

LARGE SCALE INFRASTRUCTURE

2110414 Large Scale Computing Systems
Natawut Nupairoj, Ph.D.

Outline

2

- Overview
- Hardware Virtualization
- Storage Technology

3

Overview

Trends in IT Management

4

- Server performance and storage size grow very rapidly
- Equipment become much cheaper
- Some applications exhibit “seasonal” workload demands
- Lead to server and storage consolidation



Adapted from

- P. Strassmann, “Introduction to Virtualization”, George Mason University, 2008
- M. Behrens, “Virtualization Assessment”

Server Consolidation

- Old applications rely on many servers
 - ▣ High operation cost: maintenance, electricity, etc.
 - ▣ Heterogeneous environments
 - ▣ Difficult to migrate
- New servers are very powerful and under-utilized
 - ▣ Some resources remain idle
- Reduce costs by consolidating servers

Before Server Virtualization: Physical Isolation



Virtualization Concept


7




Capacity Utilization: Stand-Alone vs. Virtualization Servers

8

Dedicated Server



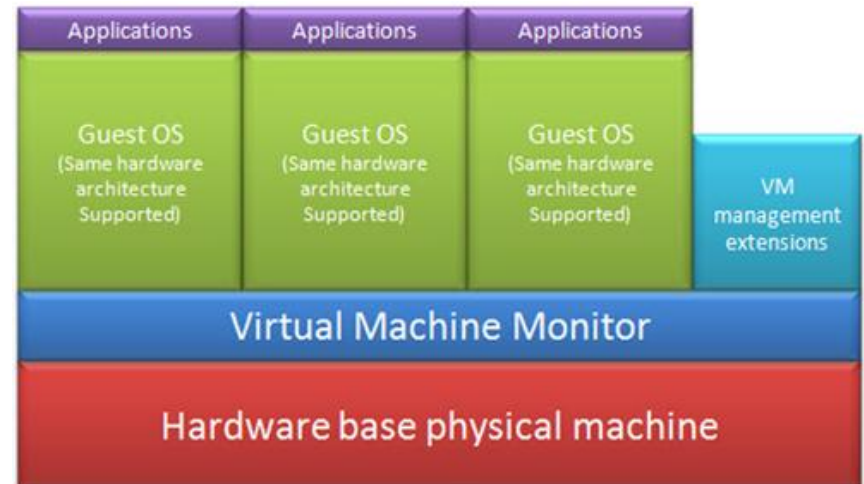
Virtualized Servers



Virtualization Approaches

9

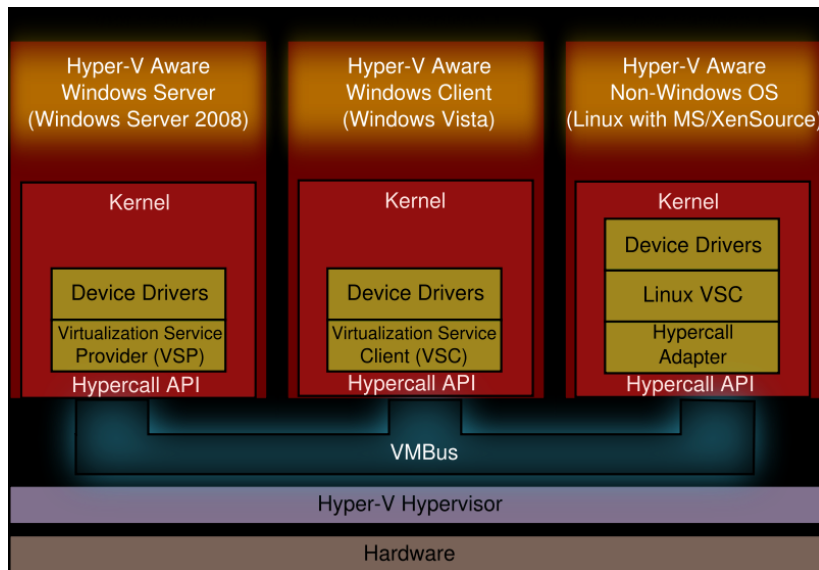
- Hardware-assisted virtualization
 - Require hardware support e.g. CPU special instruction sets
 - Accelerated virtualization, hardware virtual machine, native virtualization
- Full Virtualization
 - Guest OS is unaware of being virtualized
 - Required a special software called “Hypervisor” or “Virtual Machine Monitor” to manage the virtualization
 - May or may not required hardware support



Virtualization Approaches


10

- ParaVirtualization
 - Host OS (Hypervisor) provides a special Hypercall API to perform some functions for Guest OS
 - Guest OS kernel must be modified to utilize these APIs
 - Host OS cannot touch Guest OS directly



Benefits of Virtualization

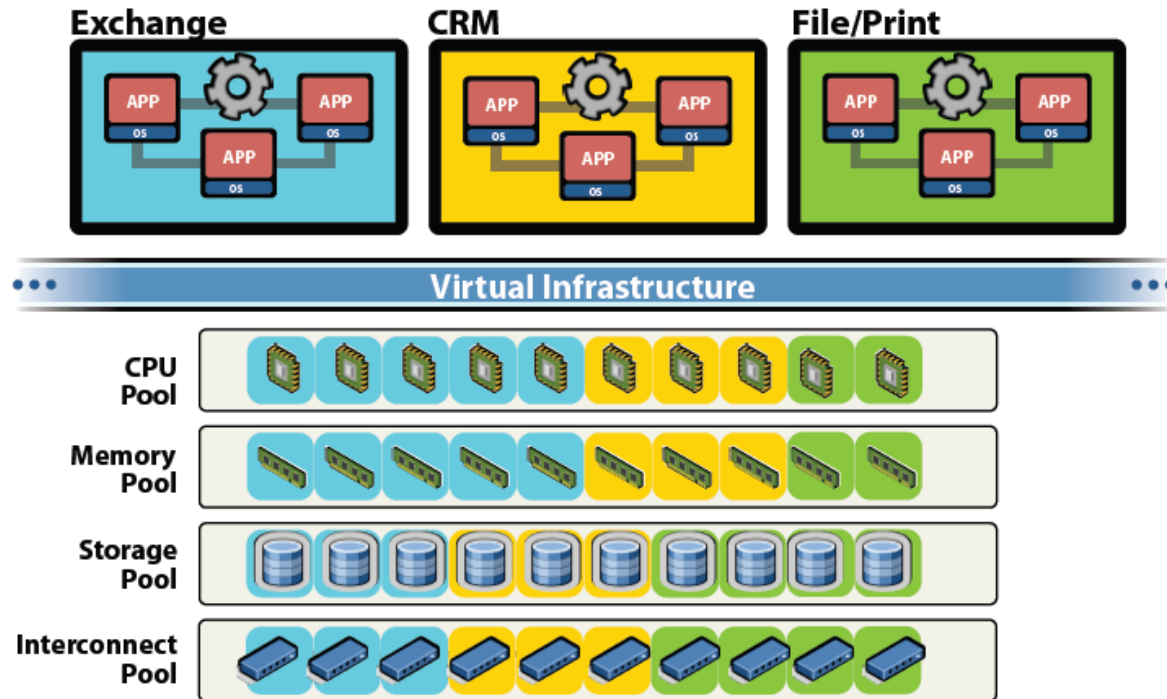
11



- Flexible Resource management
 - ▣ Server consolidation
 - ▣ Dynamic resource sharing
 - ▣ Reduce power consumption
- Simplify maintenance
 - ▣ Zero downtime maintenance
 - ▣ Live migration
 - ▣ Patch management
 - ▣ Efficient recovery

Resource Consolidation


12



- ❑ Reduce number of physical servers
- ❑ Decrease power consumptions
- ❑ Resource pooling
- ❑ Flexible resource allocation


Flexible Resource Allocation

13



Cheaper Fail-Over

14

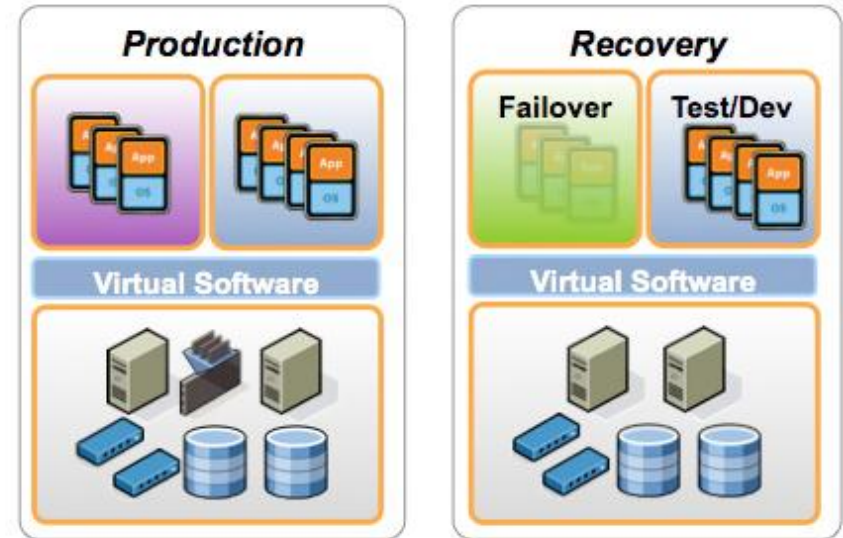


- Reduce the cost of dedicated fail-over servers
- Smooth transition when bringing fail-over servers back

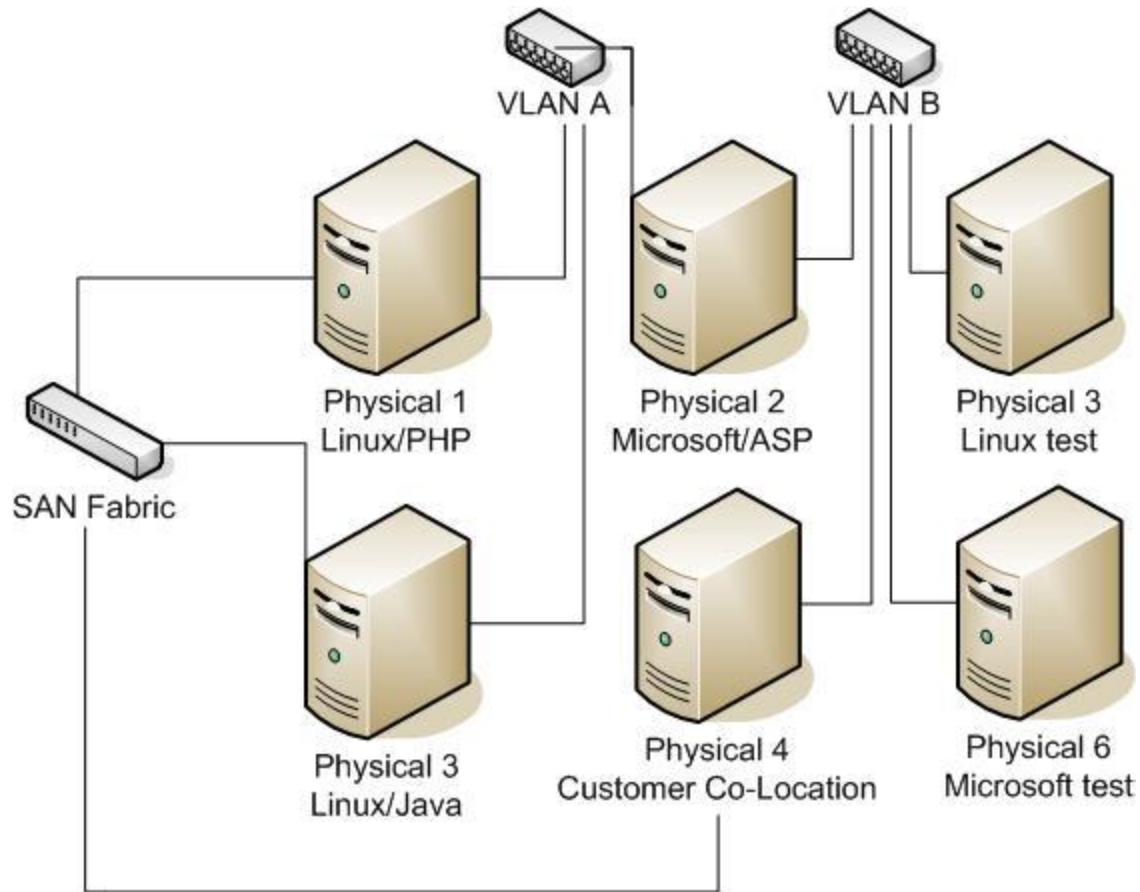
Efficient Recovery

15

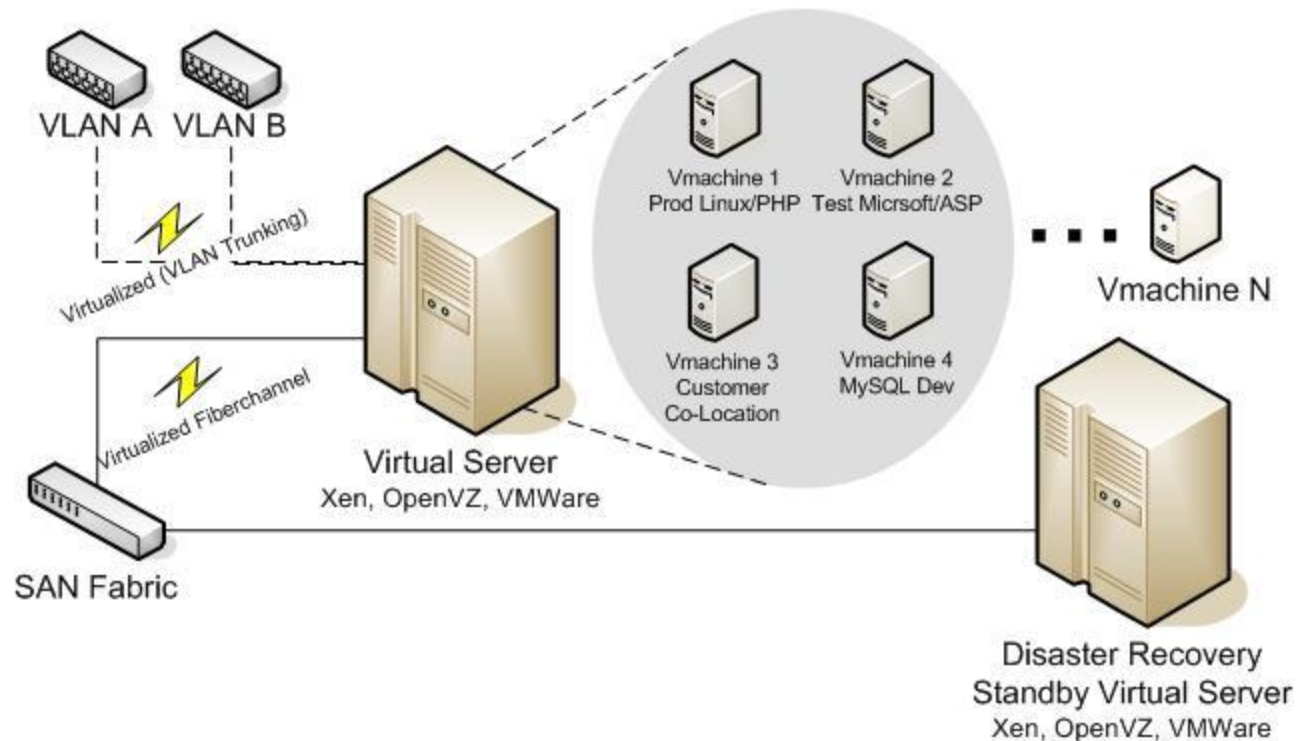
- Recovery site concept
 - For disaster recovery
 - Cold site, warm site, hot site
 - Require duplicating infrastructure (e.g. servers, networks, etc.)
- Virtualization allows a new recovery model
 - Simplify and lower cost of recovery site
 - Leverage for other workloads e.g. for testing, for other apps, etc.



Current Architecture



Virtualized Architecture




18

Storage Technology

Storage (Hard Disk Drive)

- Capacity
- Interface Technology
 - ▣ IDE (P-ATA) / SATA
 - ▣ SCSI / SAS
 - ▣ FC
- Performance
 - ▣ Seek Time
 - ▣ Data Transfer Rate (e.g. 15K RPM)



Network Attached Storage (NAS)



- File-level data storage
- Connecting directly to standard network
- Standard file-based protocols
 - ▣ NFS, CIFS, FTP, HTTP
 - ▣ UPnP, Rsync, ...
- Pros: Simple to operate and maintain, Cheap
- Cons: Performance limitation


Storage Area Network (SAN)

- Block-level data storage
- Connect to “SAN” network
 - ▣ Ethernet
 - ▣ Fiber Channel
- Storage protocols
 - ▣ Fiber Channel Protocol (FCP)
 - ▣ iSCSI
 - ▣ Fiber Channel over Ethernet (FCoE)
- Pros: High performance
- Cons: Expensive, complex

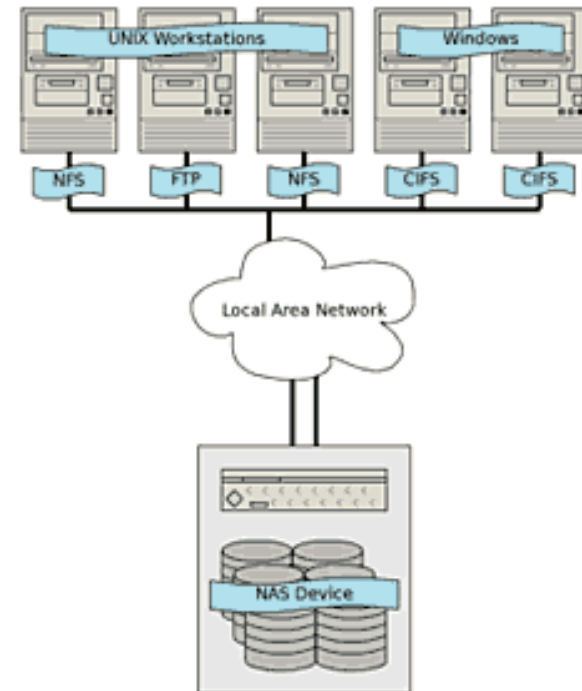


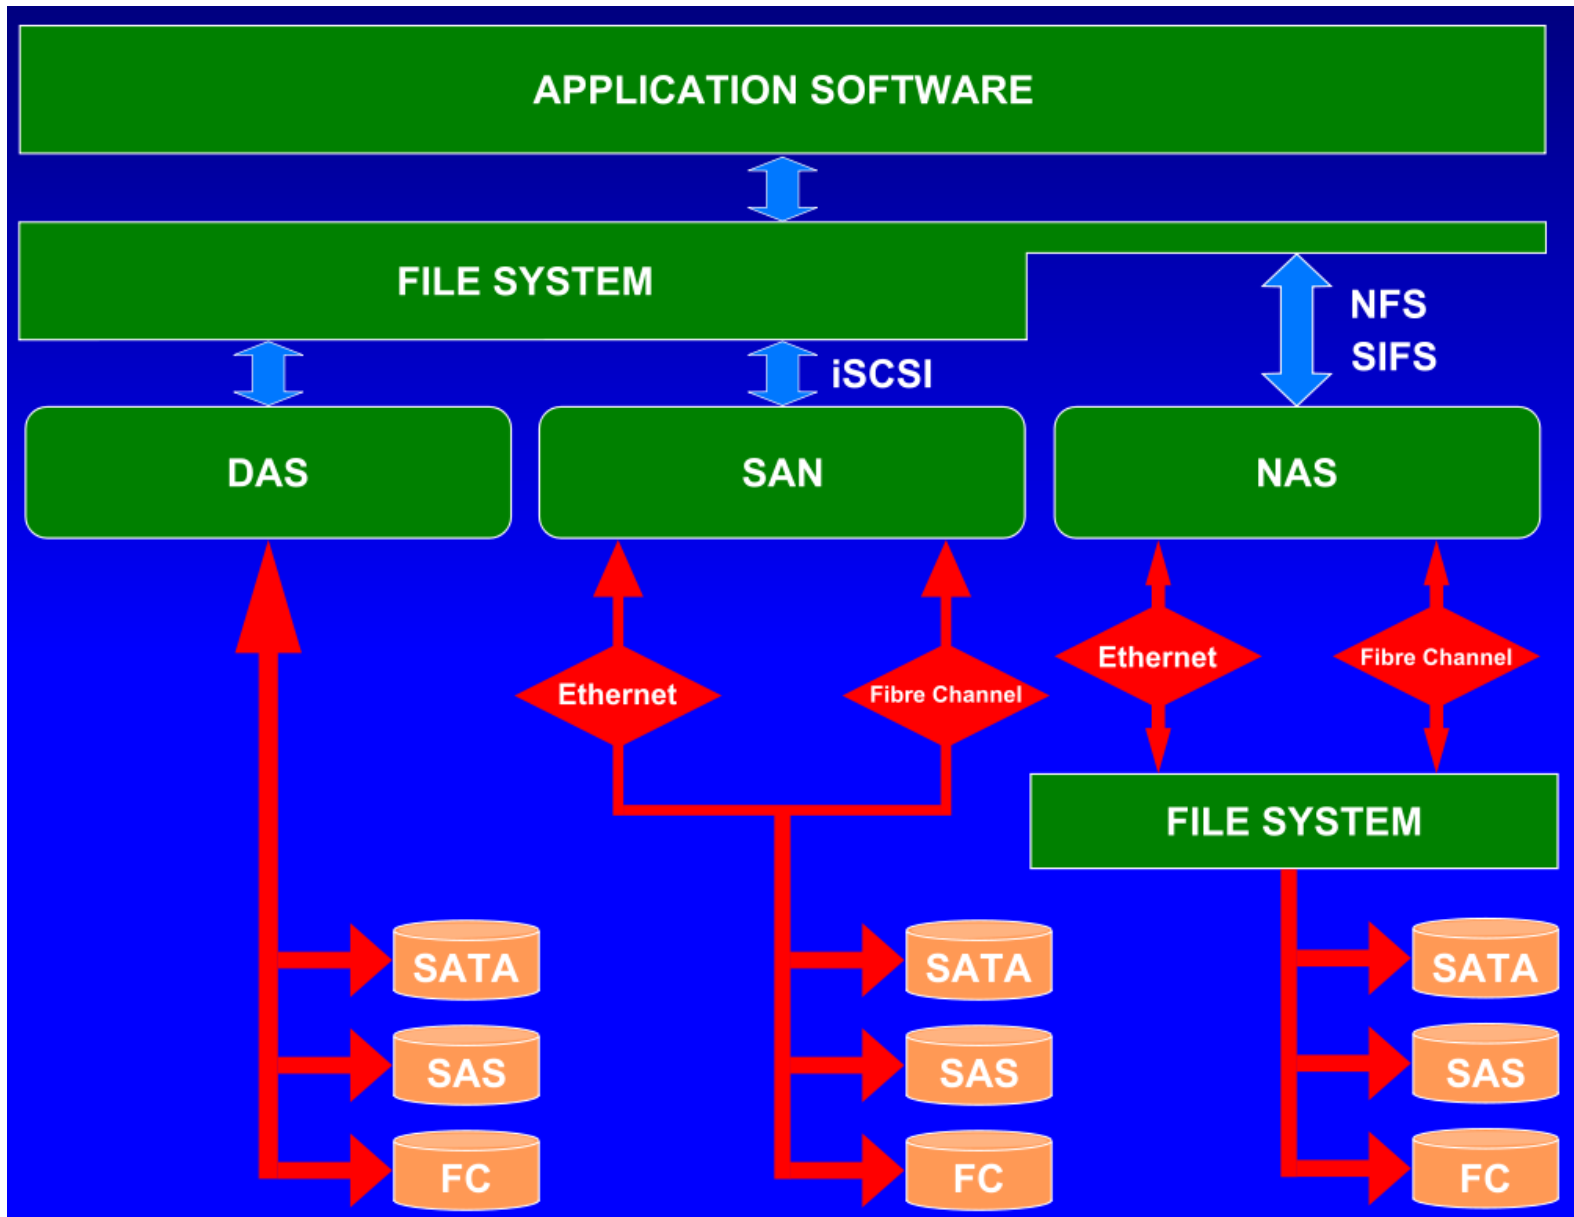
Storage Architecture: SAN vs. NAS

SAN: Storage Area Network

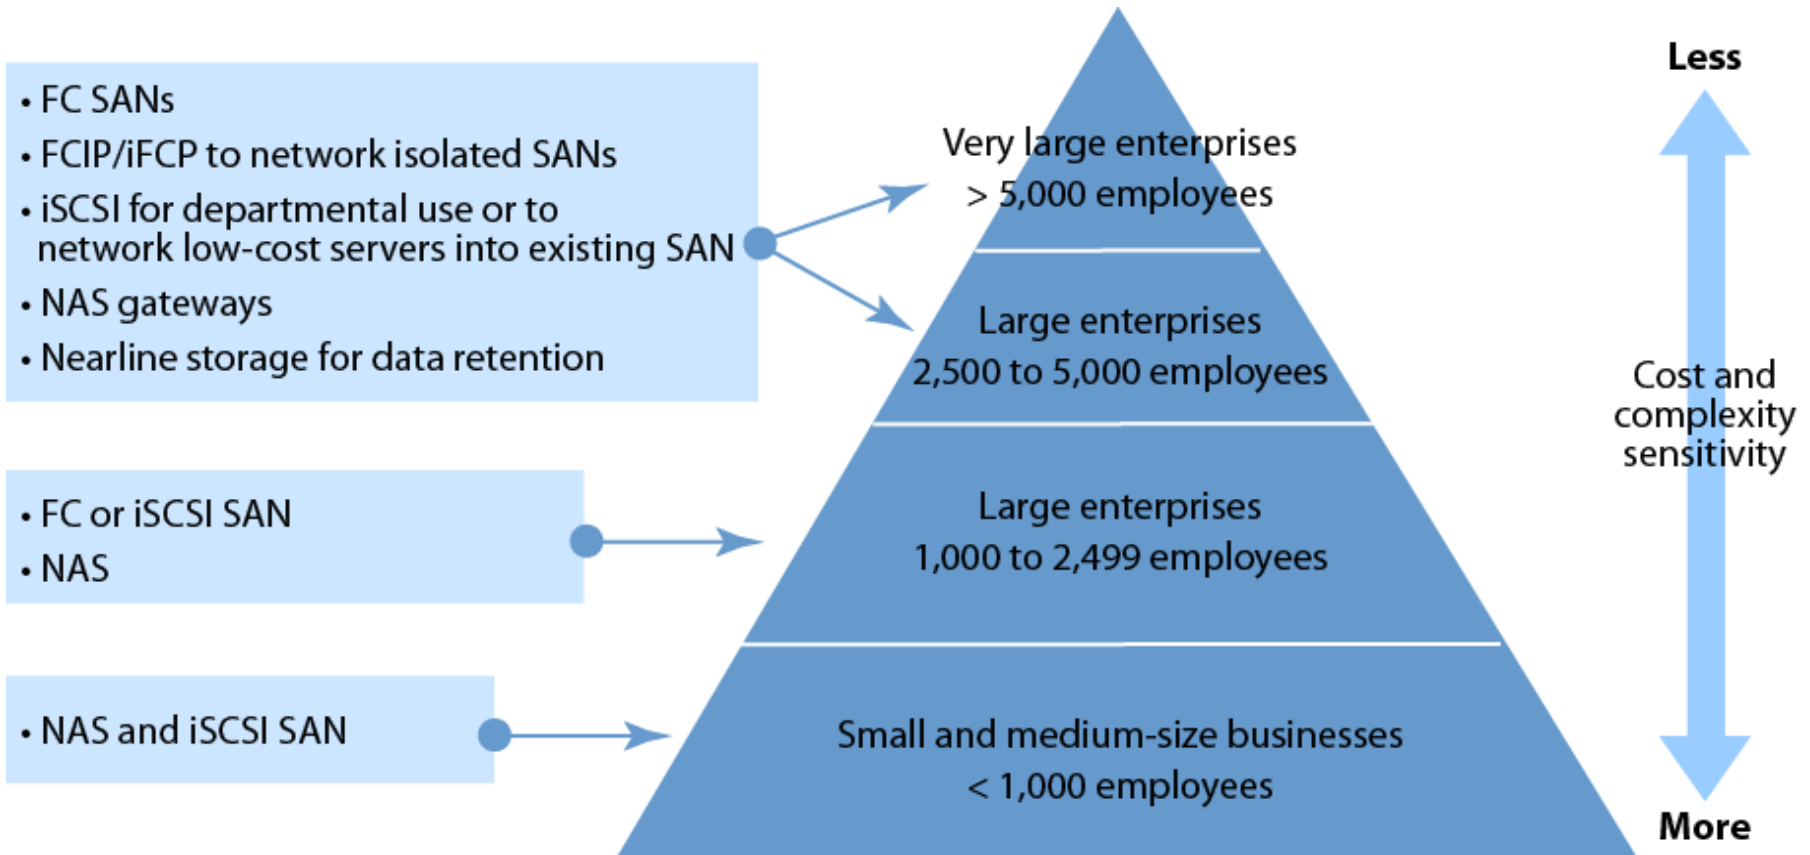


NAS: Network Attached Storage





IP-based storage adoption trends




RAID

- Redundant Array of Inexpensive Disks
 - ▣ Achieve high levels of storage reliability
 - ▣ Using low-cost and less reliable PC-class disk-drive components
- Hot-Spare
 - ▣ A drive physically installed in the array which is inactive until an active drive fails
- Hot-Swapped
 - ▣ Ability to add/remove disks without shutting down the system


Popular RAID Level

- RAID 0 – Striping
 - ▣ improved performance
 - ▣ additional storage
 - ▣ no redundancy or fault tolerance
 - ▣ N storage capacity




Popular RAID Level

- RAID 1 – Mirroring
 - ▣ Provides fault tolerance from disk errors
 - ▣ Up to one-disk failure
 - ▣ Increased read performance
 - ▣ Very small performance reduction when writing
 - ▣ 1 storage capacity



Popular RAID Level

- RAID 5 – Interleave Parity
 - ▣ Distributed data to all disks with one disk as a parity container
 - ▣ Good disk performance
 - ▣ Up to one-disk failure
 - ▣ N-1 storage capacity



References

29

- P. Strassmann, “ Introduction to Virtualization”,
<http://www.strassmann.com/pubs/gmu/2008-10.pdf>, George Mason University, 2008