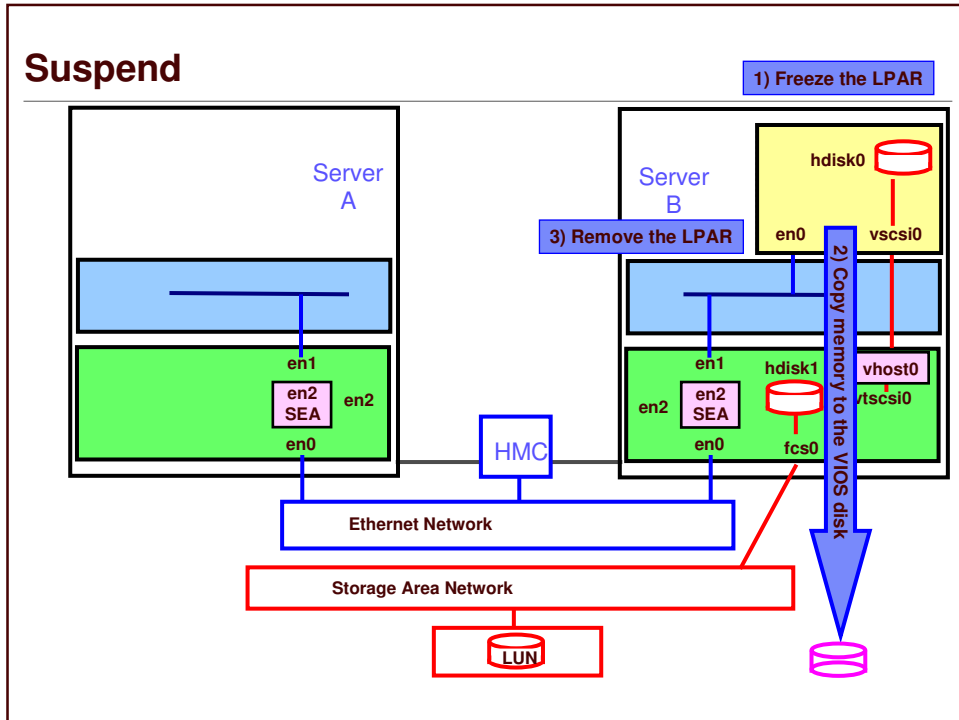
- suspend the whole machine system updates

### Debug/forensics

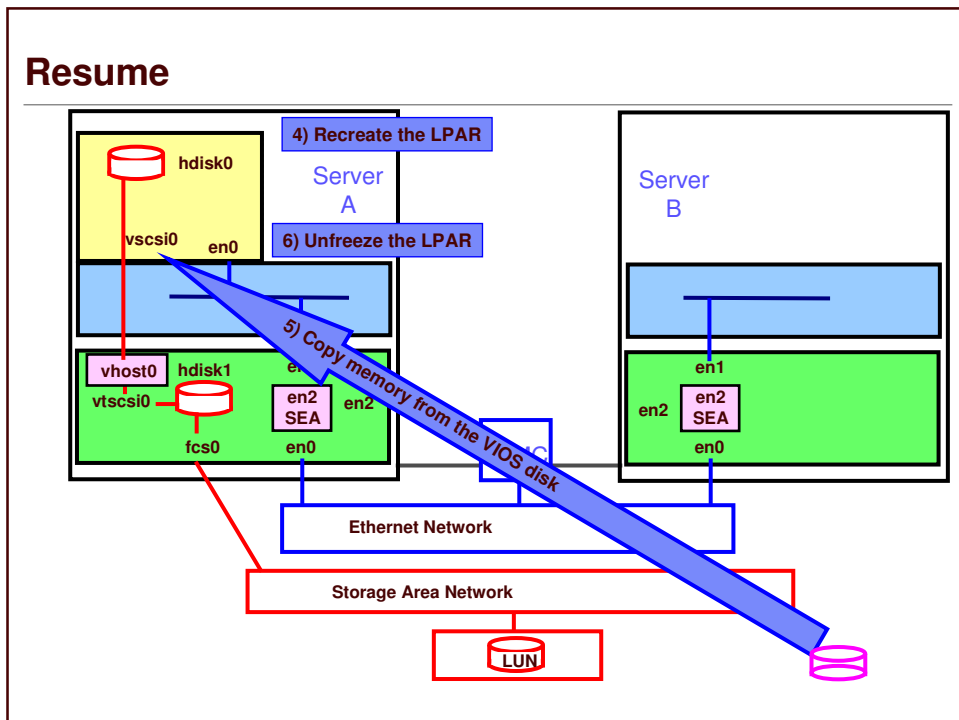
- a workload can be temporarily suspended
- a copy made for offline analysis for security or performance purposes



## Suspend



## Resume



## PowerVM Editions are tailored to client needs

**PowerVM Editions** offer a unified virtualization solution for all Power workloads

- **PowerVM Express Edition**
  - Evaluations, pilots, PoCs
  - Single-server projects
- **PowerVM Standard Edition**
  - Production deployments
  - Server consolidation
- **PowerVM Enterprise Edition**
  - Multi-server deployments
  - Advanced Functions

PowerVM Editions	Express	Standard	Enterprise
Concurrent VMs	VIOS + 2 per VMs	10 per core (up to 1000)	10 per core (up to 1000)
Virtualization Management	IVM	IVM, HMC	IVM, HMC
Virtual I/O Server	✓	✓✓	✓✓
PowerVM Lx86	✓	✓	✓
Suspend/Resume		✓	✓
Shared Processor Pools		✓	✓
Shared Storage Pools		✓	✓
Thin Provisioning		✓	✓
Live Partition Mobility			✓
Active Memory Sharing			✓



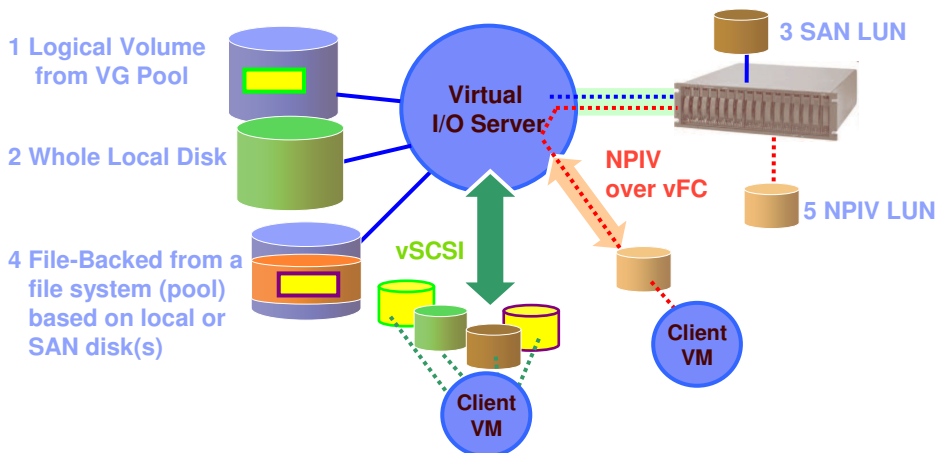
Power your planet.



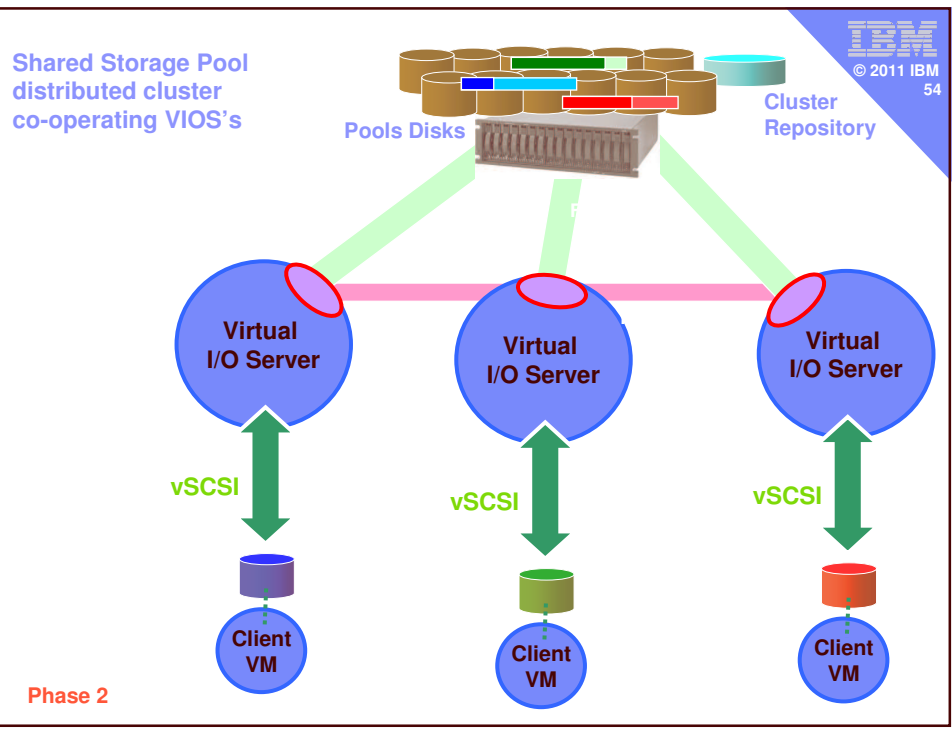
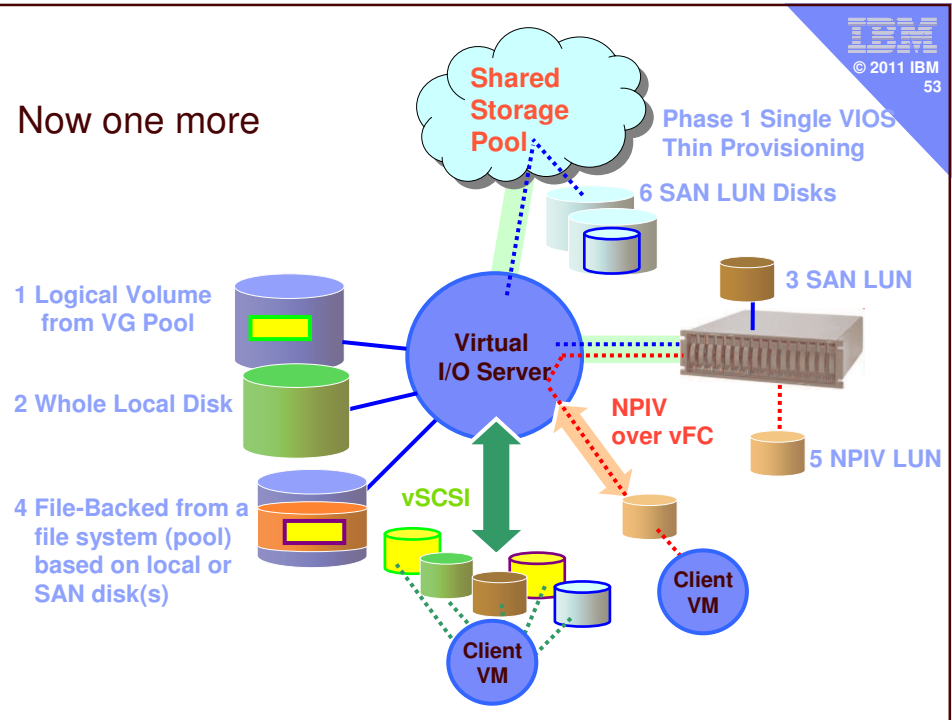
\*IBM i supports shared storage and Suspend & Resume, but does not support LPM.



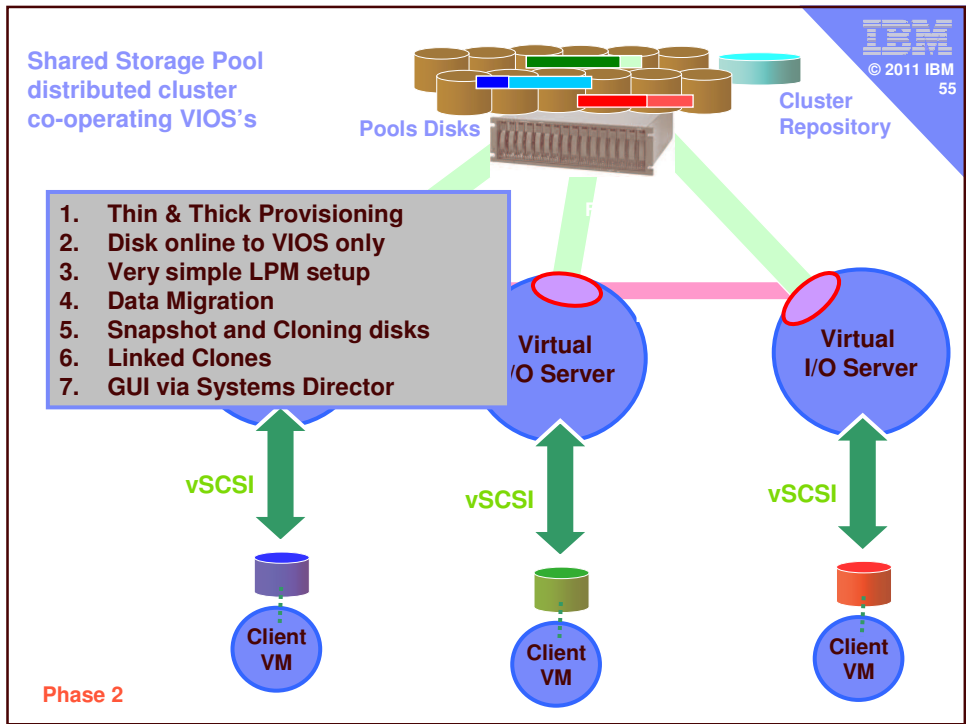
## Virtual Disks 5 type - Reminder



Now one more



Phase 2



# Power Systems

Since September 2010

© 2011 IBM 56

Power 795  
256 core

Power 780  
64 core

Power 770  
32 core

Power 750  
32 core

Power 740  
4 - 16 core

Power 720  
4 - 16 core

Power 730  
4 - 16 core

Power 710  
4 - 16 core

700/1/2 Blades  
8 or 16 core

AIX

i for Business

Linux

PowerHA

PowerVM

IBM Systems Director