Trends & Directions in IT Infrastructure

Ludwig Van den bergh
Technical Director BeNeLux
IBM Systems and Technology Group

Systems & Technology Group

Systems & Technology Group

We are adicted to technology

Adicted to technology

- **R&D budget is about 6.000.000.000 $ per year**
- **5 Nobel Price Winners**
- **WW leader in patents in 2009 (more than the sum of HP, Dell, Sun, Oracle, Intel, Apple and EMC) - leader in the past 17 years.**

→ Royalties on patents

Creating Volume Markets for technology

PowerPC Microprocessors used in ALL game consoles : Nintendo, XBOX en Playstation 2.

Latest Cell Microprocessor used in Playstation 3.

50% of the automotive market

Agenda

- ✓ Server trends
- ✓ Storage trends

Agenda

- ✓ Server trends

Virtualize everything
Multi-OS hosting
Reduce (un)scheduled downtimes

Utilization Of Distributed Servers

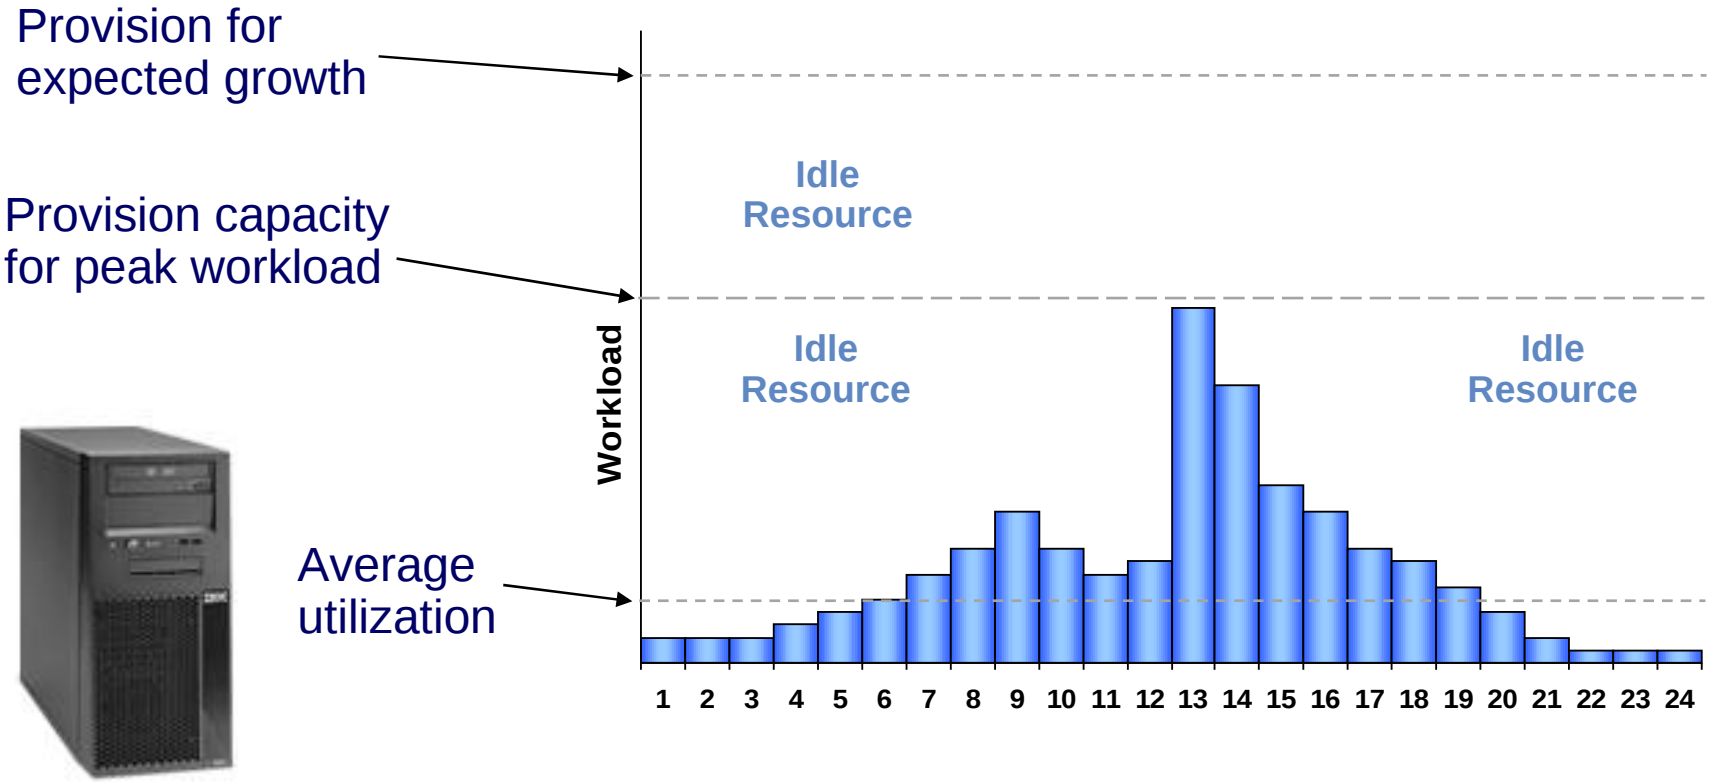

Server dedicated to one application

- Typical utilization of Windows Servers <5%
- Typical utilization of UNIX Servers 15 – 20%
- Typical utilization of System z Servers 70 – 100%

Drive for more effectiveness

Physical Partitioning

16-way SMP not partitioned

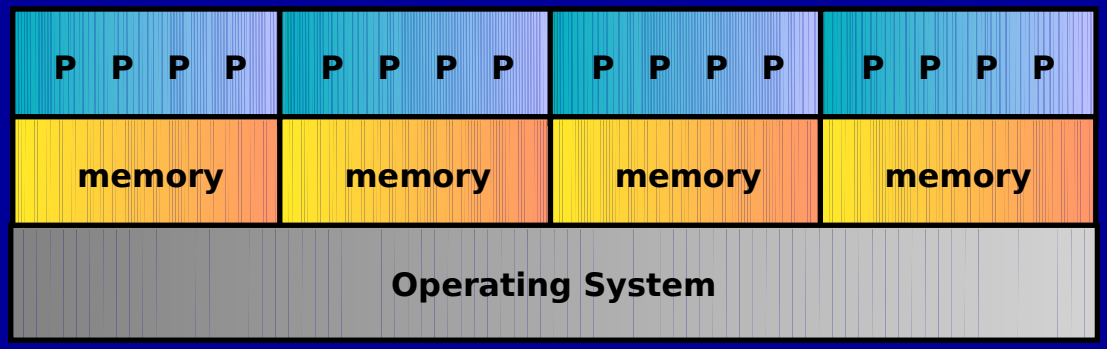

Same H/W, three systems to manage

Partitioning Steps:

1. Reconfigure system
2. Reinstall operating systems, applications, etc.
3. Reboot

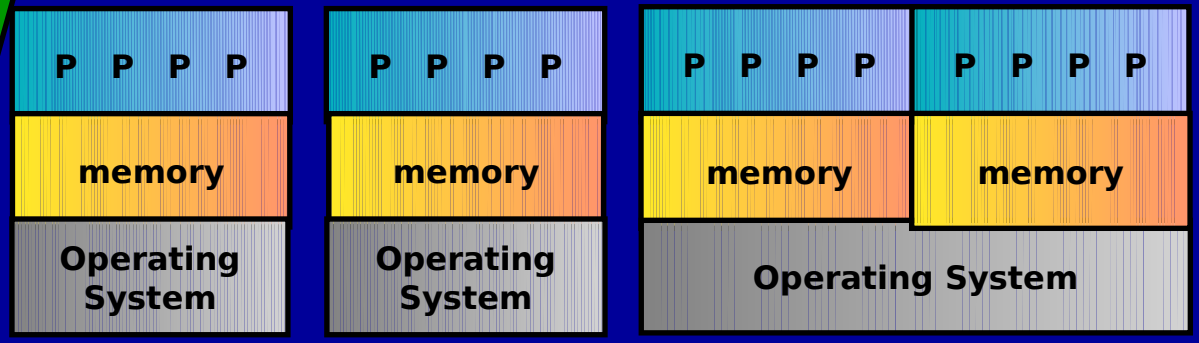

Logical Partitioning

- Resources can be dynamically moved between partitions without restart
- Can happen without operator intervention

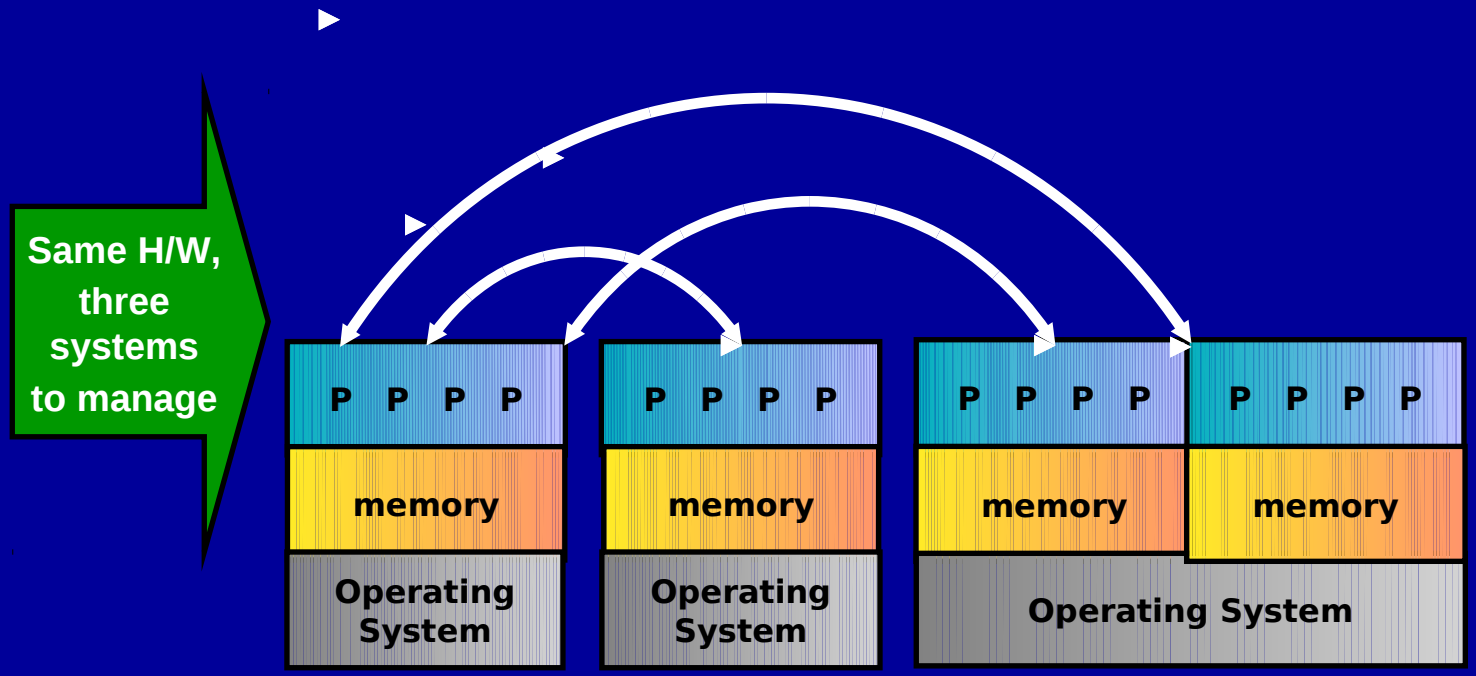

Logical Partitioning

- Resources can be dynamically moved between partitions without restart
- Can happen without operator intervention

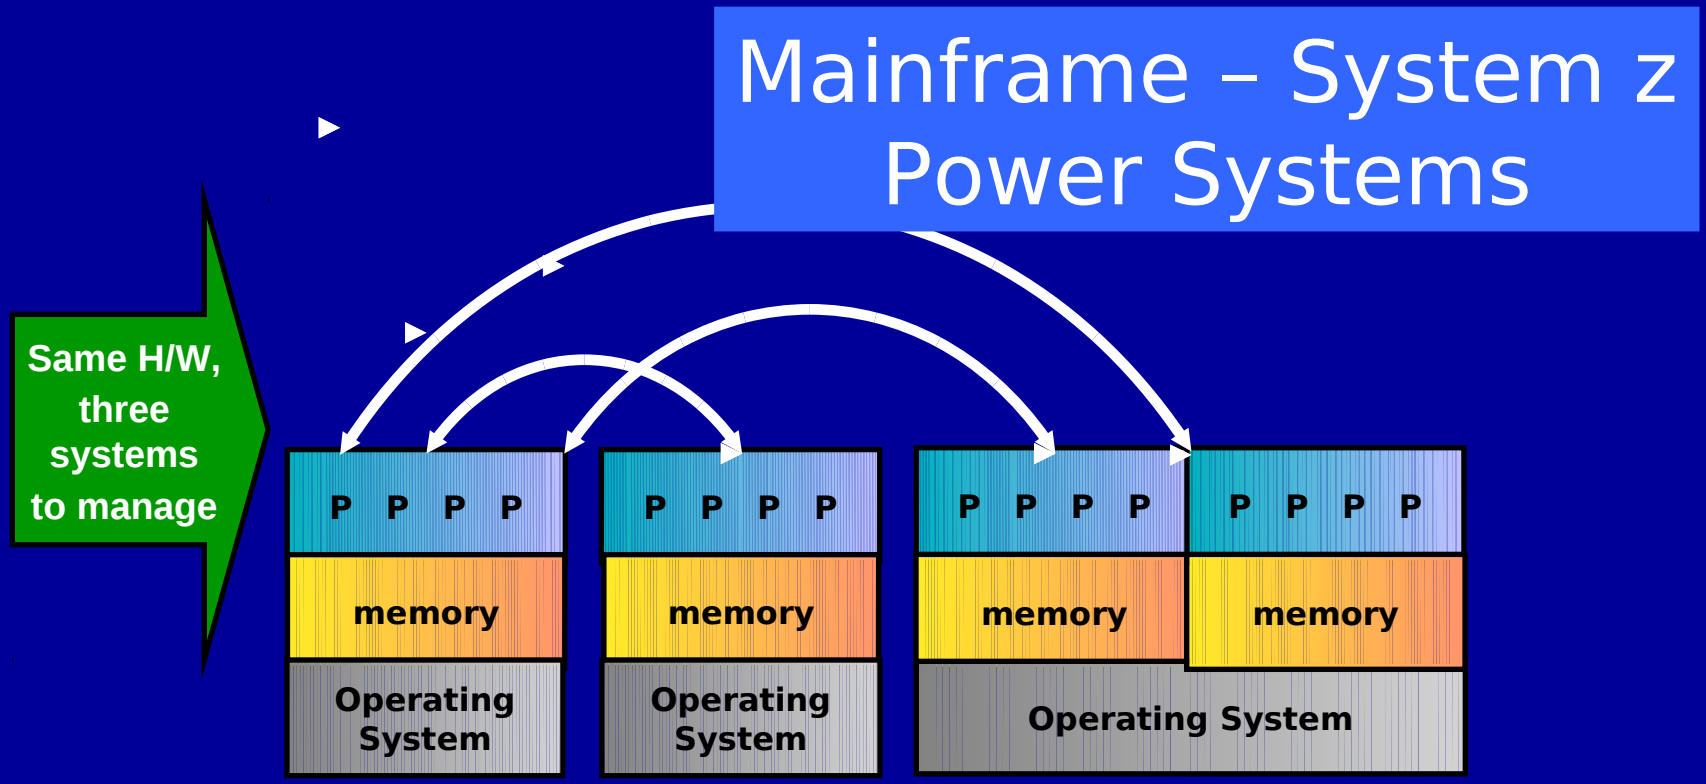

Virtualization

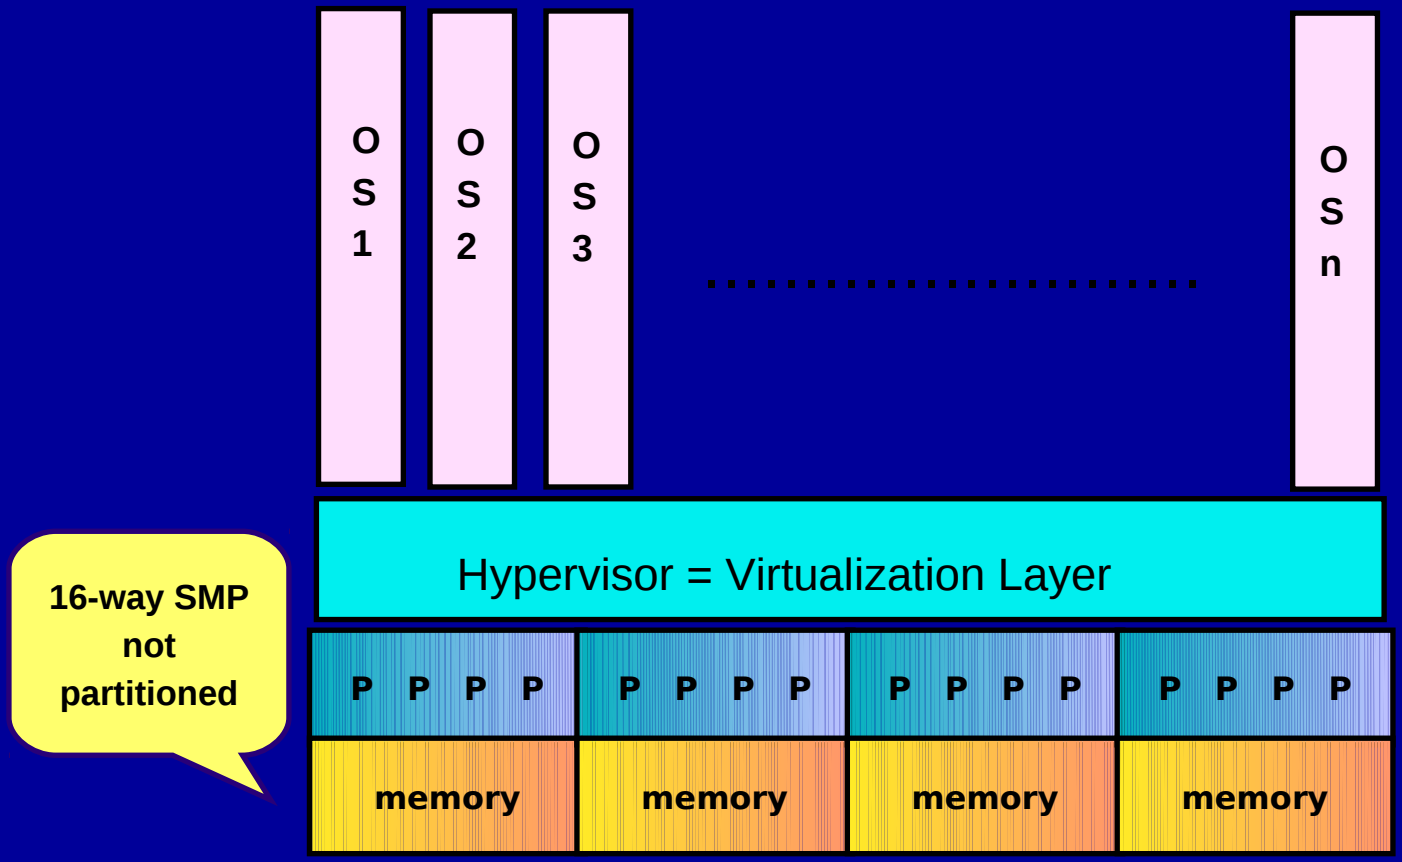

Virtualization

Mainframe – VM
x86 – VMWare, Xen, KVM, ...

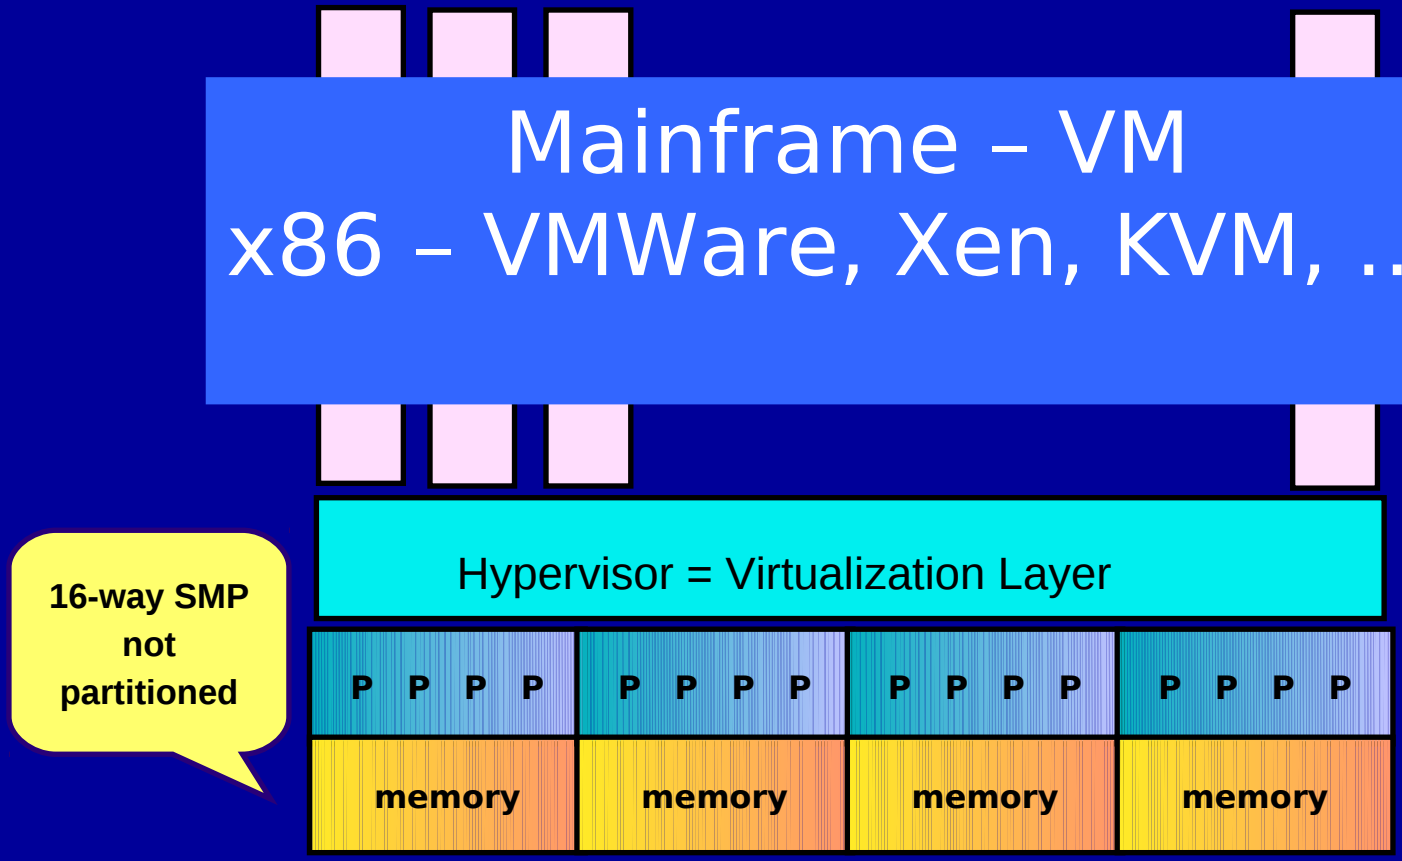

An Experiment Shows How Combining Workloads On A Shared Server Statistically Improves Utilization

- Group 1
  - Take 1 die and roll it a number of times. Count the number of times you get a 1, 2, 3, …. 6
  - Plot your results on a histogram

- Group 2
  - Take 9 dice and roll them a number of times. On each roll get the total value on the 10 dice. Count the number of times you get a 9,10, 11, 12, …., 30, 31, … 54.
  - Plot your results on a histogram.

- What do we see about the "predictability" of the result of a "roll"? (Let's roll the dice with a computer.)

After Rolling 1 Die (1 Workload)
The Distribution Would Look Like This

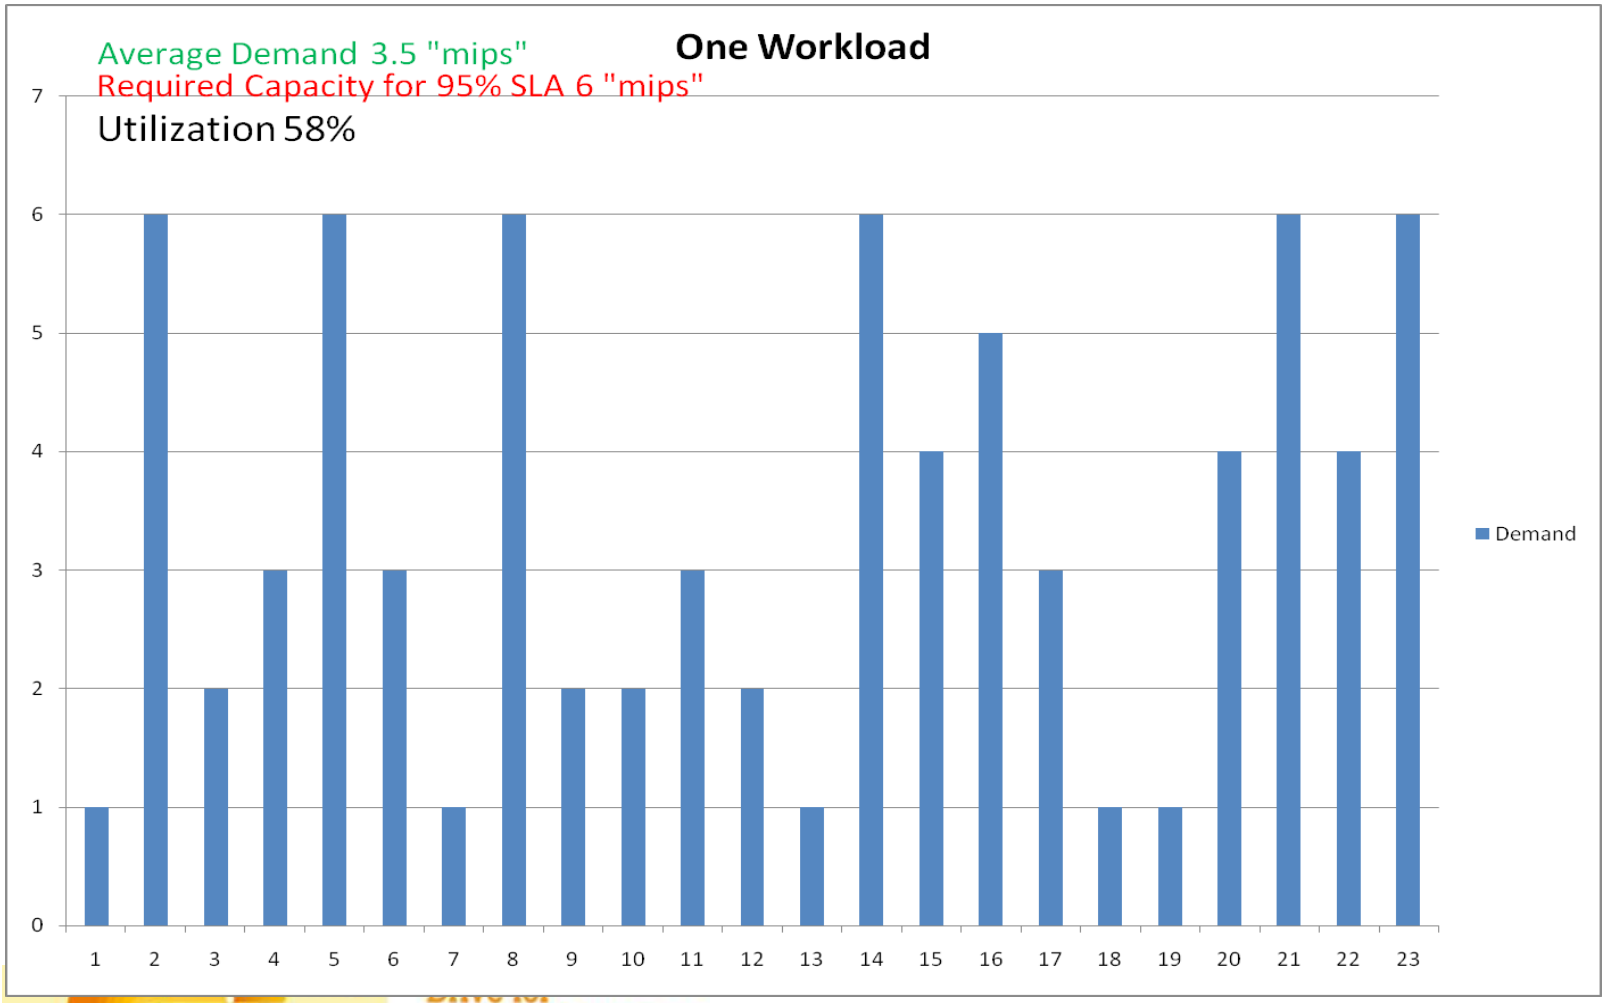

more effectiveness

After Rolling 9 Dice (9 Workloads) The Distribution Would Look Like This

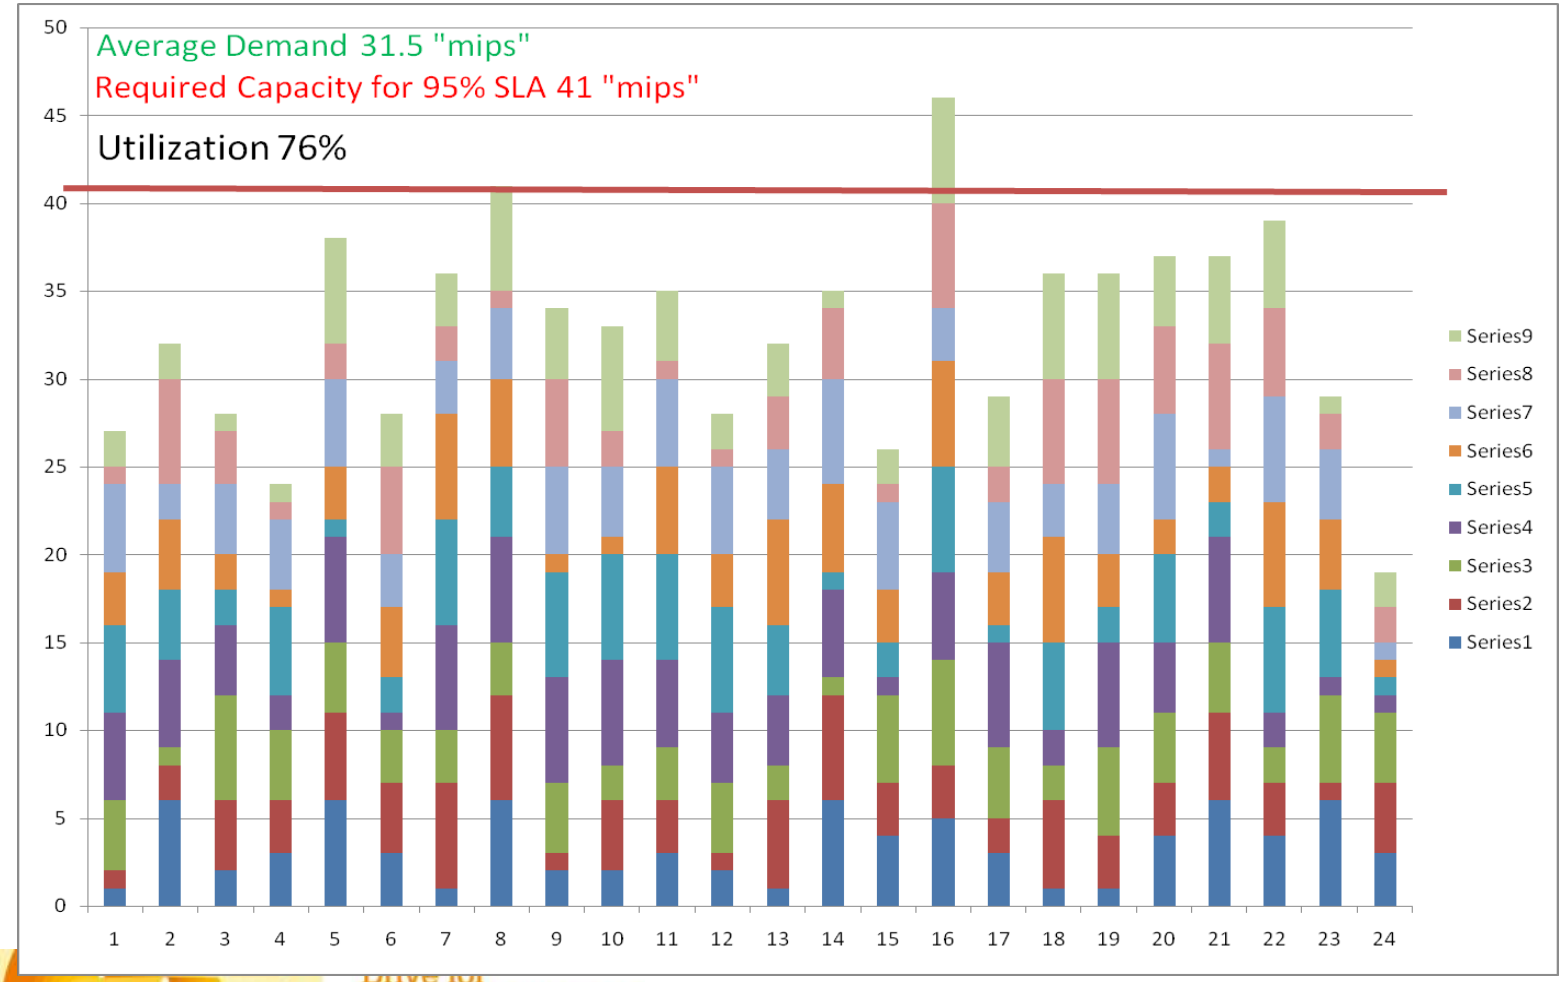

more effectiveness

After Rolling 100 Dice (100 Workloads) The Distribution Would Look Like This

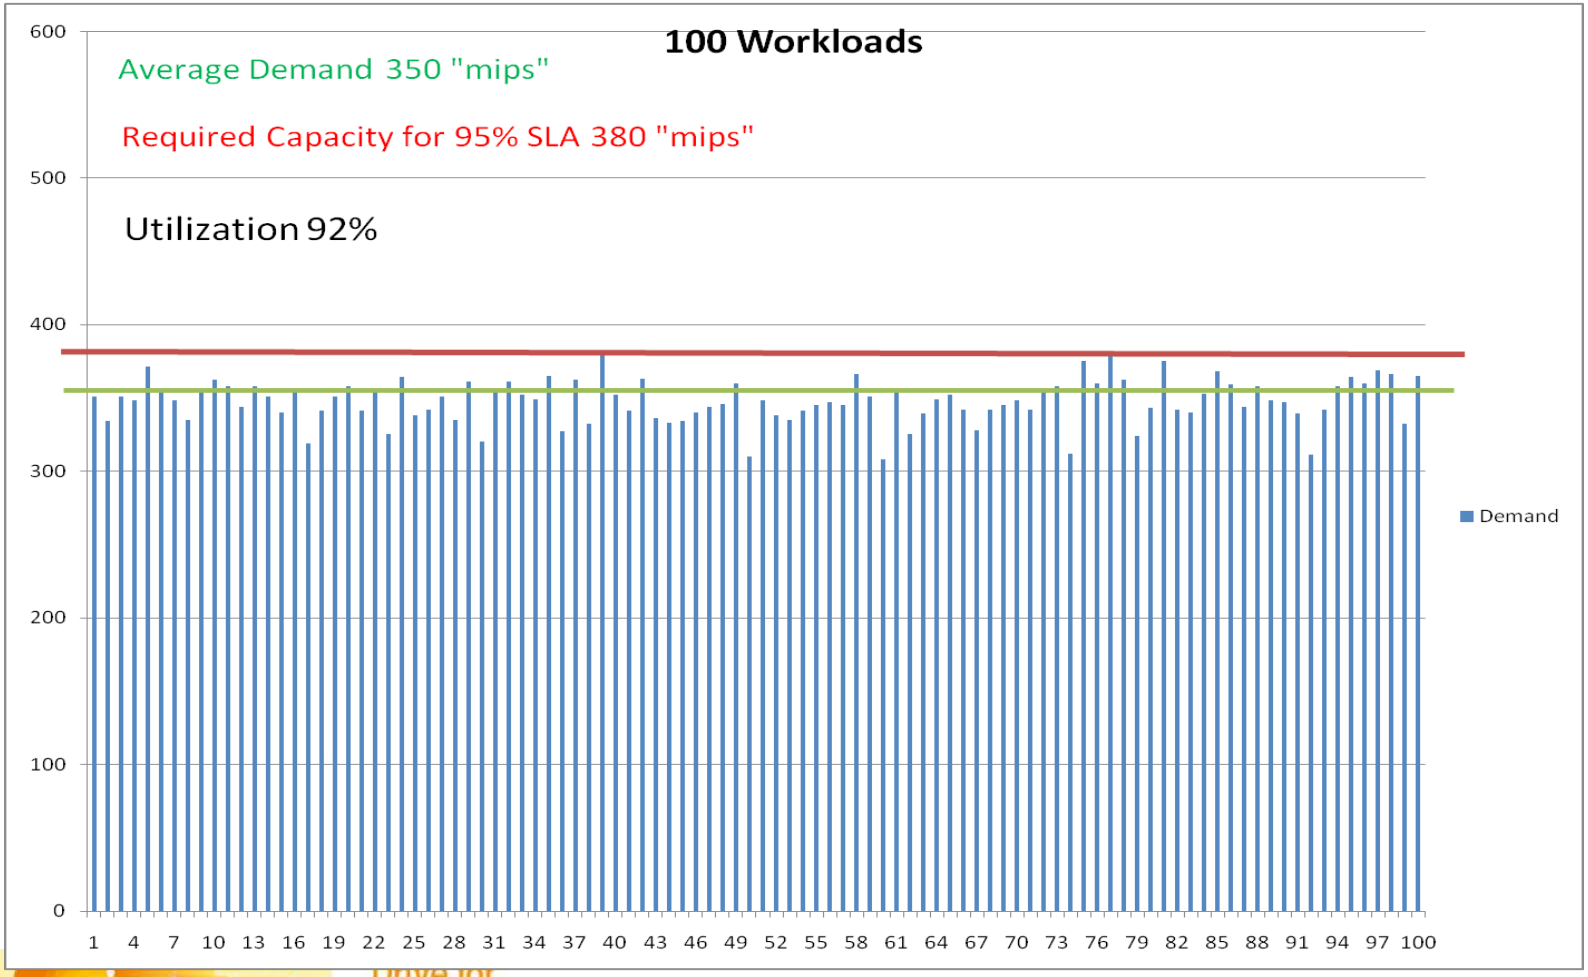

Multi-OS hosting

- ✓ Mainframe
  - ✓ Legacy
  - ✓ Linux
  - ✓ USS
- ✓ POWER
  - ✓ Legacy (AIX, i5/OS)
  - ✓ Linux
- ✓ x86
  - ✓ Win/xx
  - ✓ Linux

x86 Market

Vertical scalability

Bricks (2,4,8,16,32)
Modular servers

Classical
Rack servers

IBM ^ BladeCenter

Horizontal scalability

Trend : Consolidation , Virtualization and Multi-OS hosting Moves the market to Brick Servers and Blade Technology

VMware ESX Server

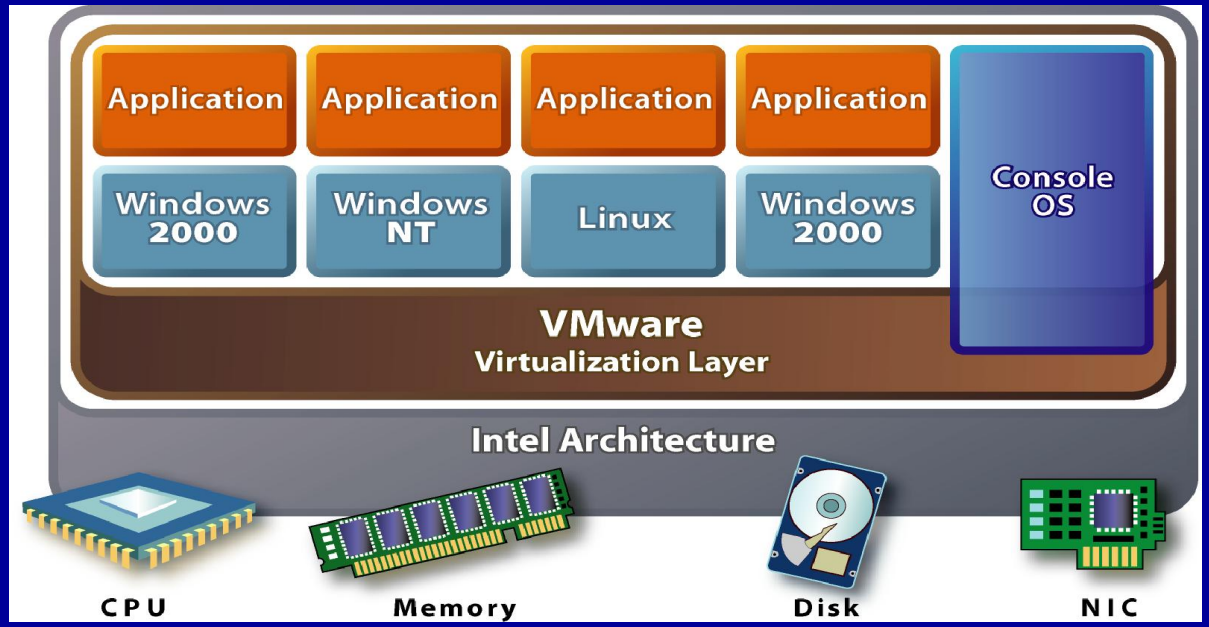

ESX Server facilitates multiple operating systems running concurrently on a single server. This lets users build and test networks of virtual machines on a single system.

What is a "Blade"?

A "server on a card" - each "Blade" has its own:

- ➔processors
- ➔ethernet
- ➔memory
- ➔optional storage
- ➔etc.

The chassis provides shared:

- ➔management console (KVM)
- ➔power supply
- ➔cooling
- ➔network switches
- ➔CD-ROM drive
- ➔diskette drive
- ➔etc.

IBM Blade - in its own ruggedized chassis

IBM Blade - with its cover on - ready for insertion into the BladeCenter

IBM BladeCenter chassis - 7U rackable - 14 Blades

Enabling an Integrated e-Business Infrastructure

1 **Web Solution (6 Blades)**
- **Caching appliance Blade**
- **Load balancing appliance Blade**
- ***Linux* Apache Blades**

2 **Collaboration Solution (3 Blades)**
- ***Windows* 2000 Domino Blades**

3 **Terminal Serving Solution (2 Blades)**
- ***Windows* 2000 Citrix MetaFrame Blades**

4 **File Serving Solution (2 Blades)**
- ***Novell* Netware V6 Blade**
- **Storage Blade**

5 **Spare Blade**

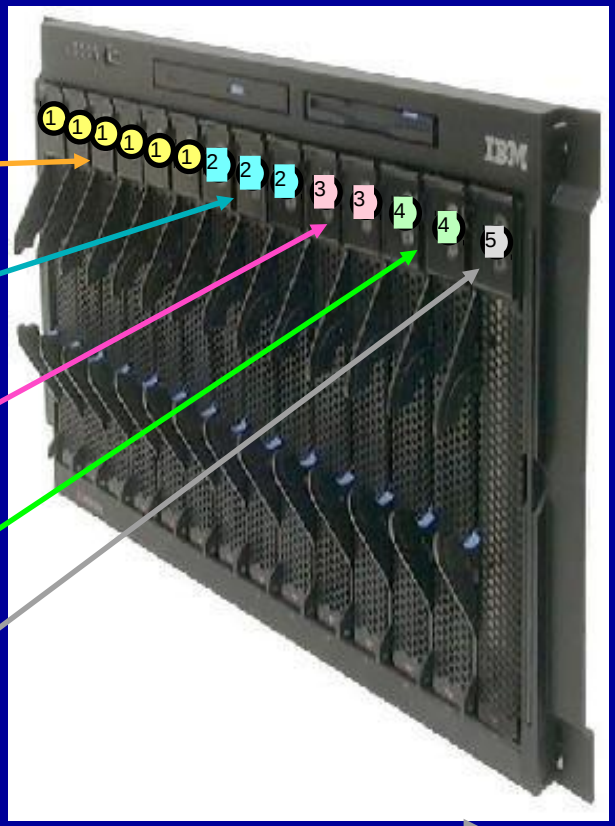

IBM ^ BladeCenter

Step 1 : Consolidate Servers

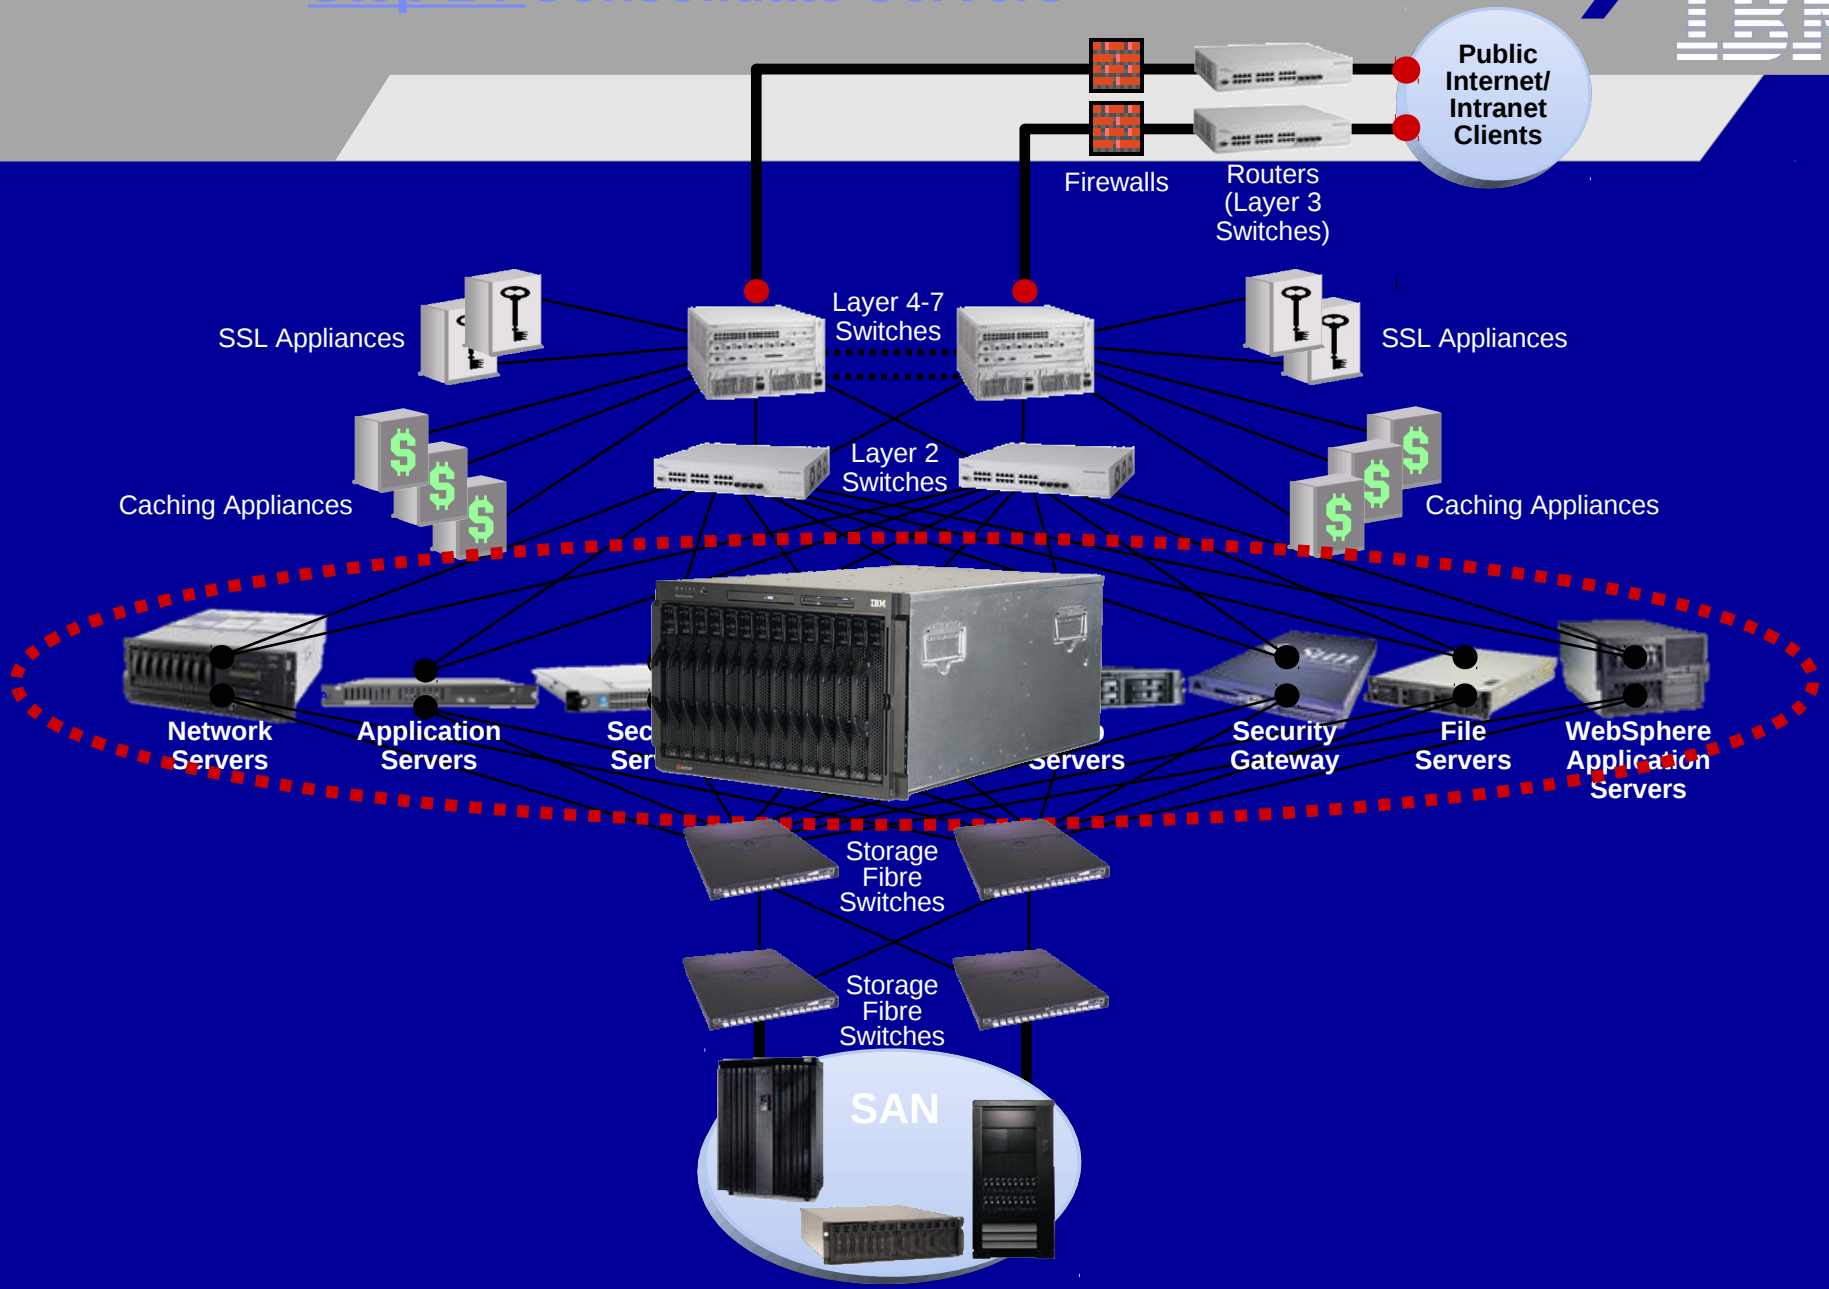

Step 2 : Integrate First Layer of the Network (L2)

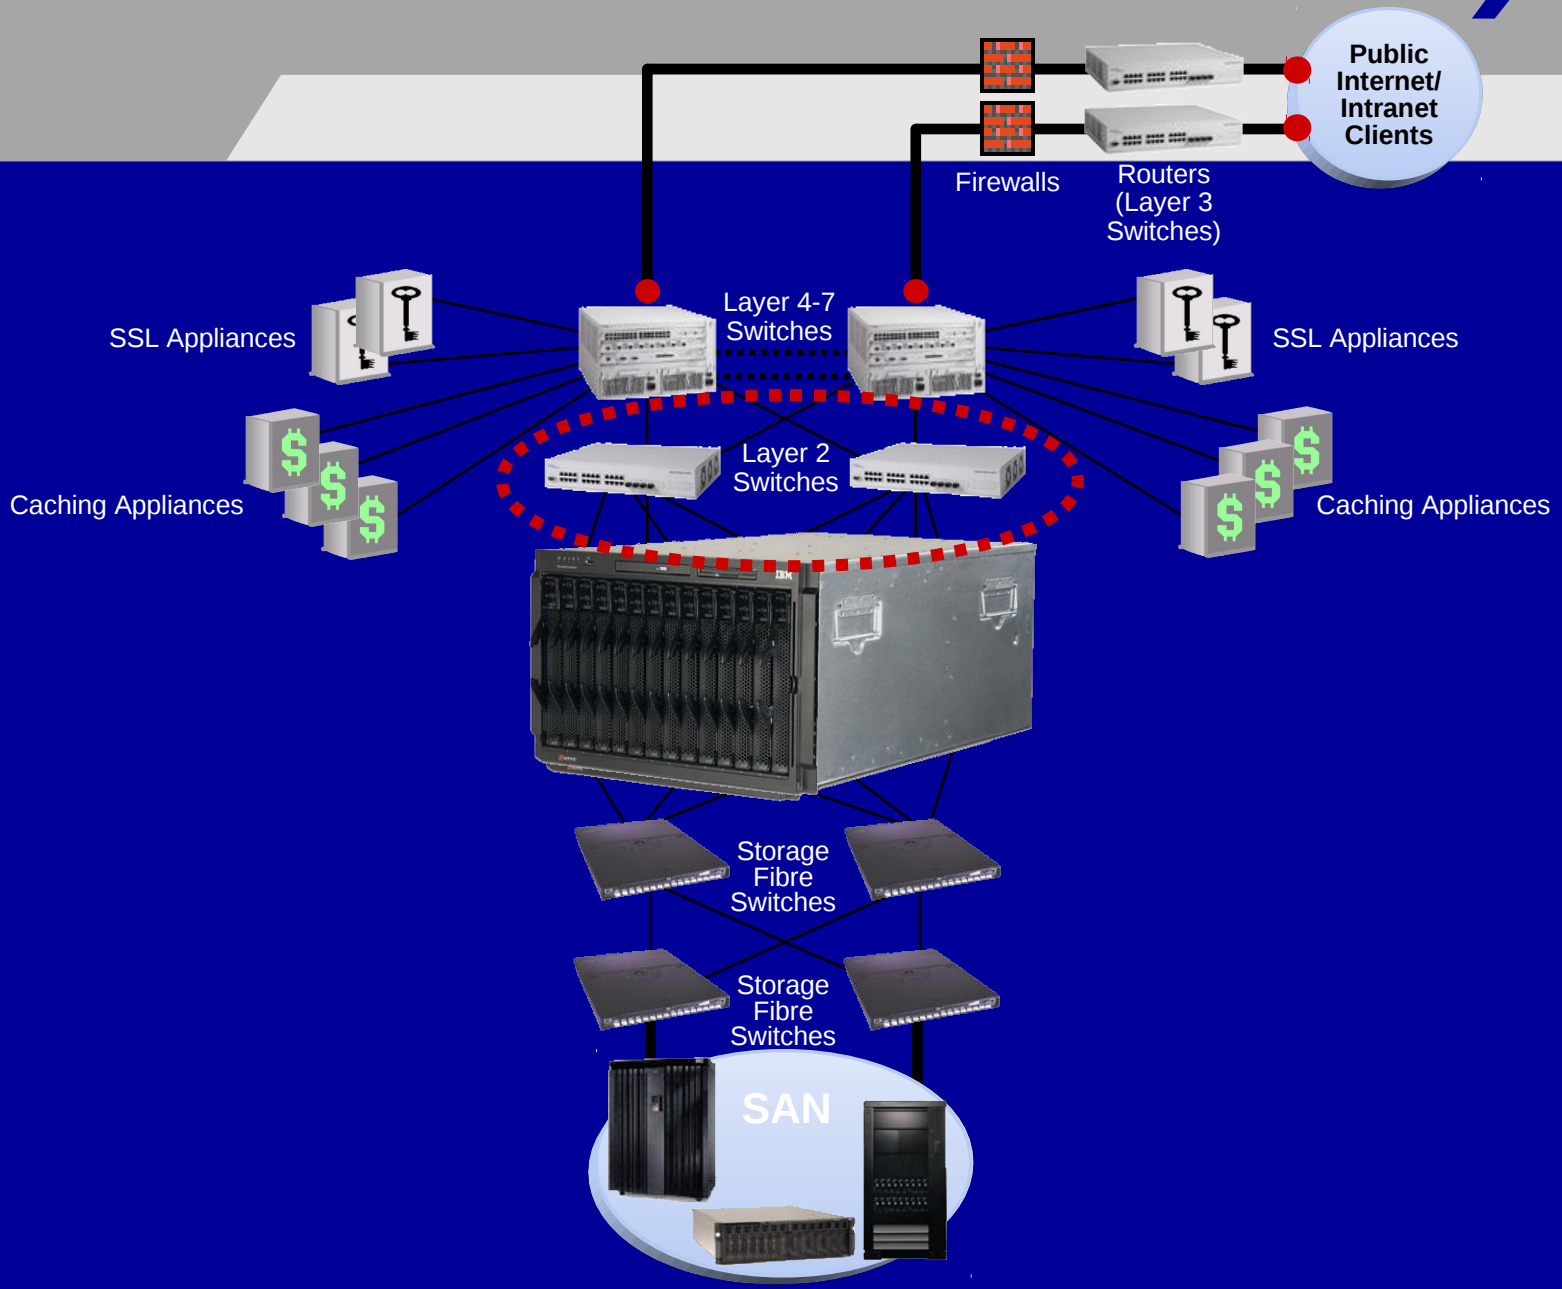

Step 3 : Integrate Storage Fabric

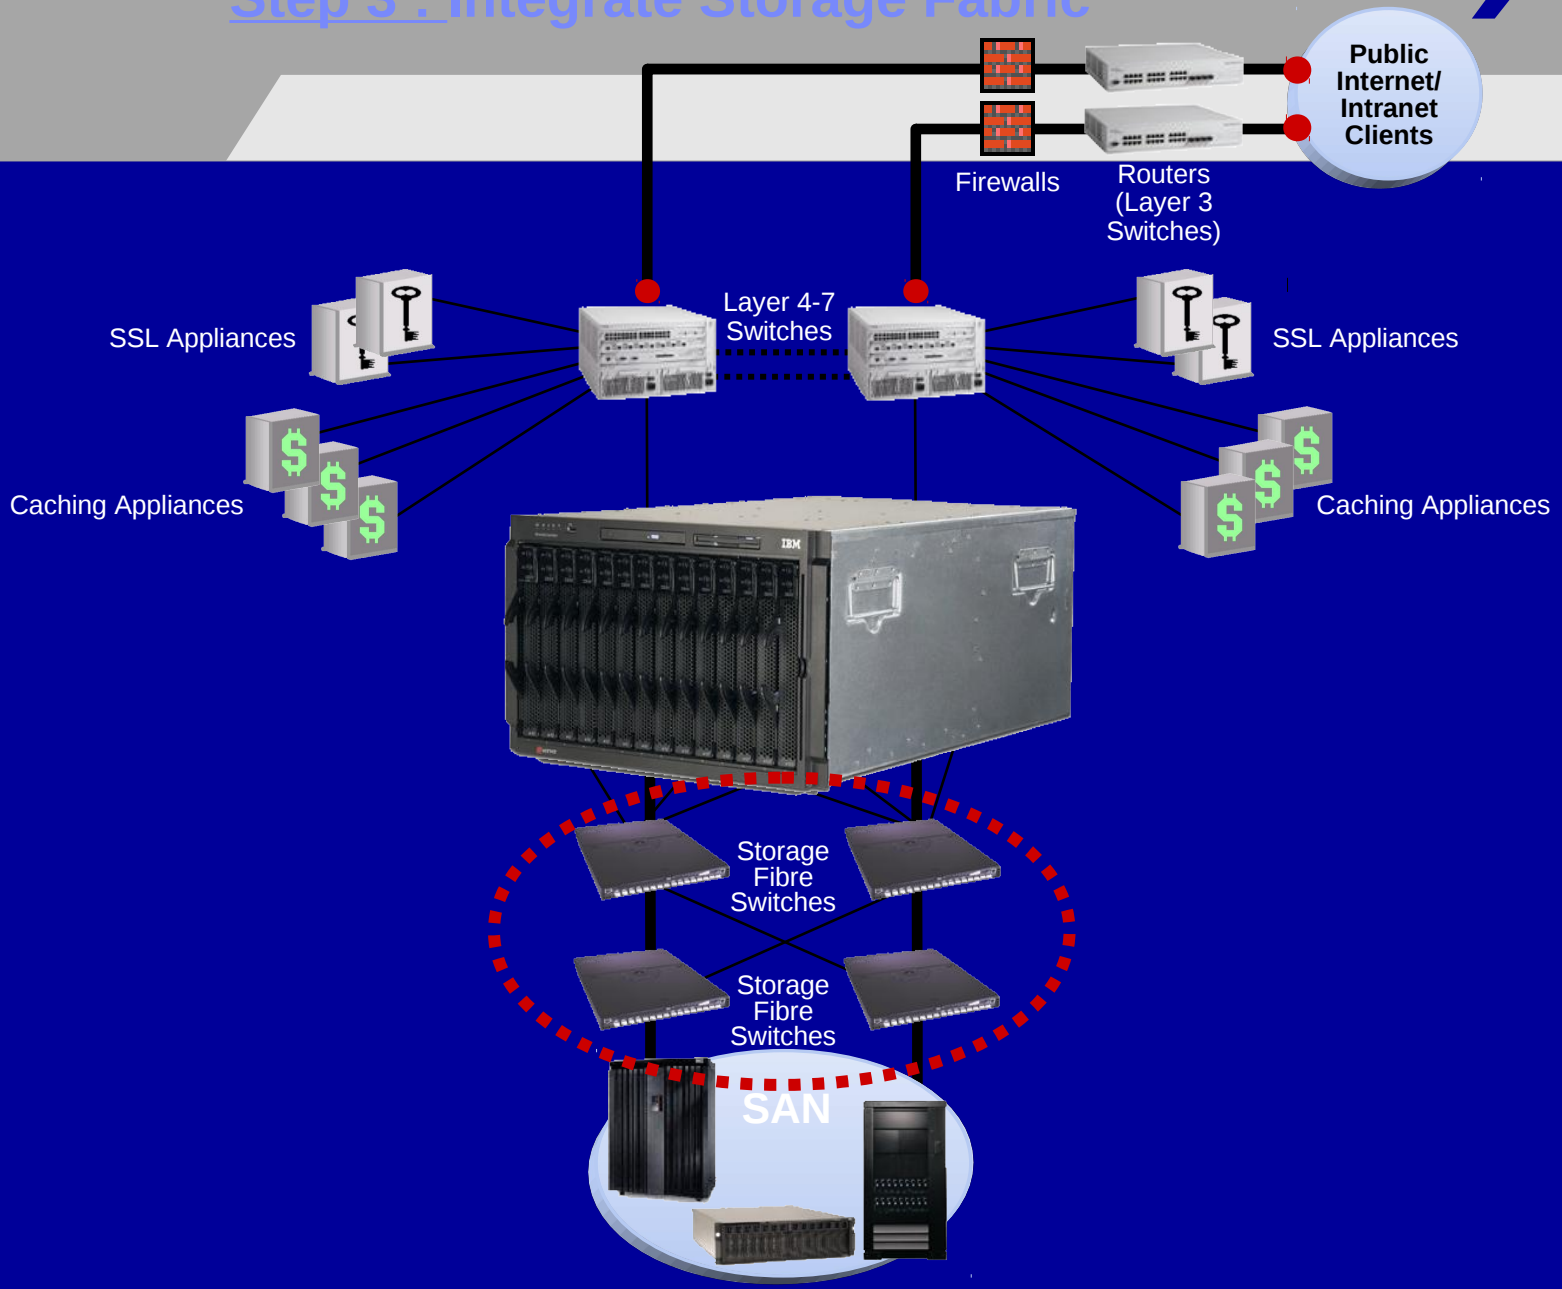

Step 4 : Integrate Second Layer of the Network (L4-7)

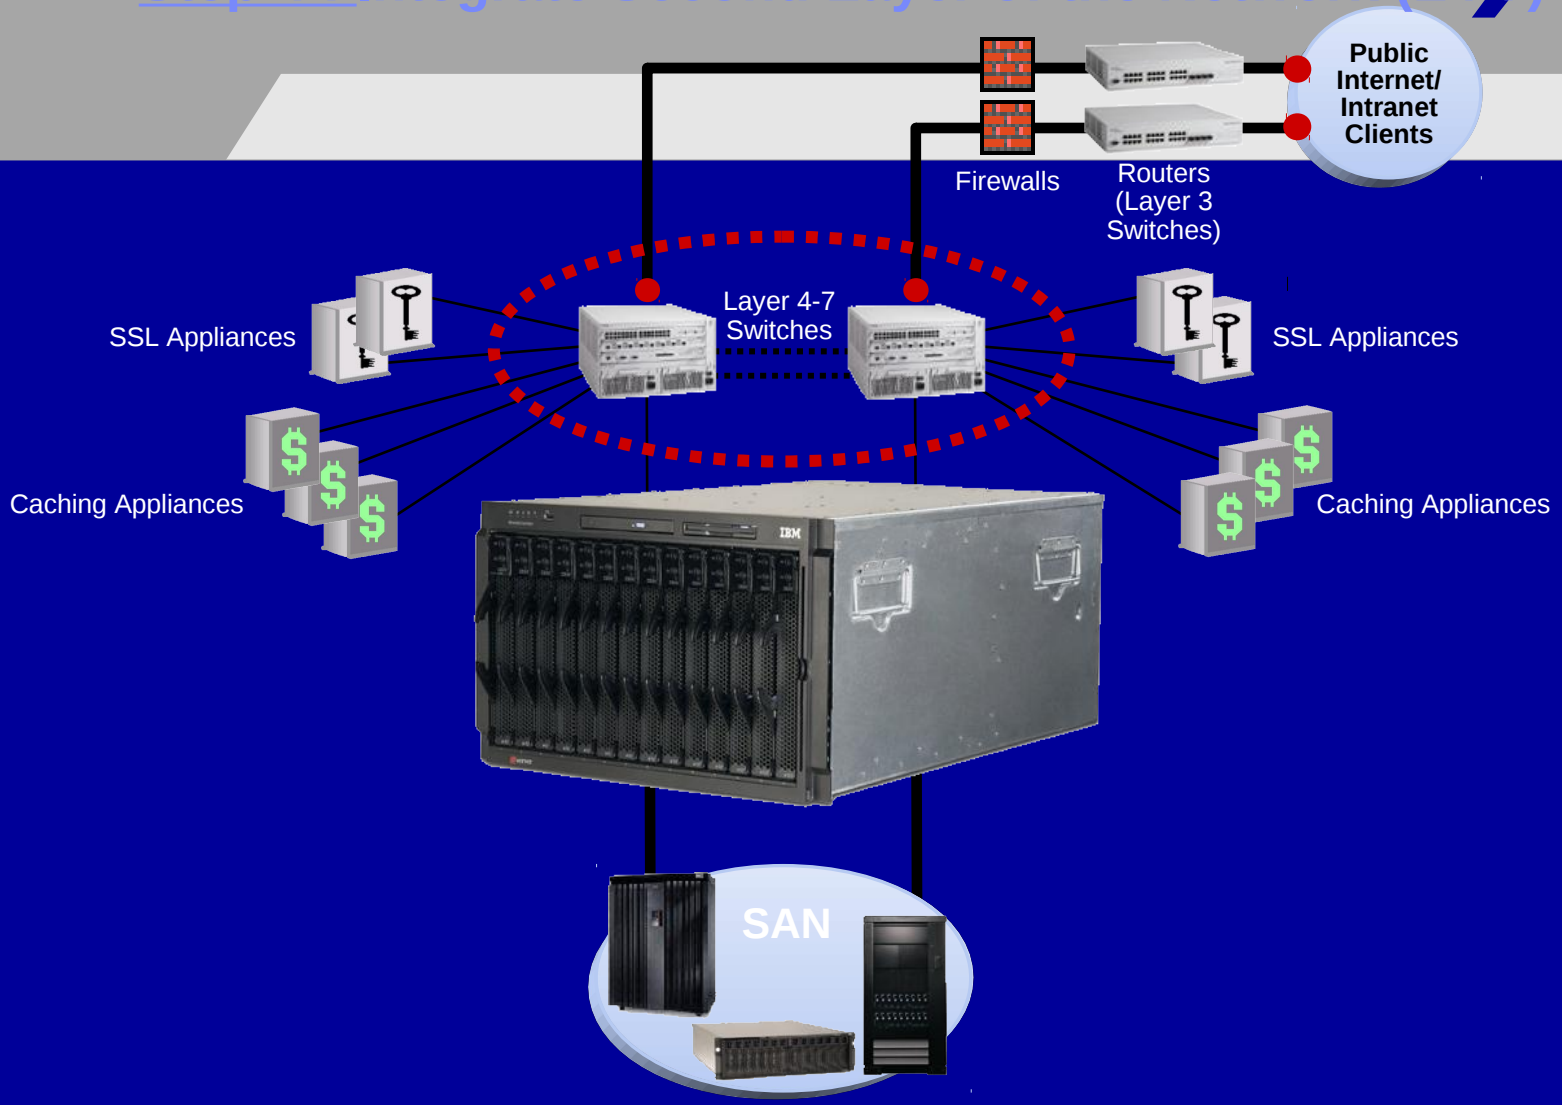

Step 5 : Consolidate Applications

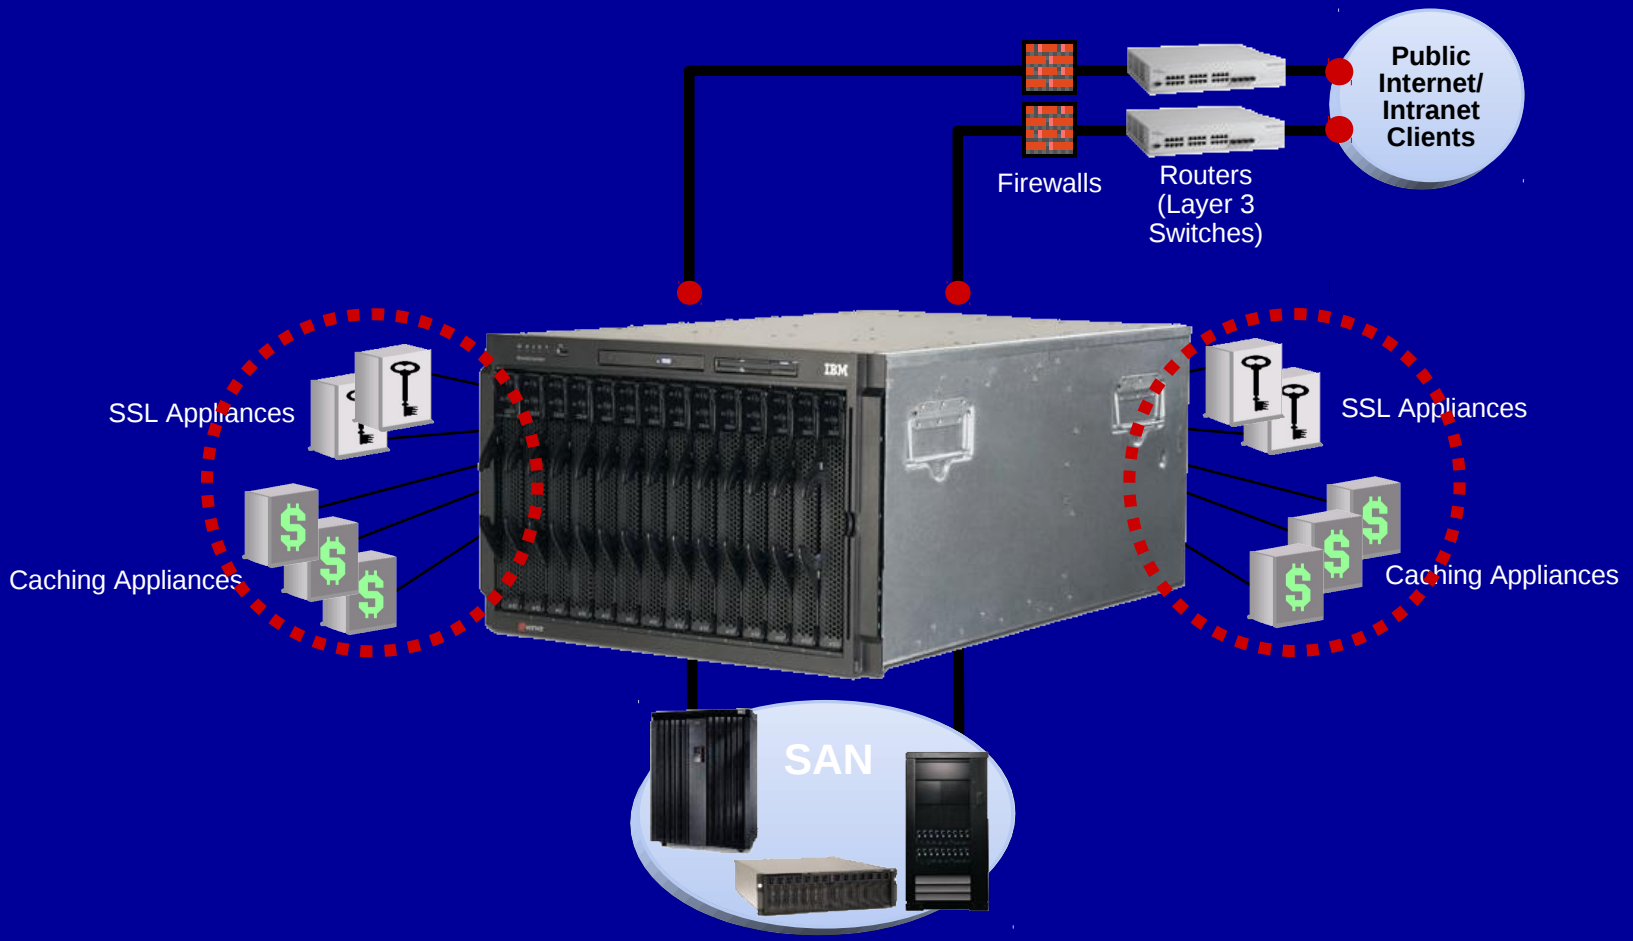

Result : BladeCenter Collapses Complexity

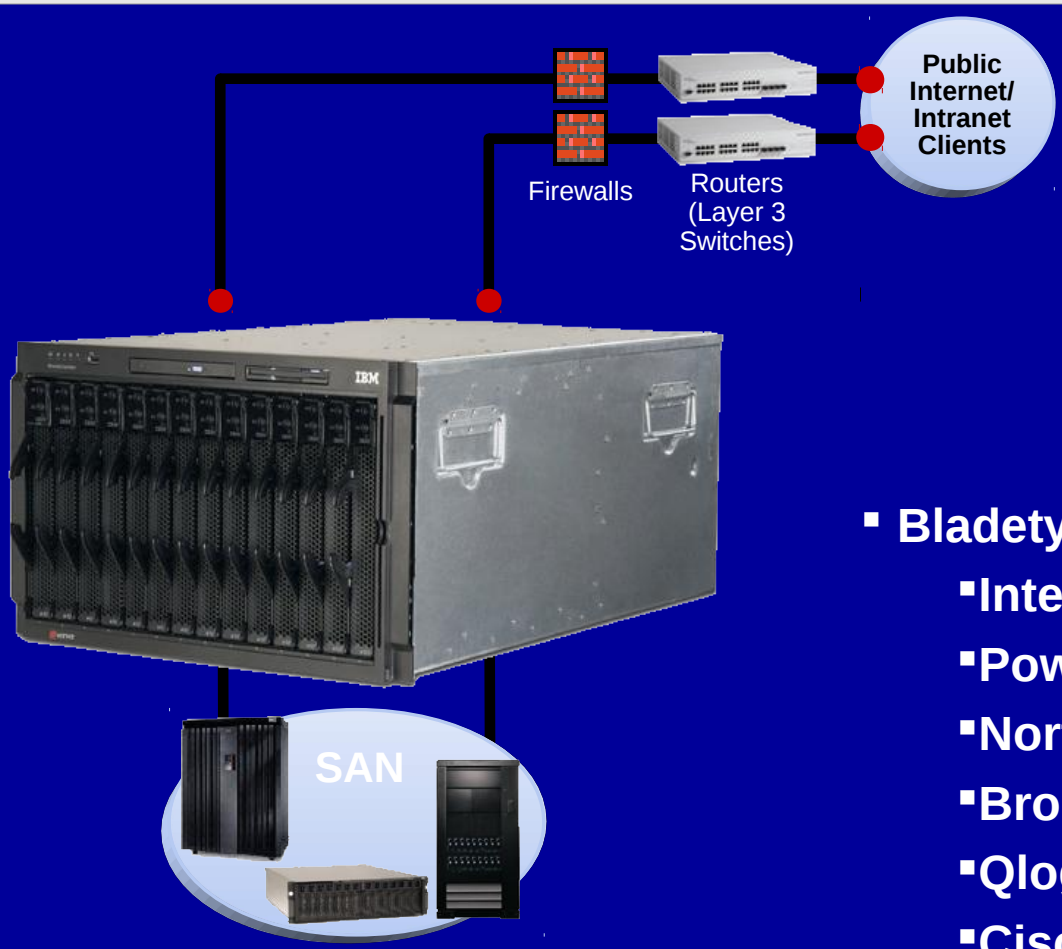

- Bladetypes :
  - Intel / AMD
  - Power
  - Nortel
  - Brocade
  - Qlogix
  - Cisco
  - Cell
  - …

BladeCentre – protection of investment

- Blade Lifecycle / Intel Server Technology Lifecycle ?
- Chassis Technology Lifecycle ?

Blade architecture value-adds

- Blade Lifecycle / Intel Server Technology Lifecycle ?

12 Months

- Chassis Technology Lifecycle ?

Blade architecture value-adds

- Blade Lifecycle / Intel Server Technology Lifecycle ?

12 Months

- Chassis Technology Lifecycle ?

  5 Years

Blade architecture value-adds

- Blade Lifecycle / Intel Server Technology Lifecycle ?

12 Months

- Chassis Technology Lifecycle ?

  5 Years

- ALL IBM Bladecenter Models house all Bladetypes
  - Protect Bladecentre investment
  - Blade manufacturers hate to redevelop their blades types every time a new chassis is developed

Reducing (un)scheduled downtimes

- The mainframe heritage

Continuous availability / highest RAS

- Continuous improvements in HW, microcode, OS and HA tooling for POWER & x86 platforms

Example – Live Partition Mobility

Move a running partition from one POWER server to another with no application downtime

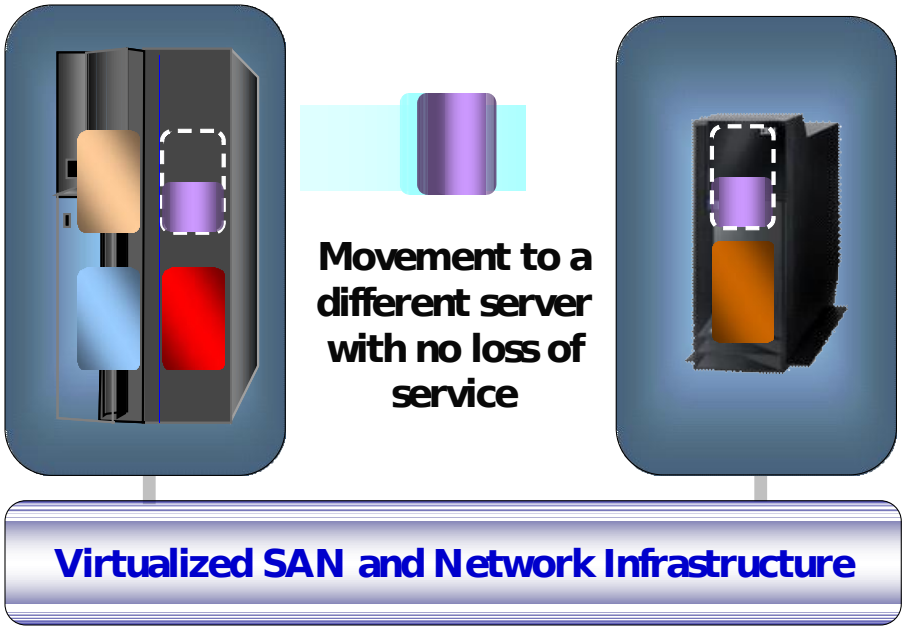

- Reduce planned downtime by moving workloads to another server during system maintenance

- Rebalance processing power across servers when and where you need it

Agenda

- ✓ Storage trends

Solid State Drives
Networked Storage
Virtualize everything
Store Less

Memory Hierarchy

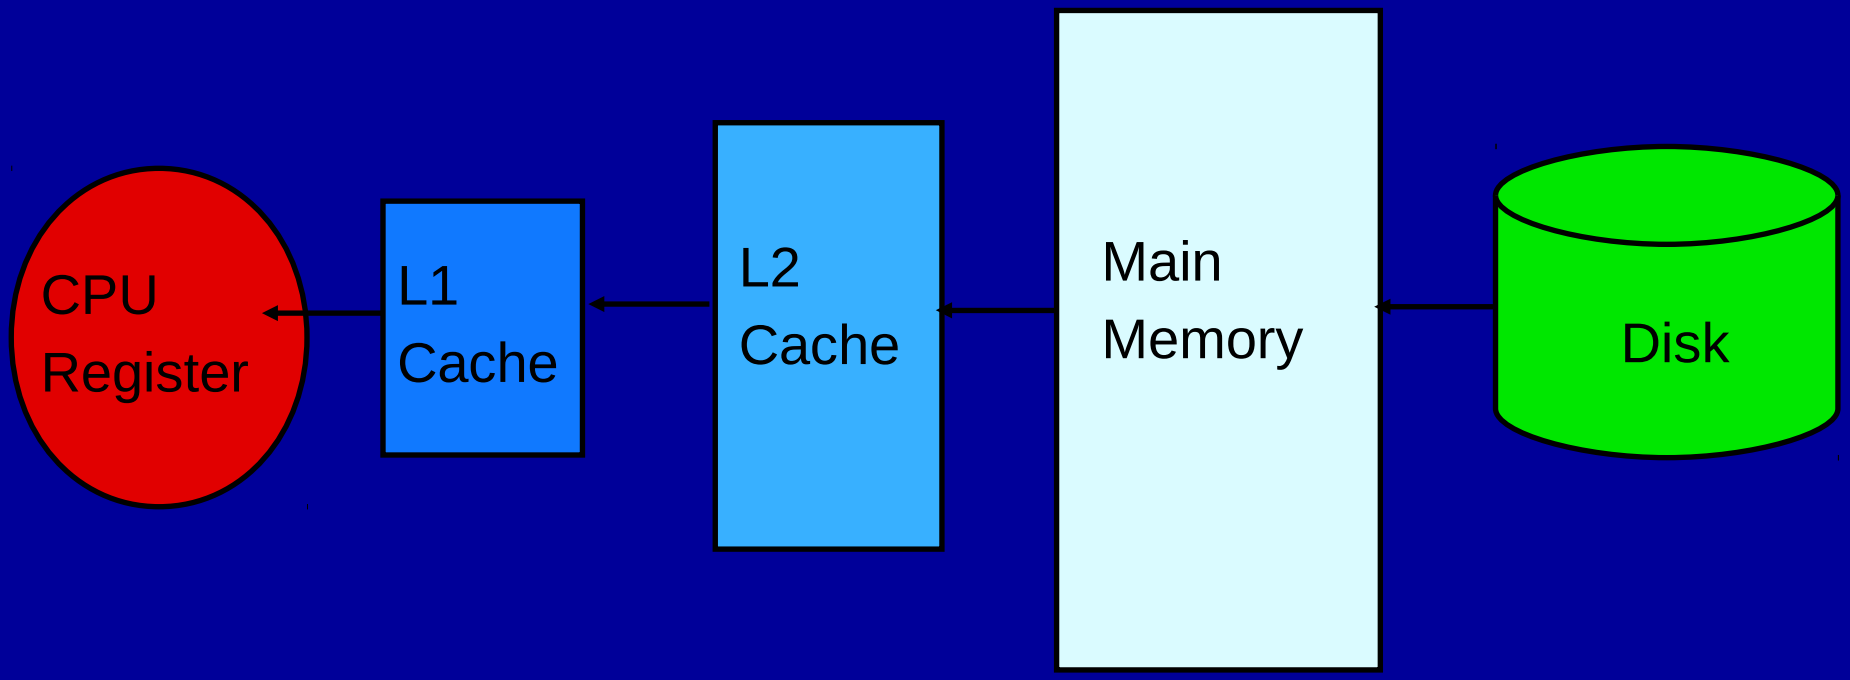

CPU Wait cycles	1-2	3-10	20-50	1000000+

Solid State Drives (SSD) Matching Applications' Need

- Today's applications can often benefit with a faster storage option
- SSD high speed can really help get rid of I/O bottlenecks, bridging the gap between memory and disk speeds
  - Improve performance
  - And save space, energy at the same time

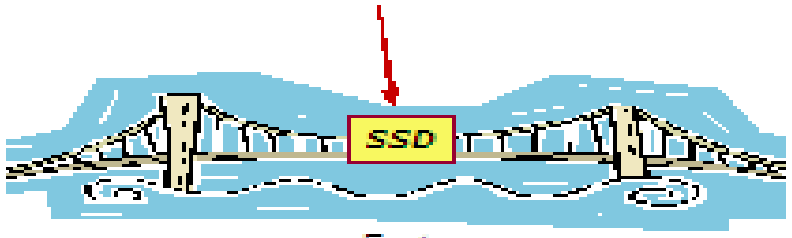

Processors	Memory		Disk
Very, very, very, very, very fast	Very, very, very fast		Very, very slow comparatively
≤ 10's ns	~100 ns	~200,000 ns	1,000,000 - 8,000,000 ns
Access Speed			

Storage evolution

From DAS (Direct Attached Storage)

to SAN (Storage Area Network)

or NAS (Network Attached Storage)

How can you connect to it ?

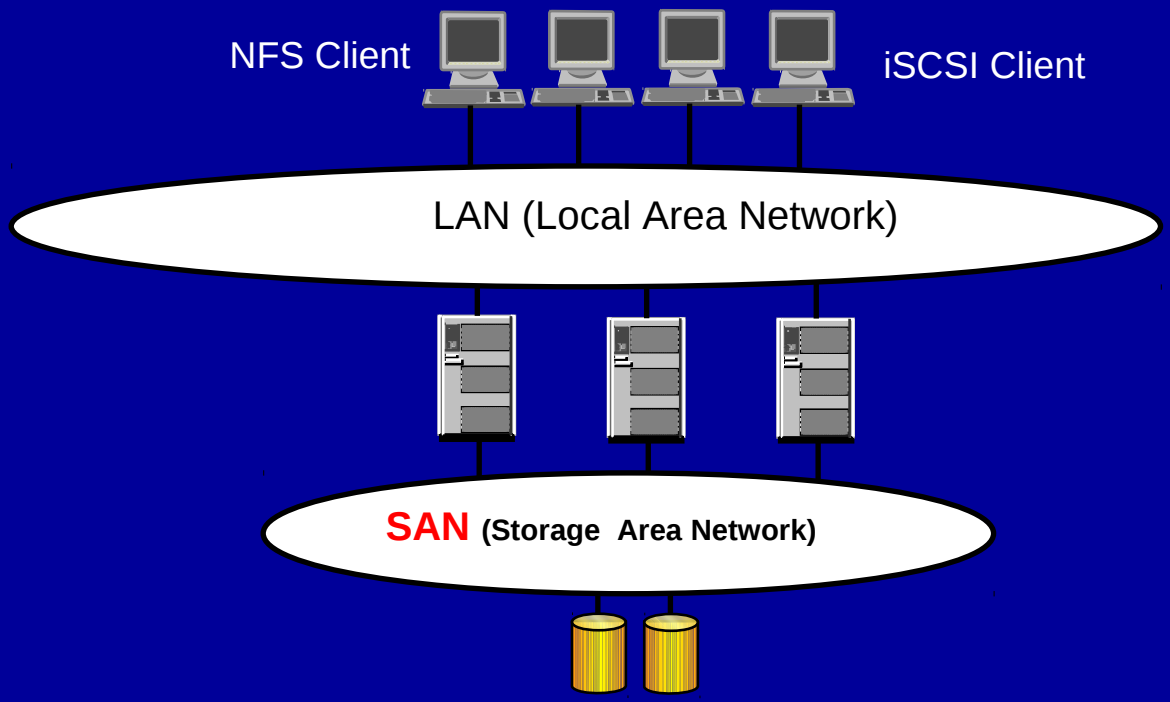

How can you connect to it ?

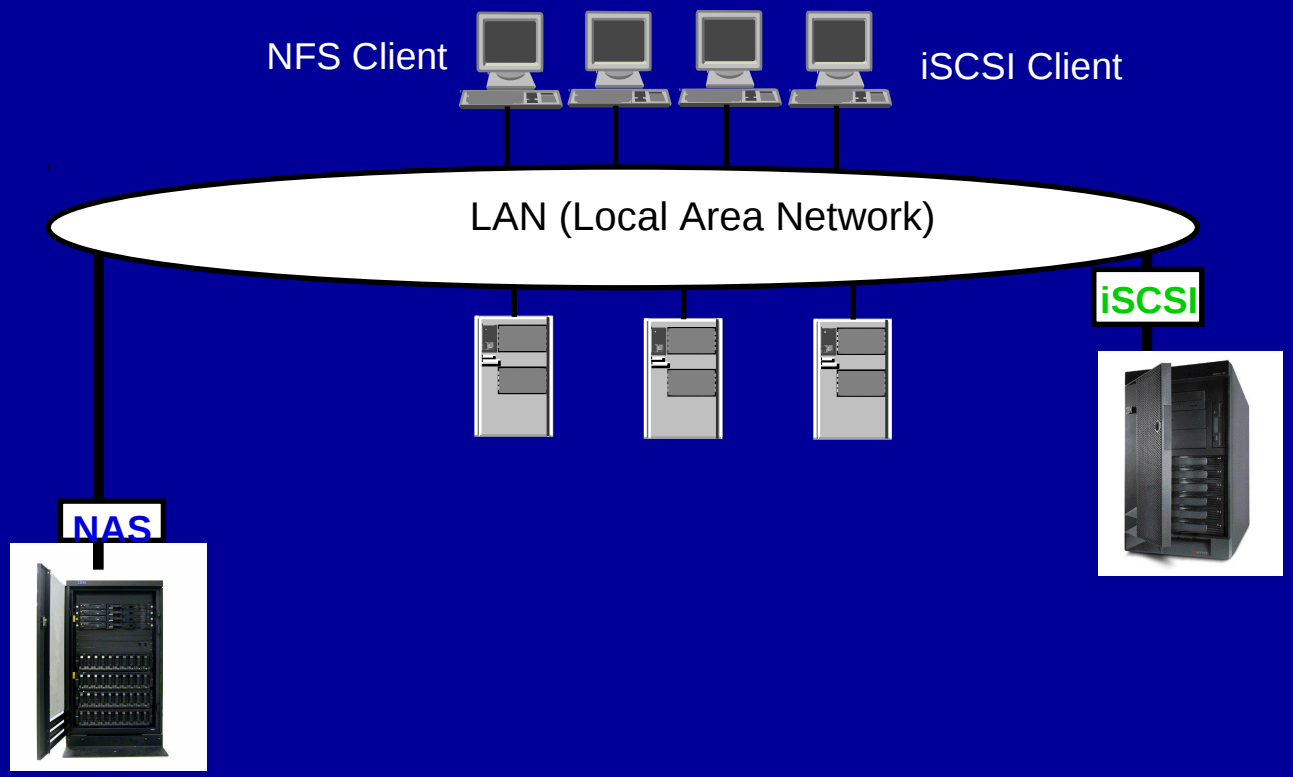

What are the main differences ?

SAN - high speed, scalable, multi-server, storage interconnect

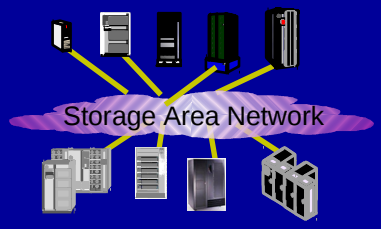

- ***Big Databases, OLTP, CRM, etc***
- ***Fibre Channel***
- ***Design and Installation time***
- ***Expensive***
- ***High IT support***

NAS

File Server

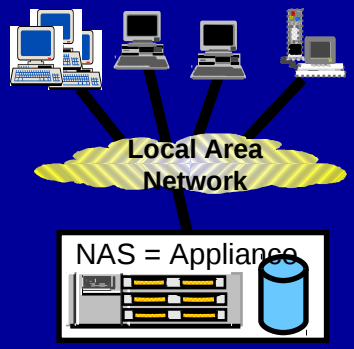

- ***Files only e.g .doc, .prz, .ppt, .nsf.***
- ***TCP/IP protocols***
- ***Quick to install***
- ***Cheap***
- ***Low IT support***

iSCSI - works like SAN storage but uses IP for the interconnect

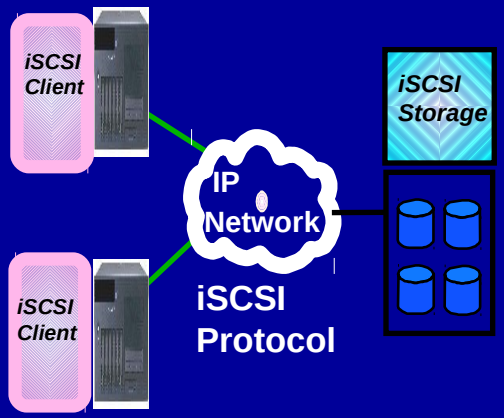

- ***Small databases***
- ***Low cost IP network***
- ***Quick to install***
- ***emerging business***

Storage Virtualization

- How do I simplify management of heterogeneous storage subsystems ?
- How can I improve the utilization of my Storage ?
- How can I continue operations during data movement between Storage subsystems ?

Flexible Storage Infrastructure with SAN Volume Controller

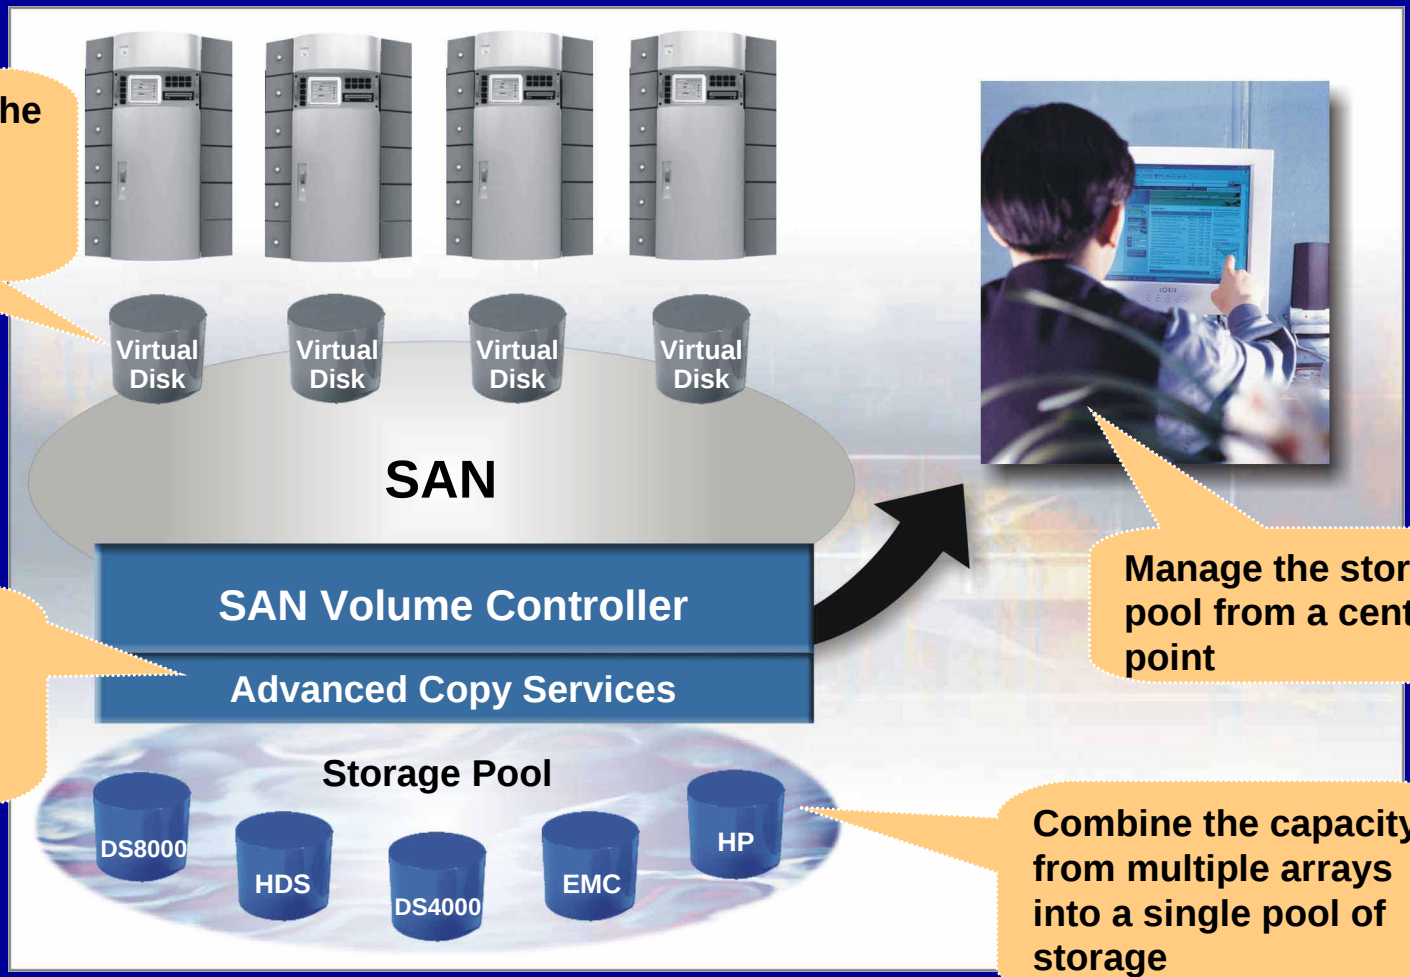

Infrastructure Simplification with SAN Volume Controller

- Traditional SAN
- **Capacity is isolated in SAN islands**
- **Multiple management points**
- **Poor capacity utilization**
- **Capacity is purchased for, and owned by individual processors**

- SAN Volume Controller
- **Combines capacity into a single pool**
- **Uses storage assets more efficiently**
- **Single management point**
- **Capacity purchases can be deferred until the physical capacity of the SAN reaches a trigger point.**

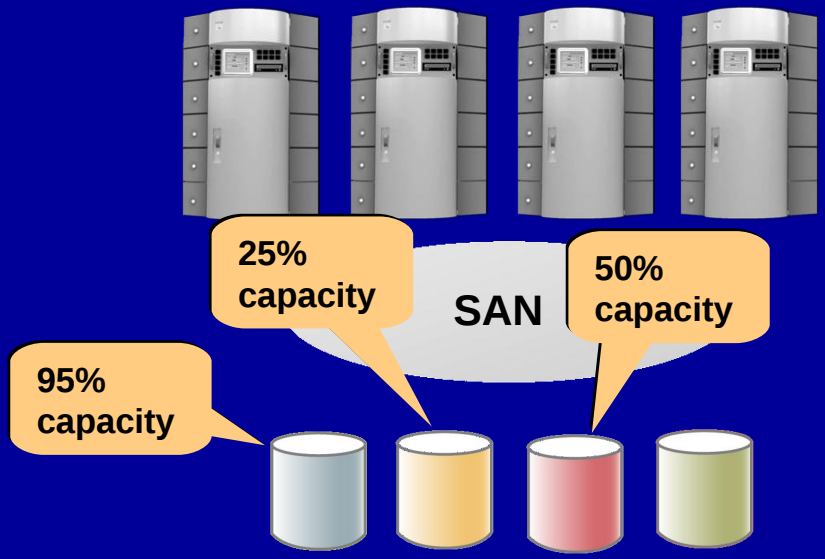

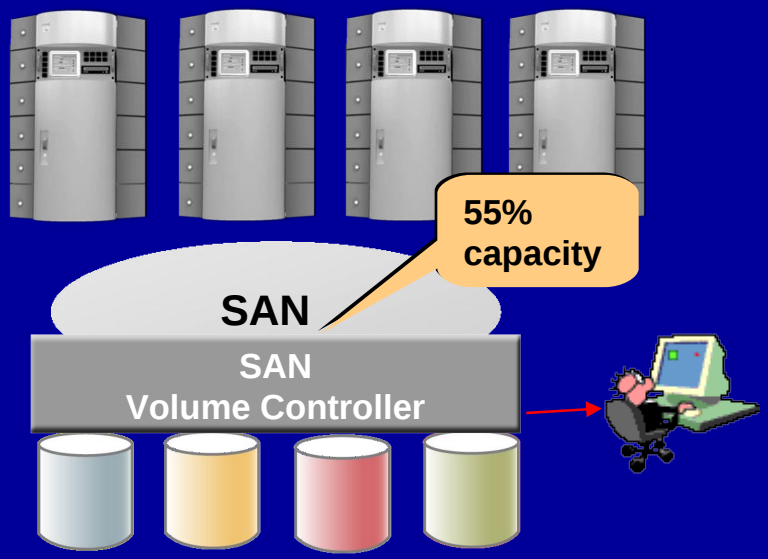

Non-disruptive Data Migration with SAN Volume Controller

- Traditional SAN
- **Stop applications**
- **Move data**
- **Re-establish host connections**
- **Restart applications**

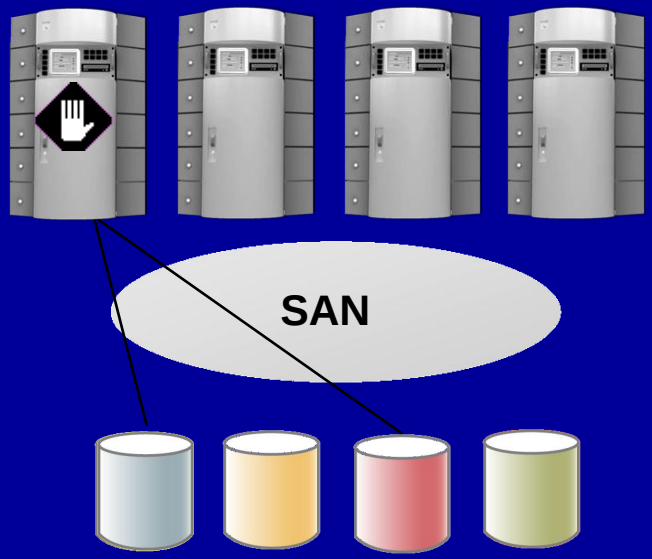

- SAN Volume Controller
- **Move data**
- **Host systems and applications are not affected.**

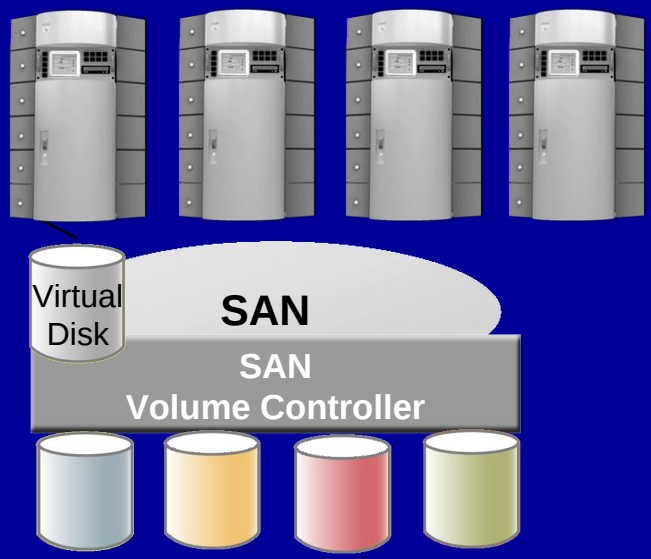

- Improving Storage utilization with Data Deduplication

Protect More. Store Less.™

ProtecTIER reduces the required backup disk capacity by

up to 25 times or more!

Drive for more effectiveness

Improving Storage utilization with Data compression

- **Compresses data on the initial write**
  - Less data stored on primary storage (50% to 80%)
  - Primary data compression shrinks backups, too
  - Transparent to applications and processes
- **Compresses data *before* writing to storage arrays**
  - Storage cache holds more data (50% to 80%)

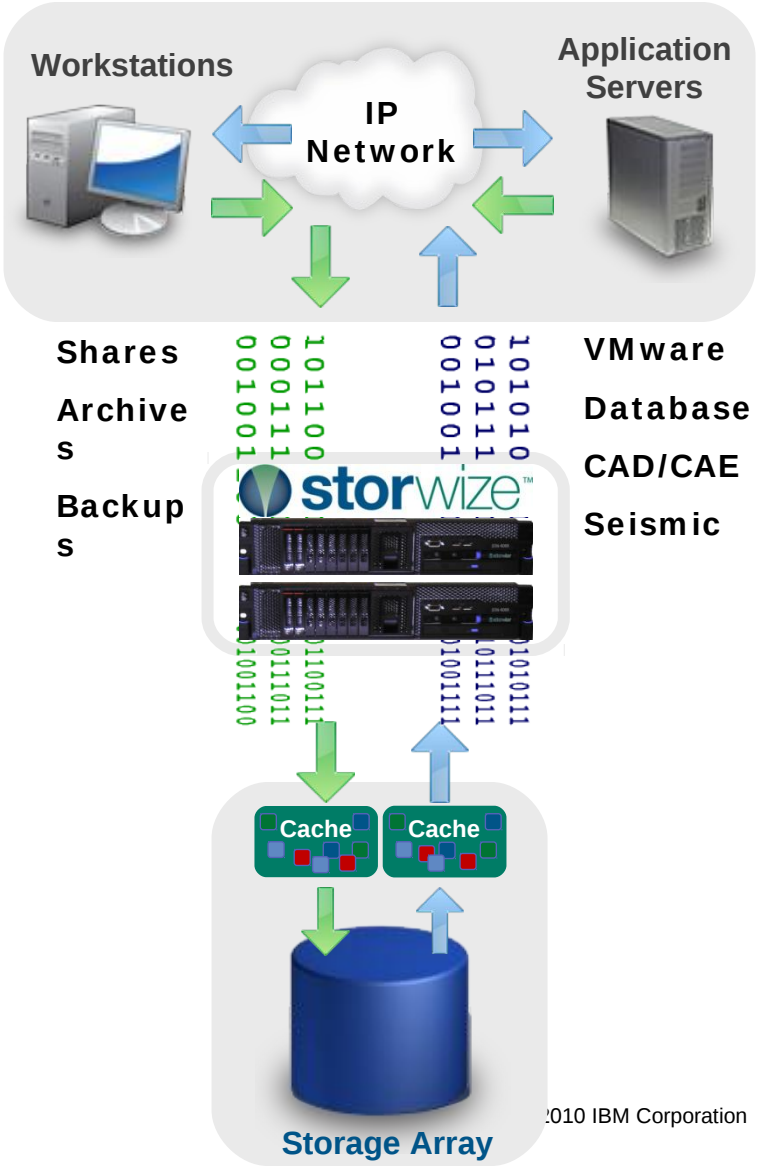

Drive for more effectiveness

Summary

- Server evolution
  - Virtualization &consolidating
  - multi – OS support
- Networked Storage and Solid State Drives
- Storage virtualization reduces complexity