.NET and Hardware: Deployment considerations

Steve Tramack
Sr. Solutions Engineering Manager
HP

HP WORLD 2003
Solutions and Technology Conference & Expo

Agenda

- .NET overview
  - Windows Server 2003 overview
- Server technology
  - Processor
  - Memory
  - I/O
  - Packaging (Blade technology)
- Storage technology
  - Performance
  - Availability

Microsoft .NET

HP WORLD 2003
Solutions and Technology Conference & Expo

Software foundation of XML Web services

Built to connect people, systems and devices

XML Web Services

XML Web Services

Connected Business

Agility to connect more deeply with customers, integrate with partners and empower employees

Connected Development

Quickly build solutions that interoperate across languages, platforms and devices

Microsoft .net connected

Connected Experiences

Information anytime, anywhere, any device

Connected Systems

Integrate new and existing systems across the organization, and externally with partners.

Microsoft .NET

HP WORLD 2003
Solutions and Technology Conference & Expo

Web services support across the Microsoft platform

.NET Web Services Today

Emirates Airline Group

Mobile Device

Microsoft® .net™

Dynamic Internet Solutions Architecture

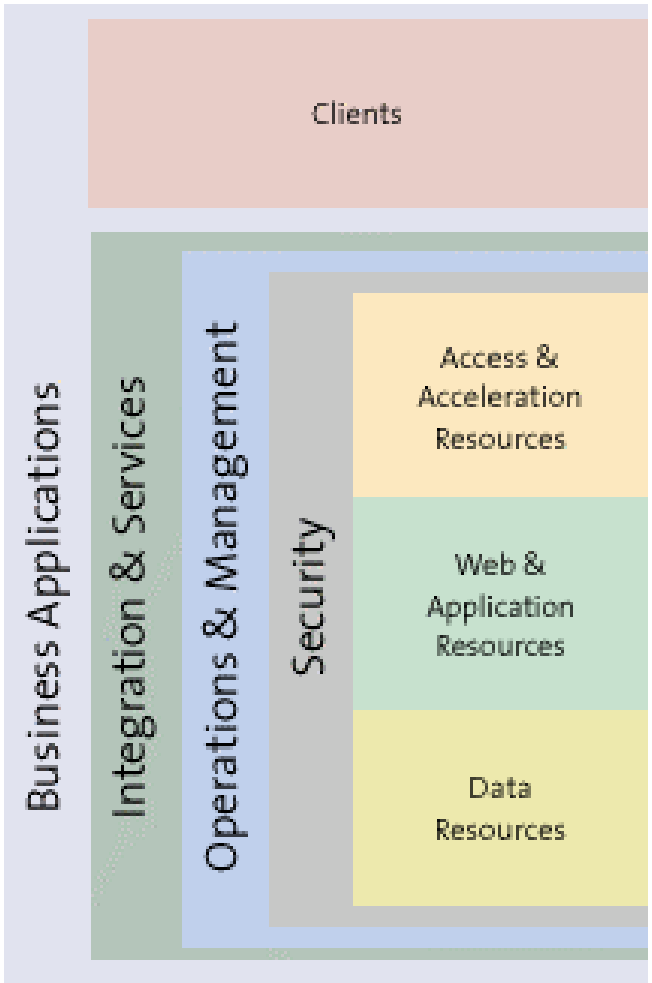

Windows Server 2003
…at the heart of a .NET infrastructure

- New versions
  - Web: 2P, 2GB
  - Standard: 4P, 4GB
  - Enterprise (32 and 64 bit): 8P, 32GB, 8-node clusters
  - Datacenter (32 and 64 bit): 64P, 64GB, 8-node clusters

  Hyper-Threading aware

  New Active Directory

  New management features, including headless deployment

  VSS API (online snap / clone)

  IIS 6.0

  Better reliability, availability, performance
- .NET Framework integration (XML, SOAP, UDDI, WDSL)
- Real-Time Collaboration features

Windows Server 2003

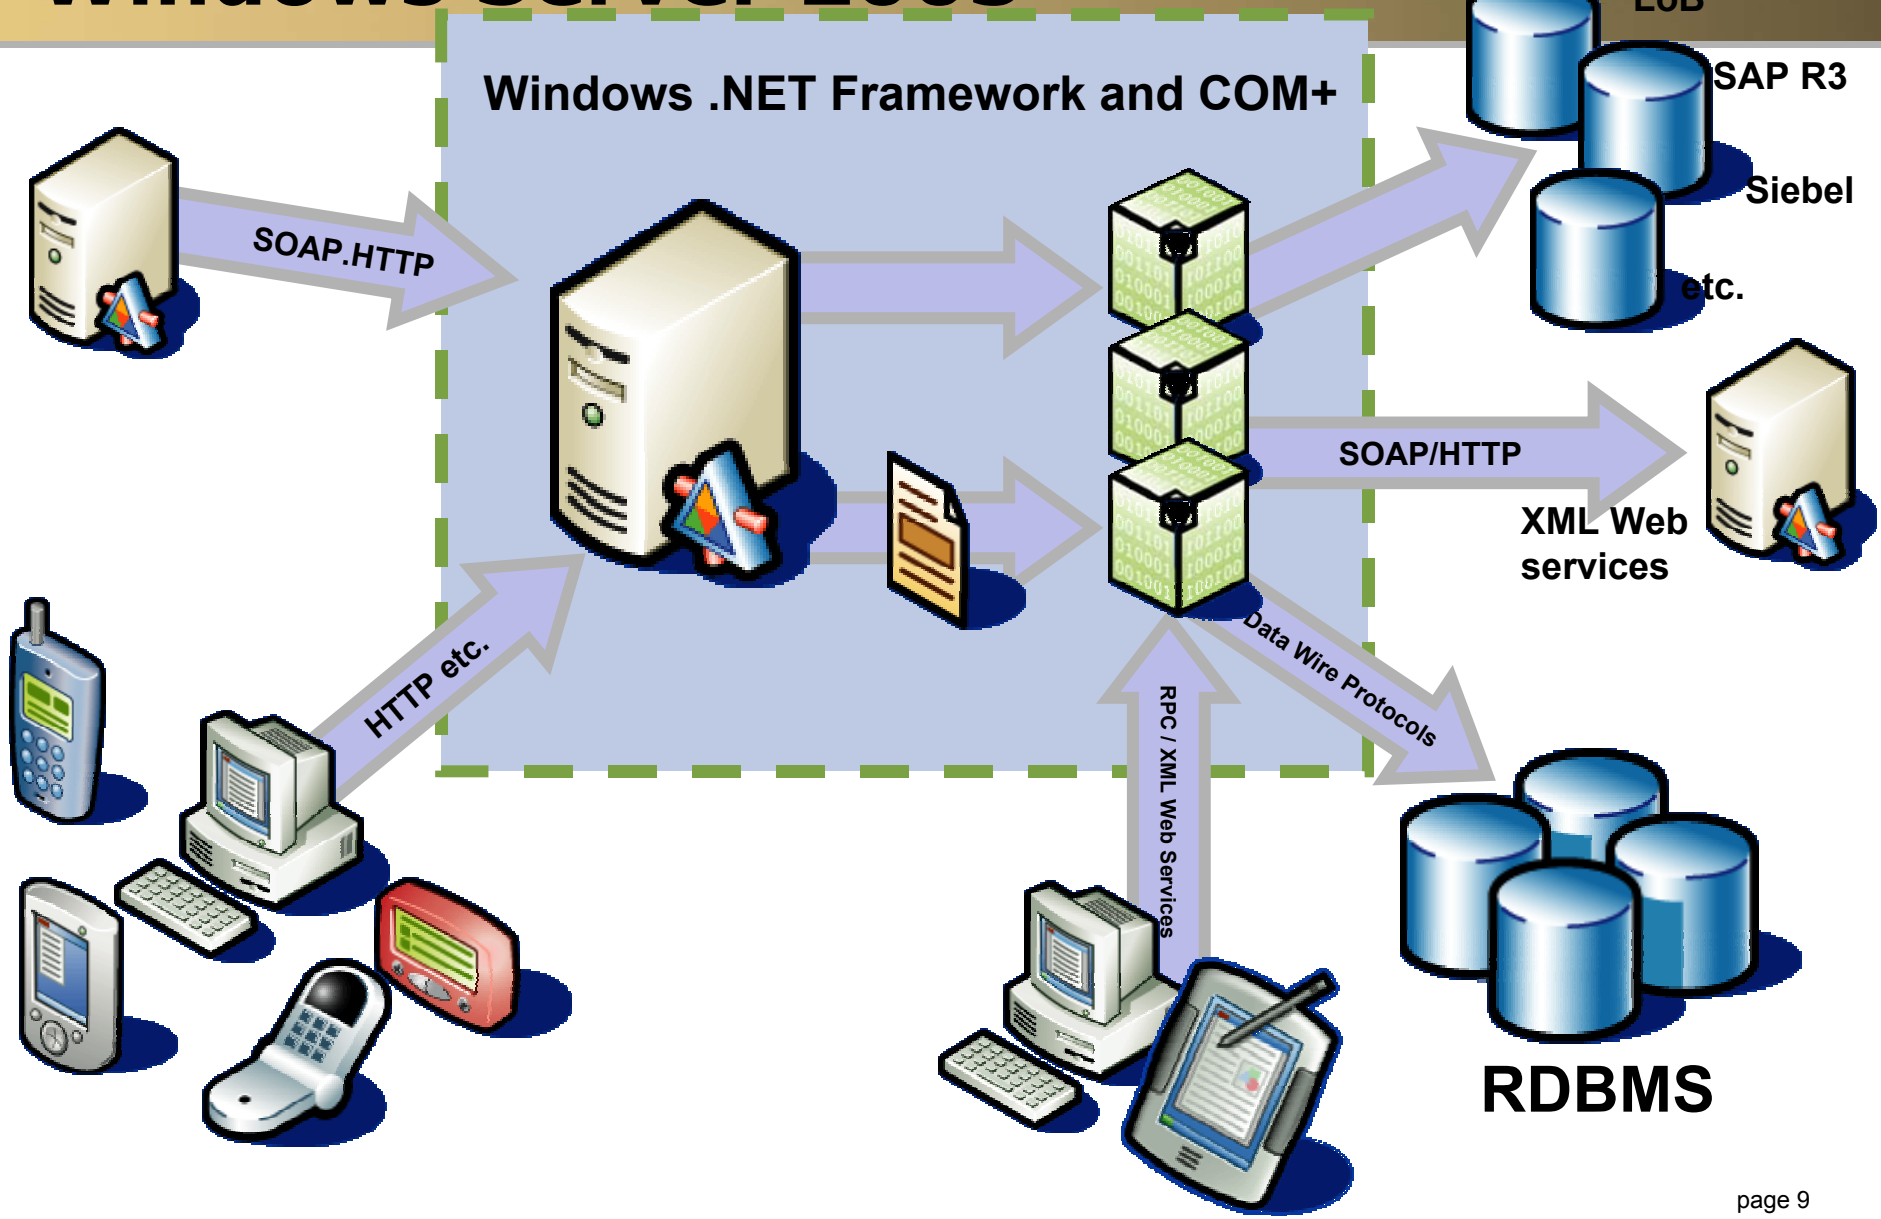

Server hardware technology

- Performance and Availability
- Elements
  - Processors
  - Memory
  - PCI-X
  - Server packaging technology (blades)

"Tualatin"

Intel Pentium III Processor

- .13 micron
  - 1.4GHz
  - 800MHz ultra low voltage
    - 800 MHz version consumes 11.2 Watts of power
    - Replaced by Pentium M processor up to 1GHz, 1MB L2 cache

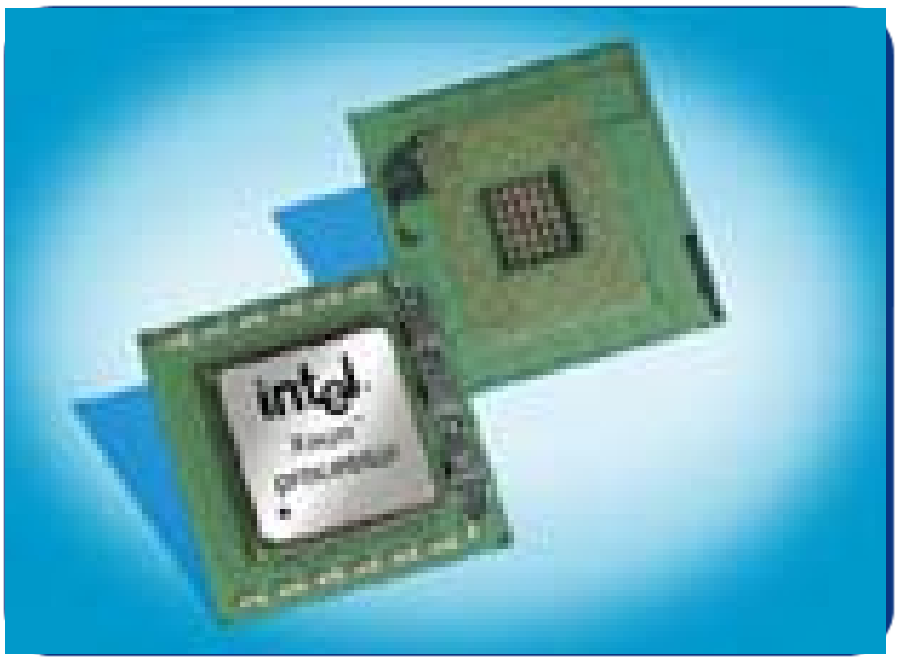

- 133 MHz System Bus
- 512K L2 ATC
  - Advanced Transfer cache transfers at core clock speed

"Cascade"

Intel Pentium III Xeon

- .18 micron technology
  - 700 – 900MHz (1.0GHz DP version)
- 133MHz FSB
  - 1.06GB/s to and from processor

- 1 – 2MB L2 ATC
  - Advanced Transfer cache transfers at core clock speed
- Greater than 2P configurations

"Prestonia"

Intel Xeon Processor

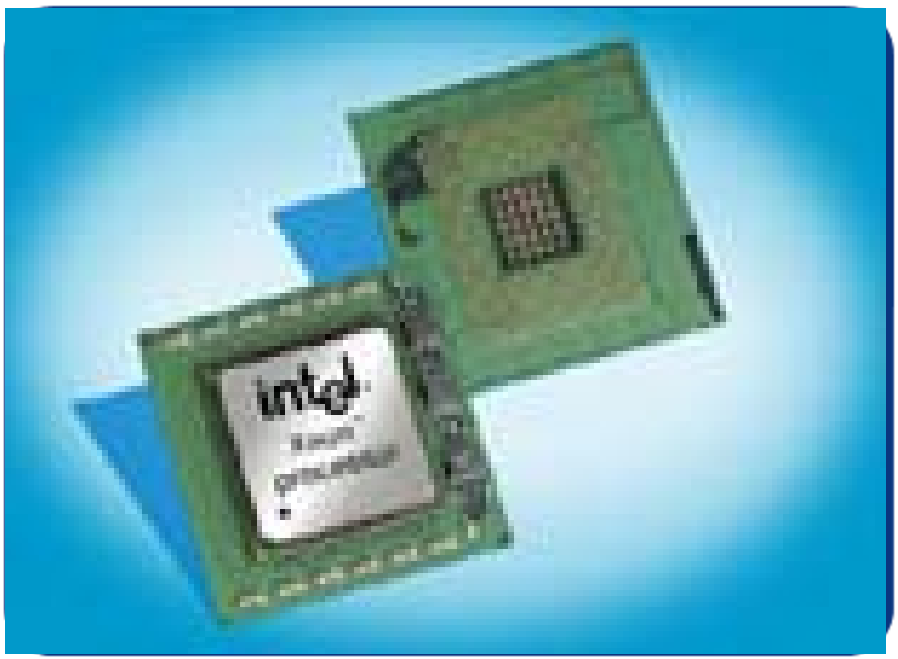

- .13 micron
  - Faster speeds and performance
  - 2.4 – 3.06GHz
- 400/533 MHz FSB
  - Up to 4.3GB/sec to and from processor
- 512K/1M L2 ATC
  - Advanced Transfer cache transfers at core clock speed
- Hyperthreading (Jackson)
  - Parallel processing
  - Each processor processes 1.5 procs-worth of instructions

"Gallatin"

Intel Xeon Processor MP

- .18 micron
  - 1.4 – 2.8 GHz
- 400 MHz FSB
  - 3.2GB/sec to and from processor
- NetBurst micro architecture

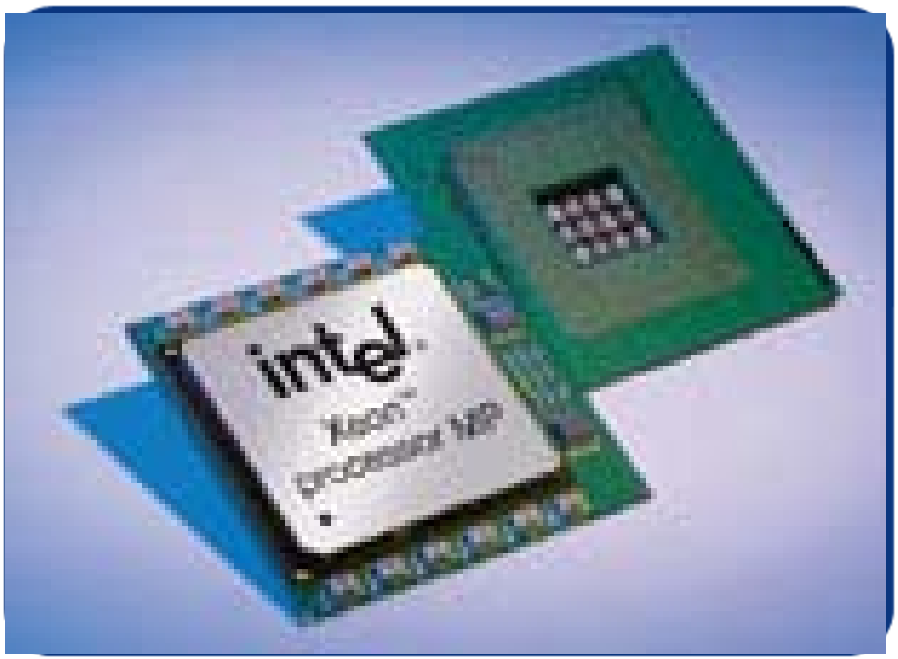

- 256K L2 ATC
- 512K / 1MB / 2MB integrated L3 cache
- Hyperthreading (Jackson)

Processors: hyper-threading

- Introduction of hyper-threading
  - Two architectural states per processor sharing a single execution unit
    - Operating system sees two logical processors, enabling execution of two threads simultaneously on each physical processor
  - Helpful in intensive threaded environments, such as Microsoft Exchange
  - Does not help in other products/environments (example: SQL TPC-C)
- Careful with licensing issues
  - Windows 2000 considers virtual processors to be real
    - Operating system uses first <n> processors (4 with Server, 8 with Advanced Server)
  - Windows 2003 differentiates between logical and physical processors

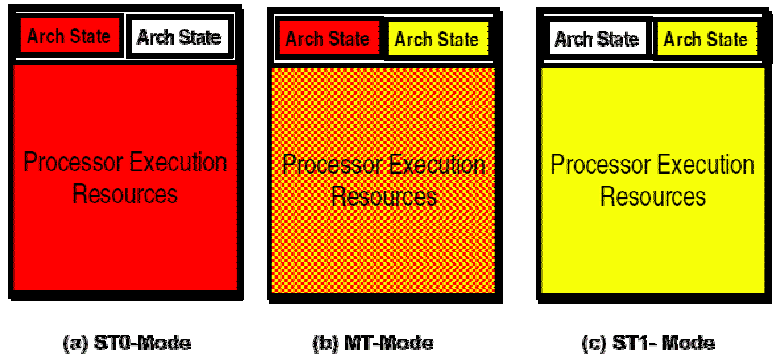

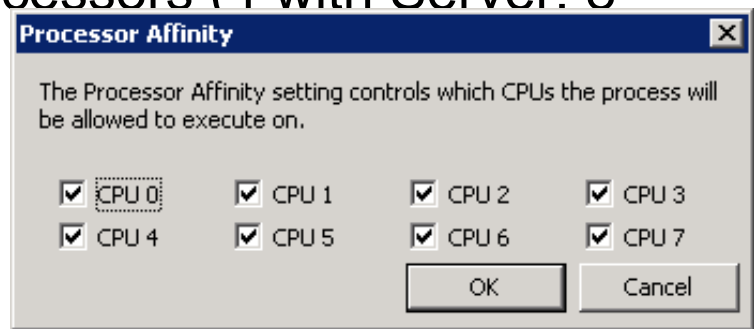

"McKinley"

Intel Itanium 2 Processor

- 1.3 – 1.5 GHz
- Based on EPIC architecture
- 400 MHz, 128 bit system bus
  - 6.4GB/sec to and from processor

- 64-bit address space
- 256K L2; 3, 4 and 6MB iL3 cache
- Enhanced Machine Check Architecture (MCA) with error correction

Microsoft Exchange benchmark results

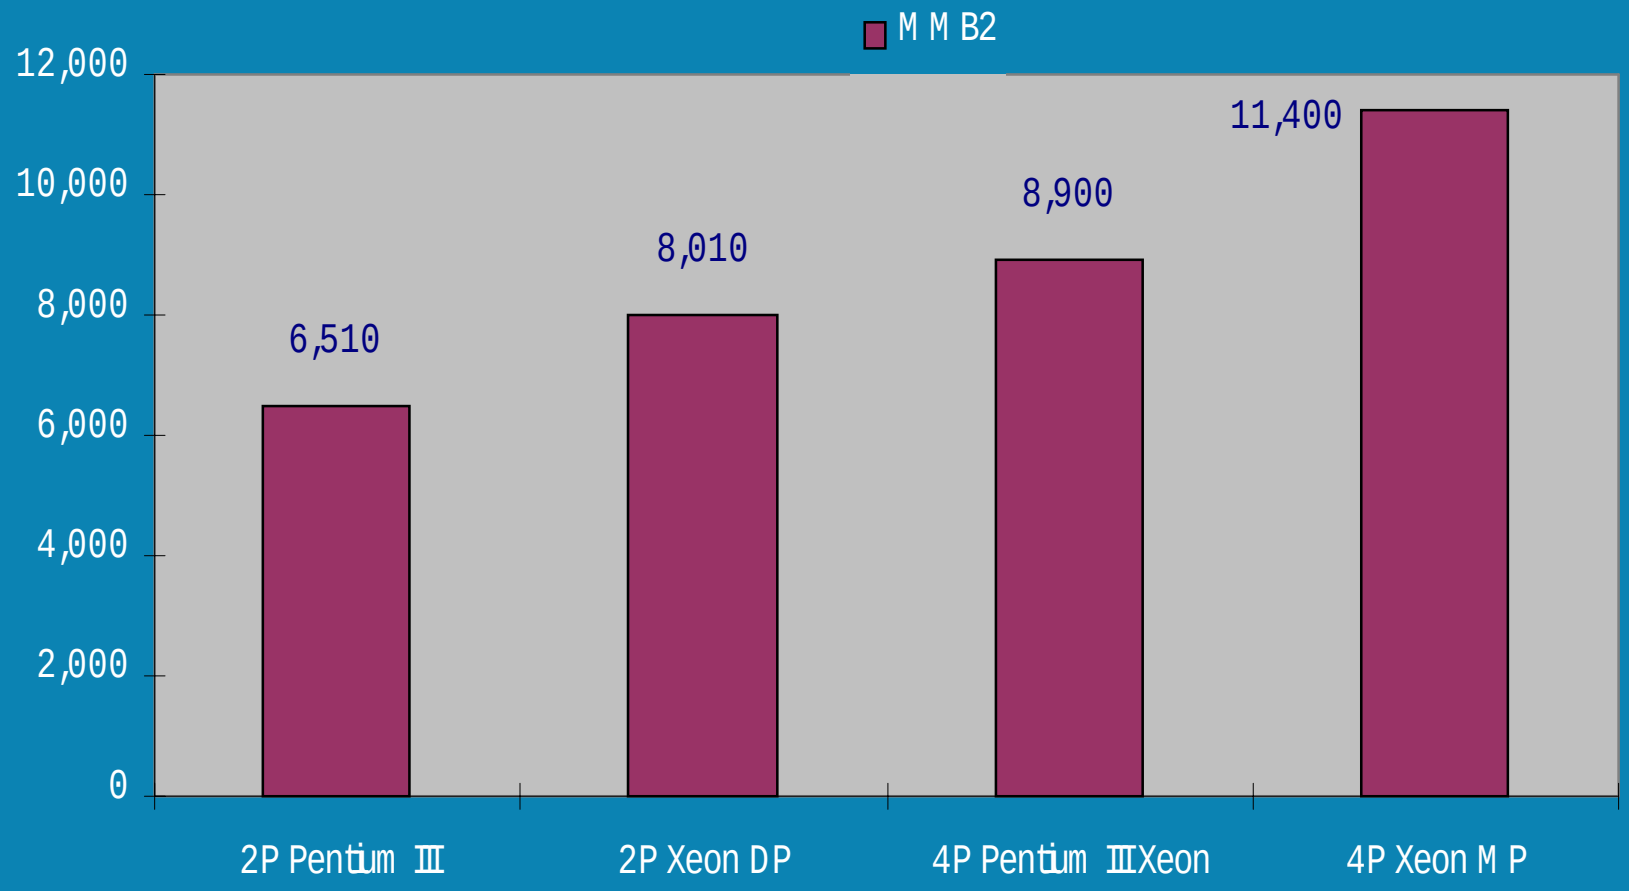

Microsoft Exchange benchmark results

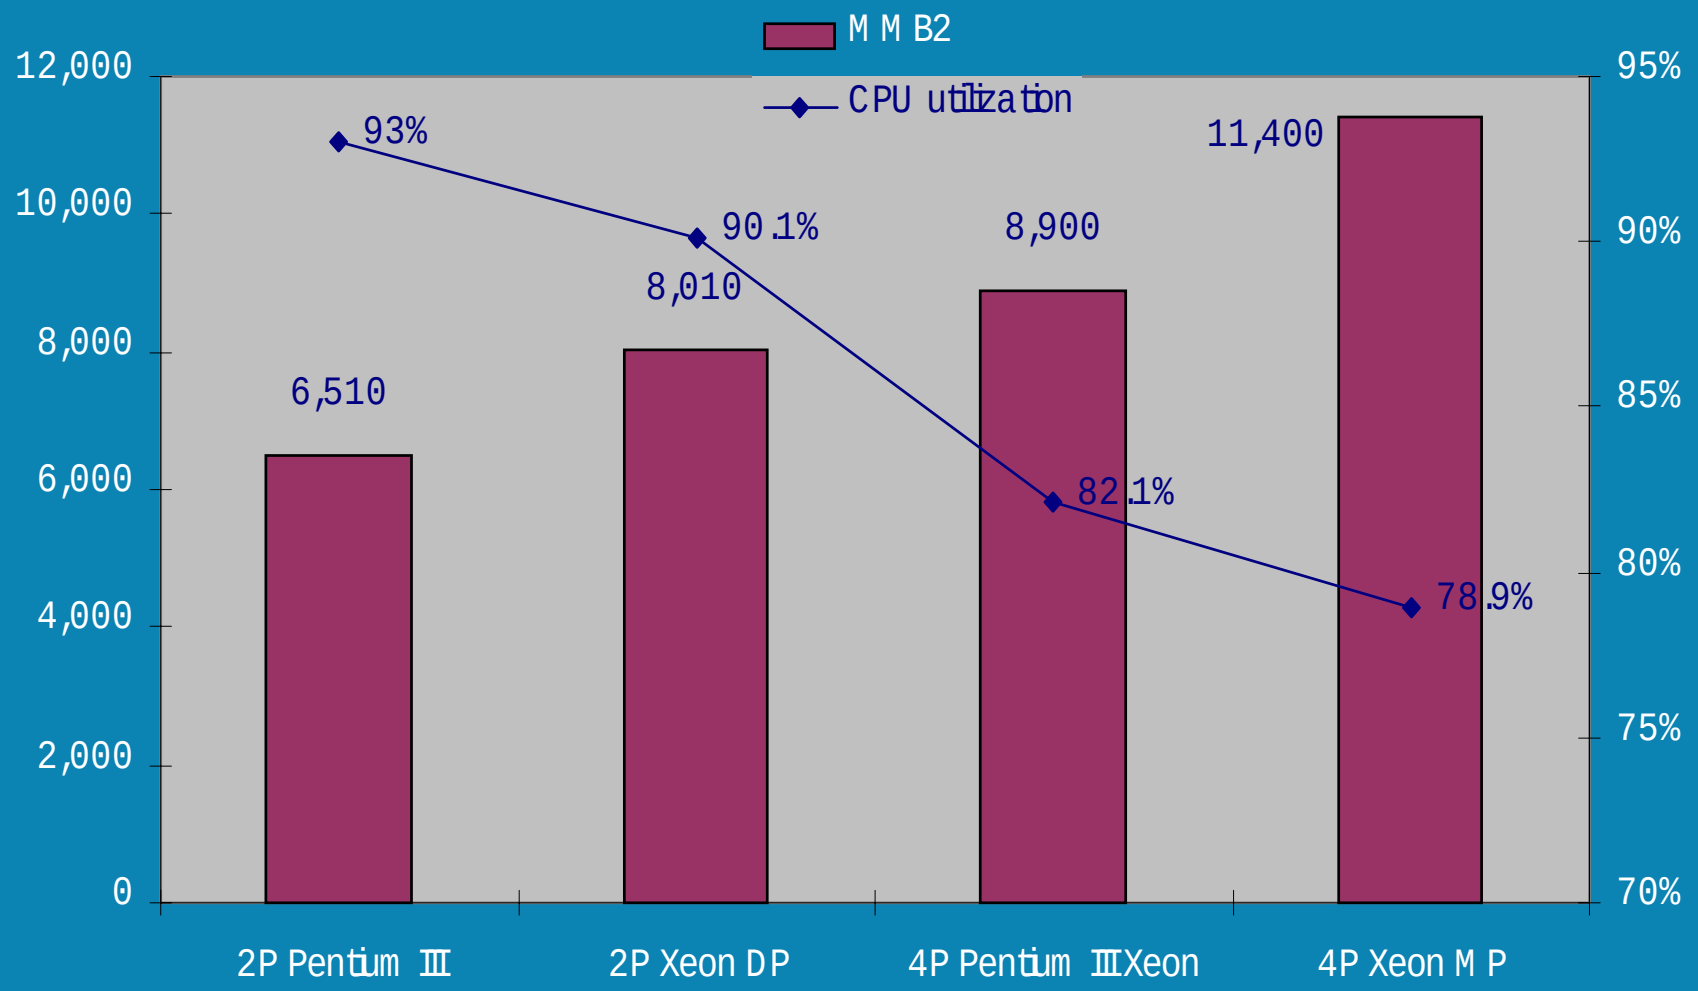

DDR Memory Technology

- 266MHz DDR SDRAM (PC2100 registered ECC)
  - transfer two data words per clock cycle
  - Memory bus runs at 533MHz to match the front side bus frequency of the processor

- 2:1 and 4:1 Interleaved Memory
  - Increases the amount of data accessed during a given memory read
  - Number of wait states is decreased, further enhancing performance

- 32GB Memory Configurations in Foster-based systems

Hot Plug Mirrored Memory

Now

Hot Replace

- Replace a failed DIMM while the server is running
- OS / Application transparent

New

Hot Add - Upgrade

- Add a new bank of DIMMs while the server is running
- OS / Application enabled

Microsoft .net Enterprise Servers

Hot Plug RAID Memory - Intelligent Fault Resilience

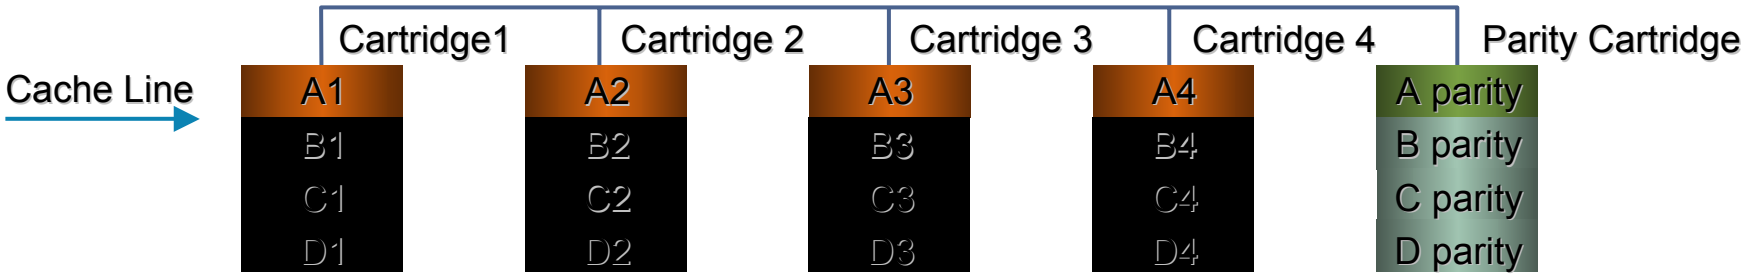

Single-bit Memory Errors
Detect & Correct

Multi-bit Memory Errors
Detect & Correct

Greater than DRAM Errors
Detect

Error Condition	Standard ECC	RAID Memory
Single-bit	Correct	Correct
Double-bit	Detect	Correct
4-bit DRAM	Detect	Correct
8-bit DRAM	May detect	Correct
Greater than DRAM	May detect	Detect

Memory

- 4GB is still the maximum memory you can use with 32-bit apps (e.g. Microsoft Exchange)
  - Introduction of /USERVA=xxx to fine-tune the kernel mode space size (recommended: xxx=3030)
- Improvements regarding VM fragmentation issues
  - Top-down approach to thread memory allocation
  - Bottom-up approach for database cache and instance

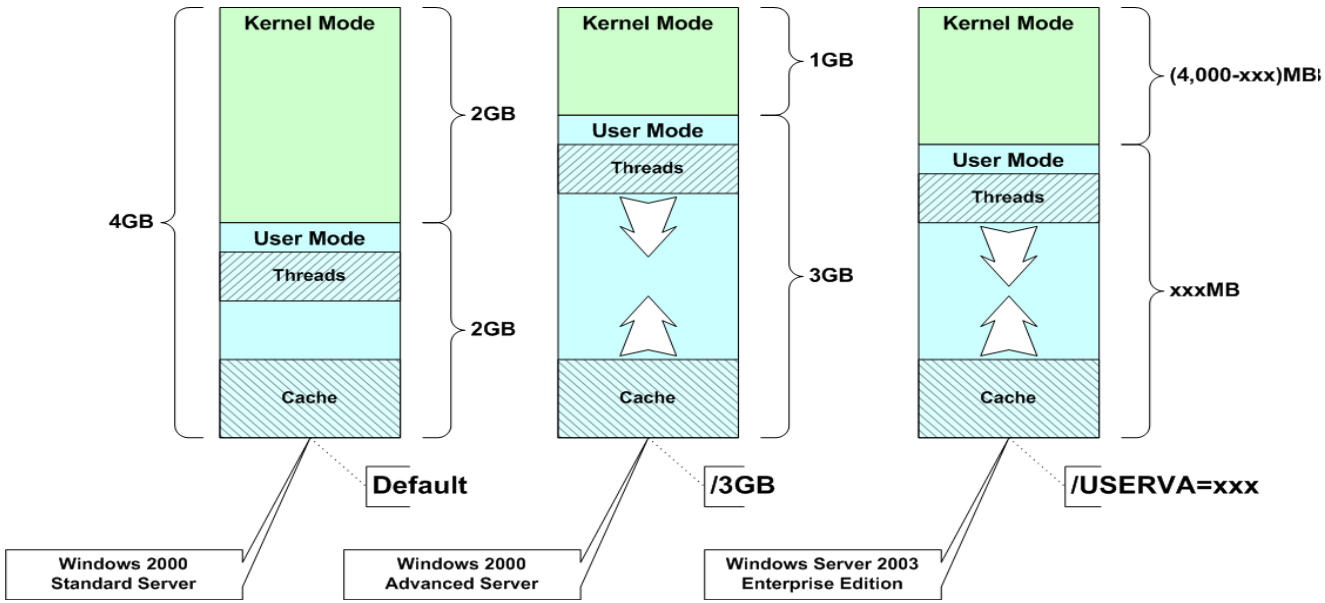

I/O Bus Evolution

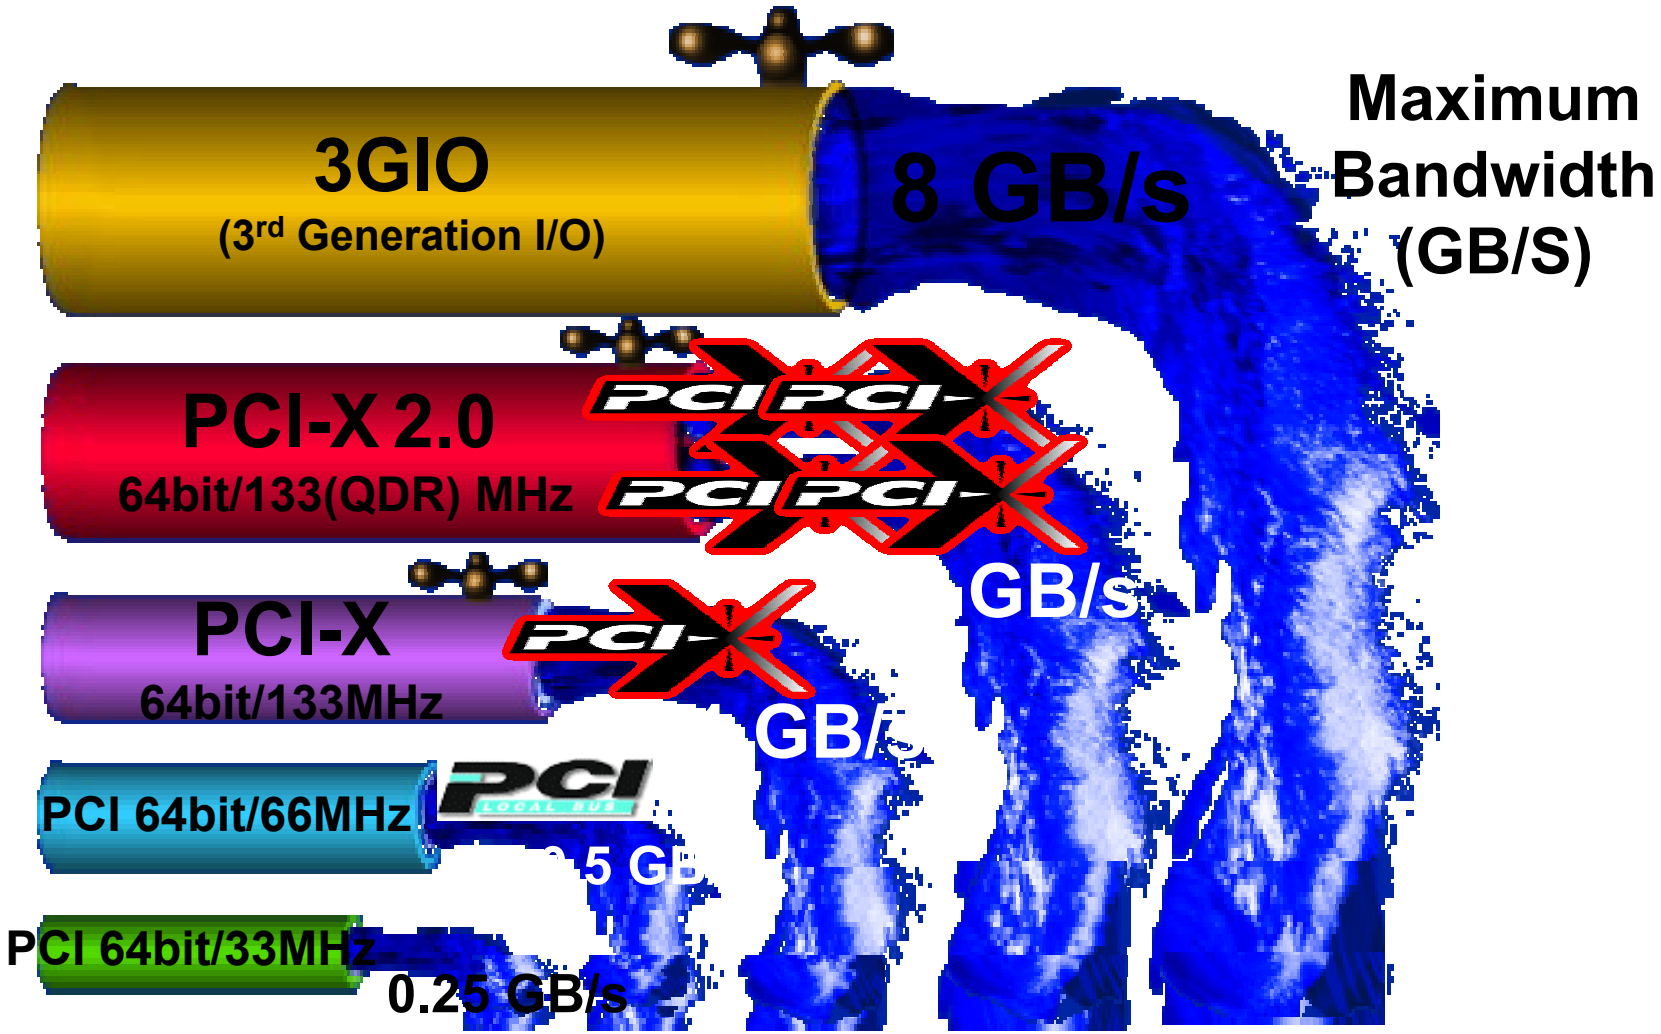

I/O: Disks

- Disks have improved
  - $/GB
  - GB/ft
  - MB/s
  - NOT I/O/s
- Still considered slow devices
  - 50-70ns to memory
  - 5-15ms to disks
  - Additional CPU overhead

I/O: Network

- 10-100Mb/s prevalent
- 1Gb/s
  - E2K backup over network
  - iSCSI
- Requires processing of the OSI stack
  - Some NICs will offload
- NIC teaming

Server packaging – blades

ProLiant BL e-Class
power-efficient edge server blades

- static web server
- utility apps *(firewall, gateway)*
- computational cluster node

ProLiant BL p-Class product performance front-end and mid-tier server blades (ProLiant BL20p, BL20pG2)

- dynamic web services / ASP.NET
- computational cluster node
- terminal server farm
- AV, media streaming

ProLiant BL p-Class product high performance back-end server blades (ProLiant BL20pG2, BL40p)

- database server
- mail / messaging server
- fiber SAN attach
- HA failover cluster

rack-centralized power for BL p-Class products

- redundant
- hot plug

Storage technology

- Important consideration for many .NET Enterprise Servers
- Performance considerations
  - I/O planning
  - SANs vs. NAS
  - iSCSI
- Availability considerations
  - Clustering
  - RAIS
  - Virtualization
  - BCVs

Storage

- Capacity versus I/O
- Use Cache for low latency
  - Read versus Write
- Use RAID for high transaction rates
  - Impact of RAID levels
- Avoid saturation
  - Similar response between 10% and 90% of the workload

I/O Performance

What are we trying to achieve?

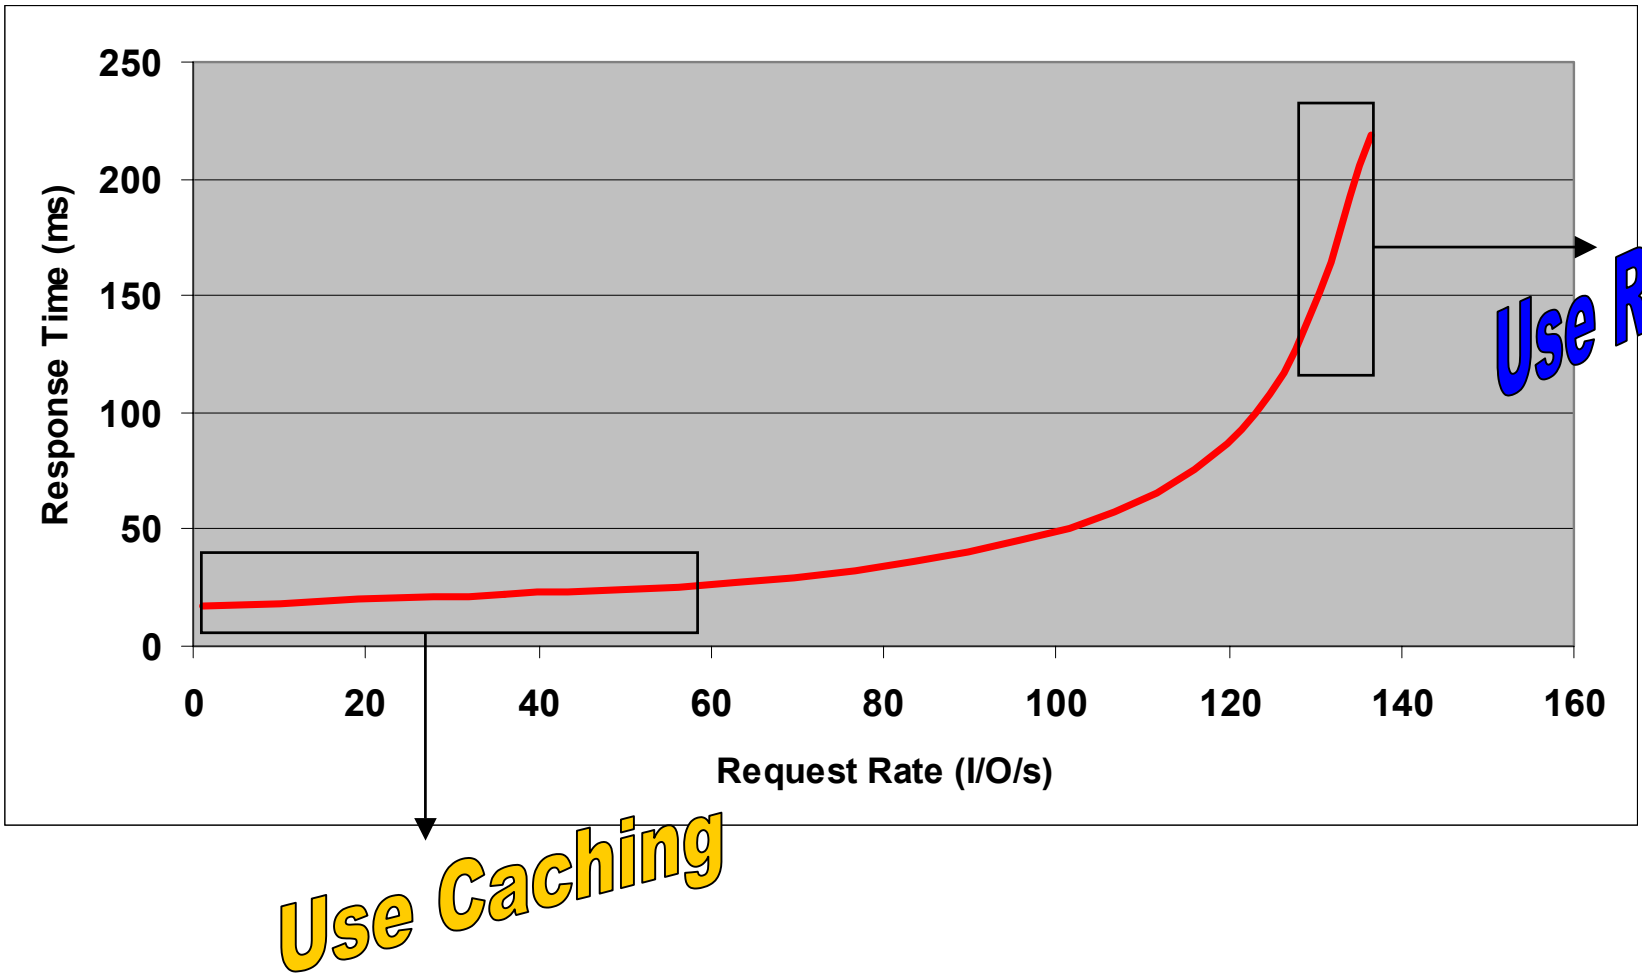

I/O Performance

Why is the I/O Profile Important?

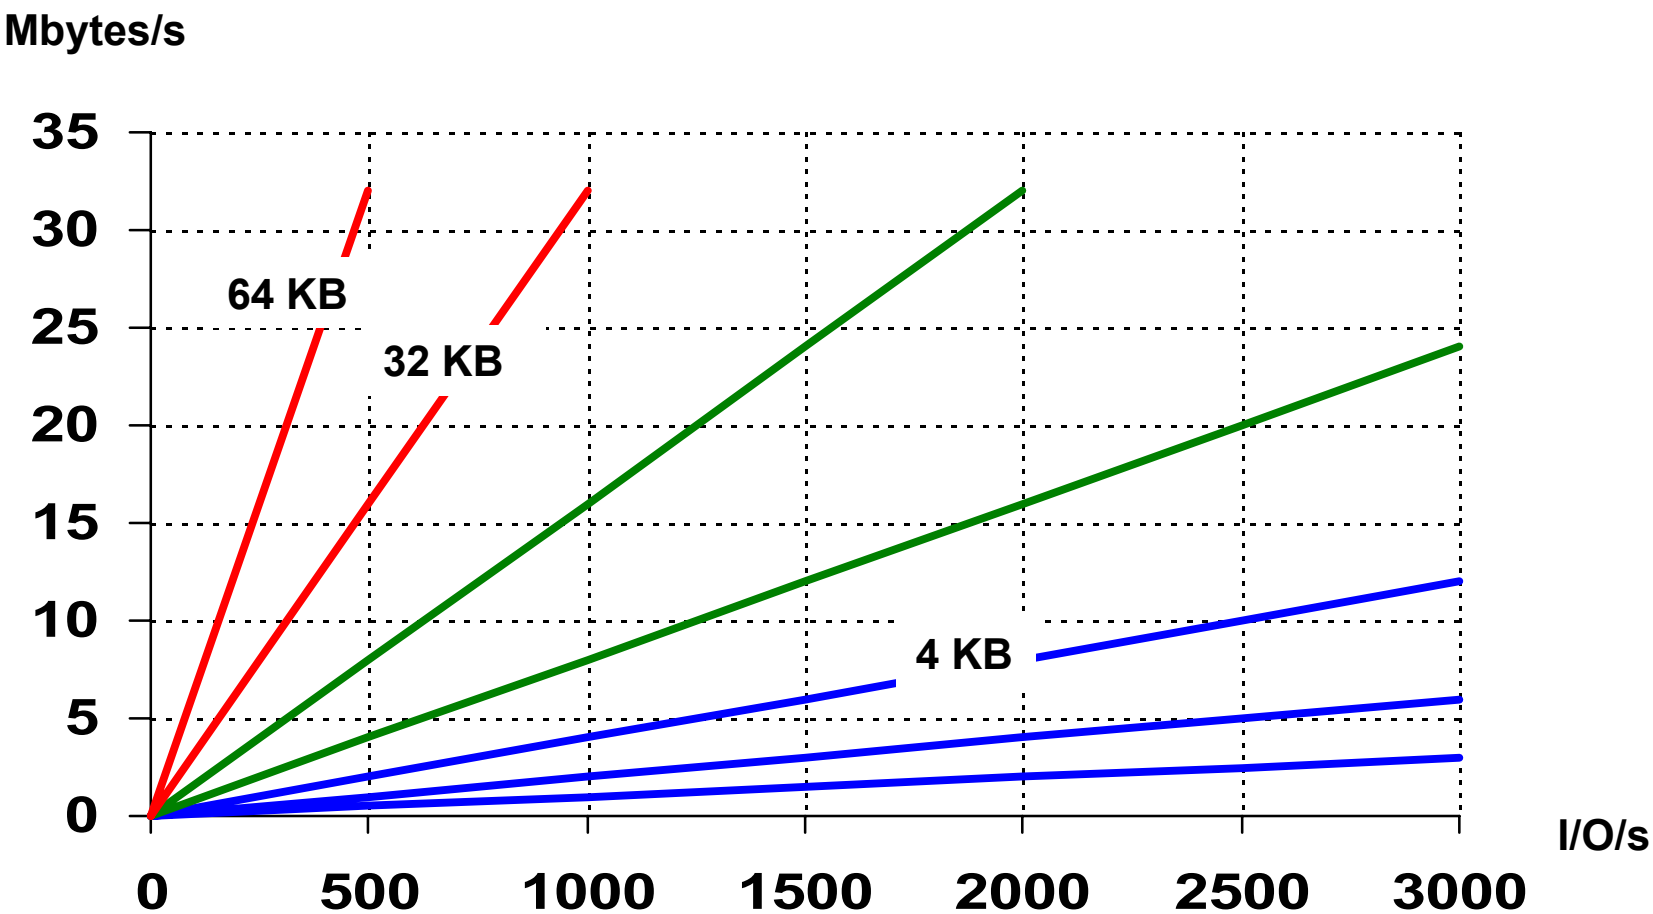

Performance Design
I/O Profiling by RAID level – Play Along at Home!

- How to Analyze Disk I/O (Formulas)
  - RAID0: READS + WRITES = I/Os per second
  - RAID1: READS + (2*WRITES) = I/Os per second
  - RAID5: READS + (4*WRITES) = I/Os per second
- BASIC Disk Subsystem Capabilities (per 10K spindle)
  - Random I/O: 80-100 I/Os per second
  - Sequential I/O: 150-200 I/Os per second
- I/O Profile Example:
  - (12) 36GB drives in a RAID5 Array for IS Store (Random I/O)
  - PerfMon: Reads/Sec: 265; Writes/Sec: 198
  - Disk Array Capability: 12*80 = 960 I/Os per second
  - Exchange Server Demand: 265+(4*198) = 1060 I/Os/sec

Reference: Compaq RAID Technology for Database Servers White Paper

SAN: Storage Area Network

- Targeted to application servers such as Exchange
- High performance, capacity
- Sharing at the cabinet/infrastructure level
  - Can have some inconvenience
- Struggle for interoperability
  - Mix various vendor equipments
  - Stretches support contracts
- Basis for today's data centers and tomorrow's .NET server deployments
- Performance and availability are not necessarily selling points
  - Flexibility, agility

NAS: Network Attached Storage

- Filers, originally
  - i.e., File Shares
- Provide storage sharing (file system level)
- Useful in cache, web environments (DISA)
  - Many front-ends share the same source of information
- Limited support with IO-intensive back-end servers
  - Most require features, APIs, access methods available in local block mode Windows volumes
  - Performance, reliability issues

Example: Storage networks with Exchange 2003

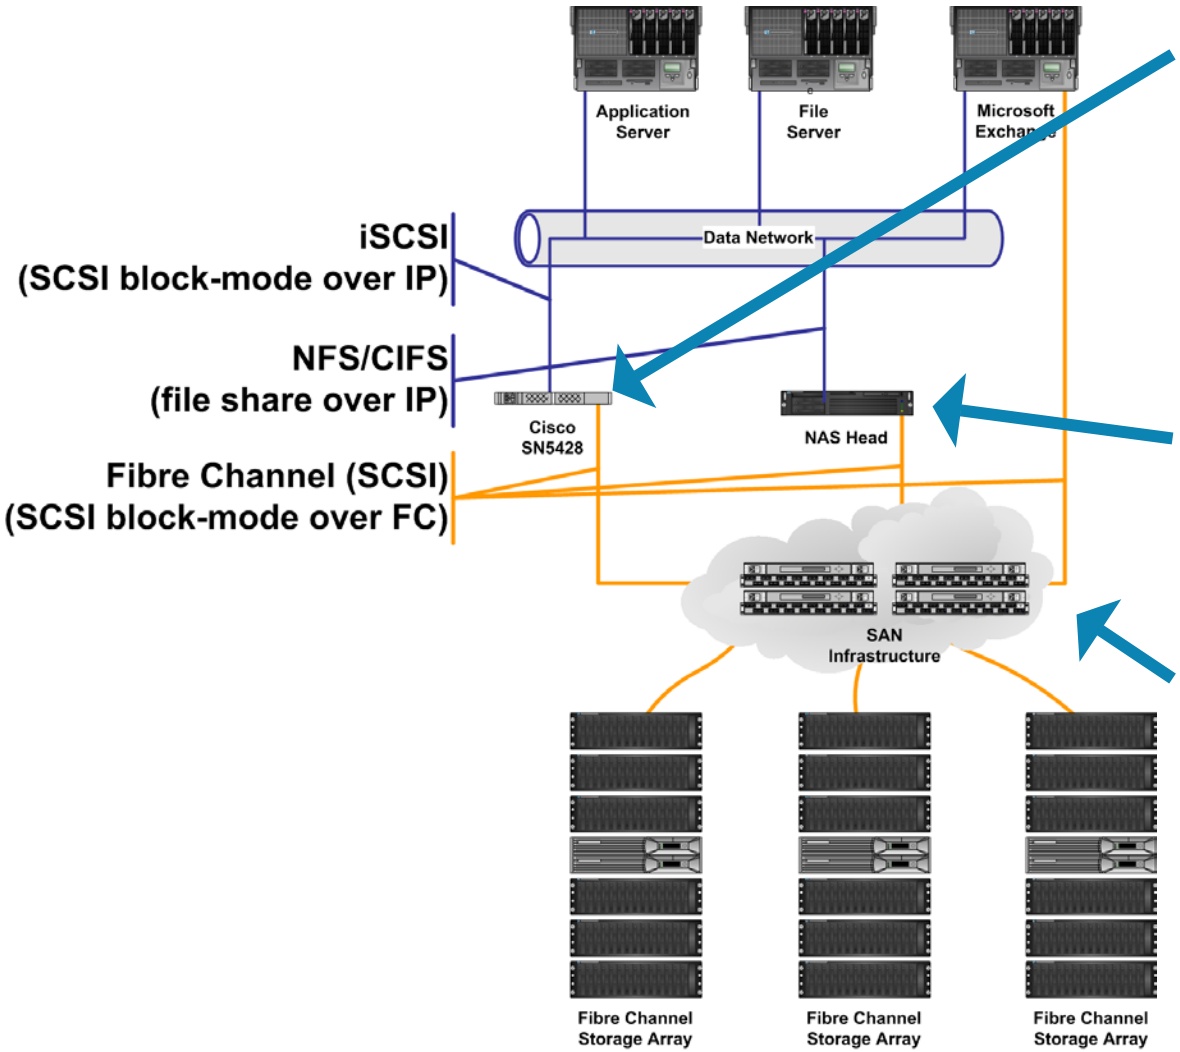

- iSCSI router: offer iSCSI support based on SAN back-end
  - Cisco is current supported HP solution
  - Drivers available from Cisco and soon from Microsoft
- NAS head: offer IP-based storage based on SAN back-end
  - Could be Windows-based using SAK
- Microsoft Exchange: use SAN (or DAS) back-end
  - Say NO to proprietary block-mode protocols
  - Microsoft Exchange 2000 and later: no support for file shares

Typical Cluster Design

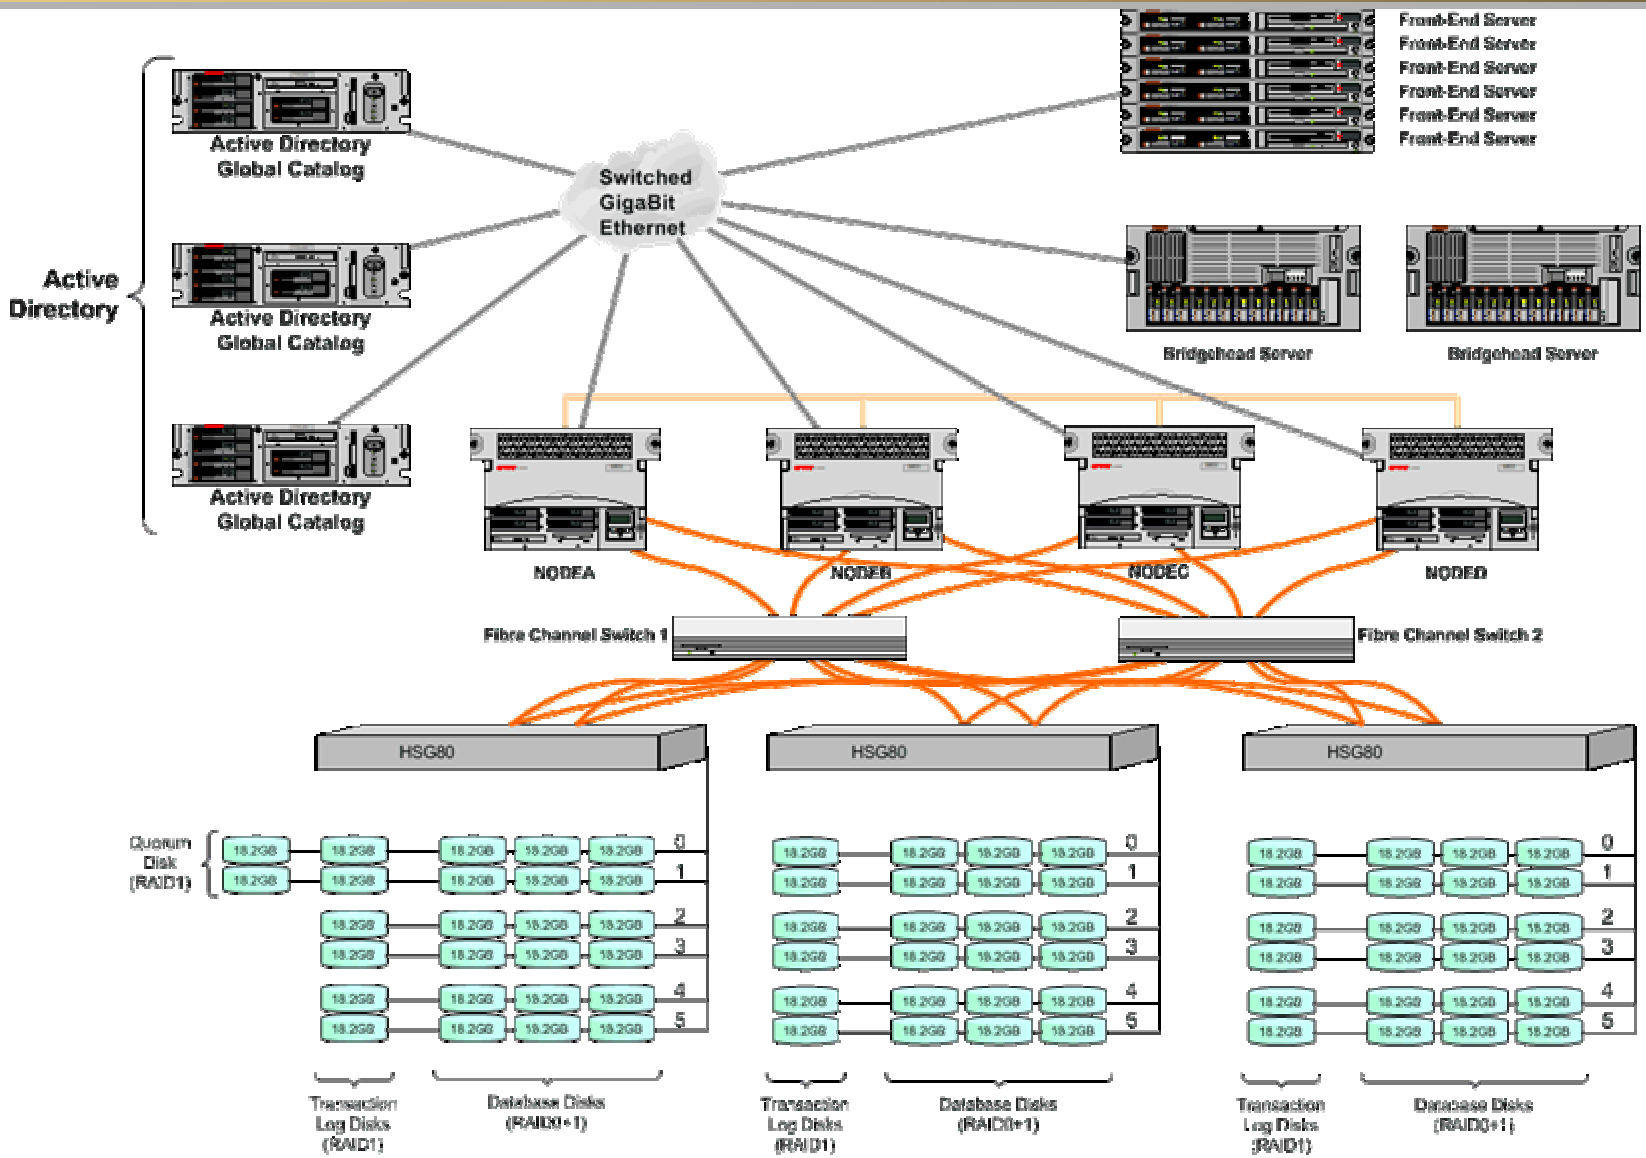

Redundant Array of Independent Servers

- Core principle
  - Boot servers in a separate storage infrastructure based on SAN technology
- If a server crashes in flames
  - Reassign SAN volumes
  - Power on a spare (cold-standby server) server
- Nothing stops you from replicating the boot and data volumes to a distant site
  - Cheap (?), yet powerful disaster tolerant solution
  - Careful about storage restrictions

RAIS: Example

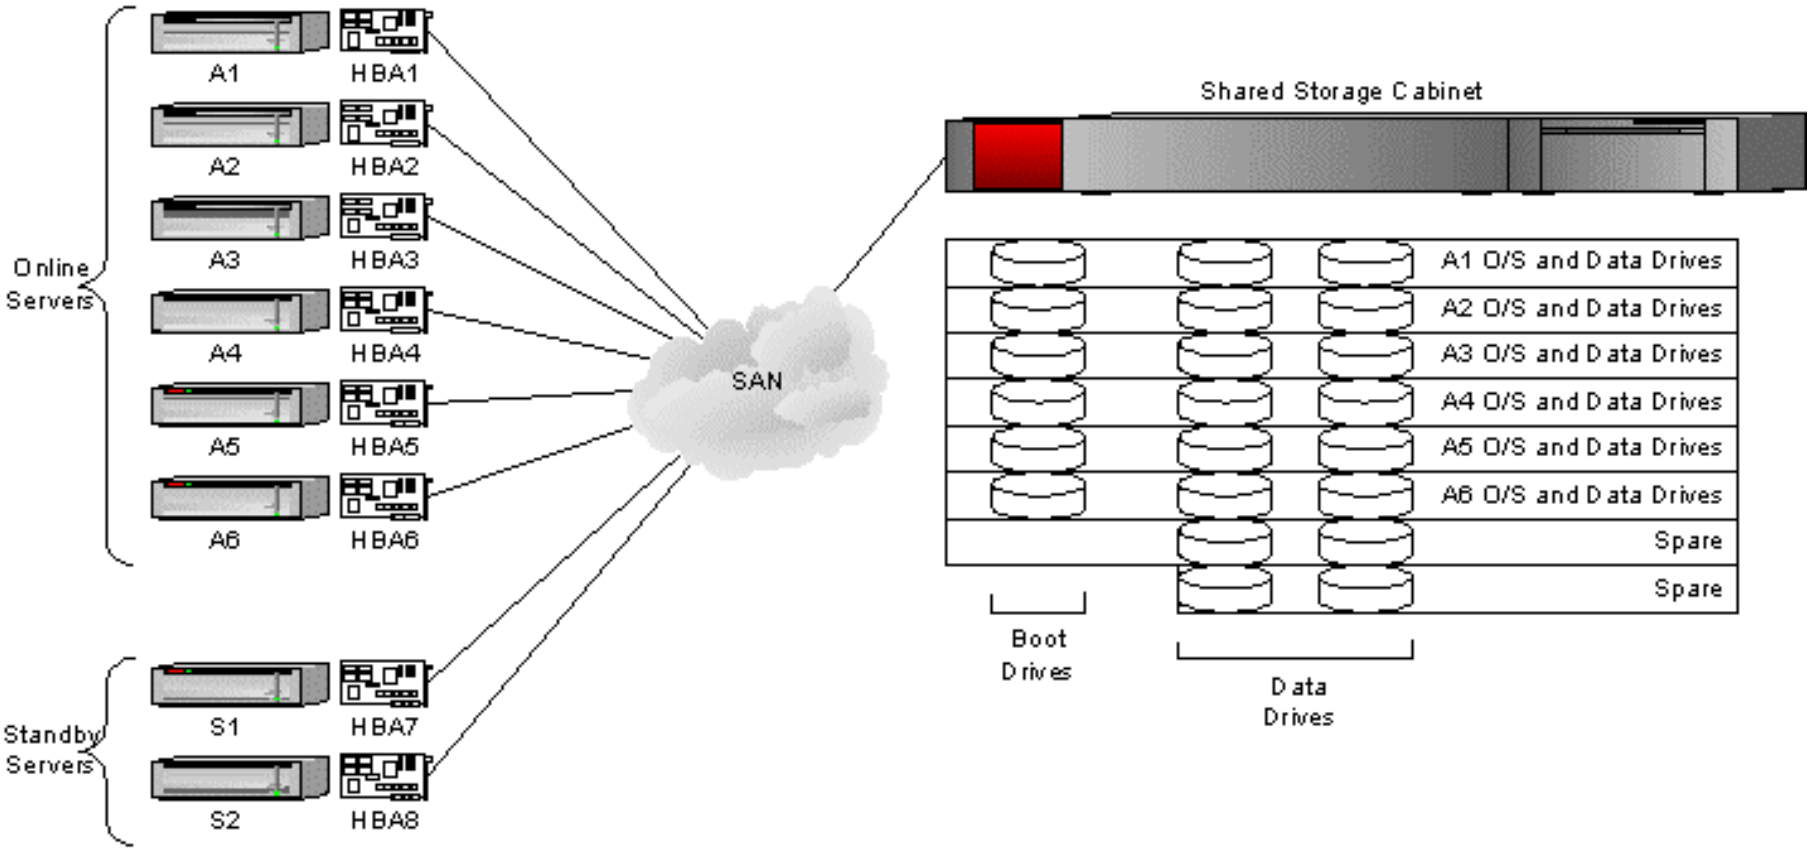

Storage Virtualization

- What is it?
  - A transparent abstraction layer between physical and virtual storage
- How it's done
  - Binds storage units into a "pool"
  - Divide "pool" into multiple "virtual disks"
  - Maintains mapping of physical to virtual disk blocks
  - Add new storage to pool online while applications continue to run

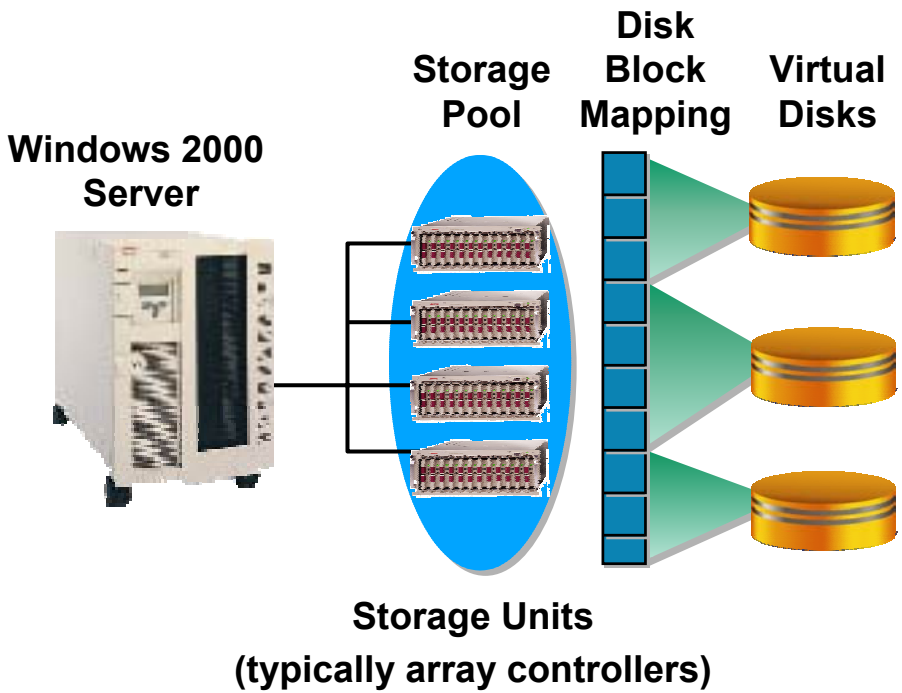

Cross-reference disk groups on EVA

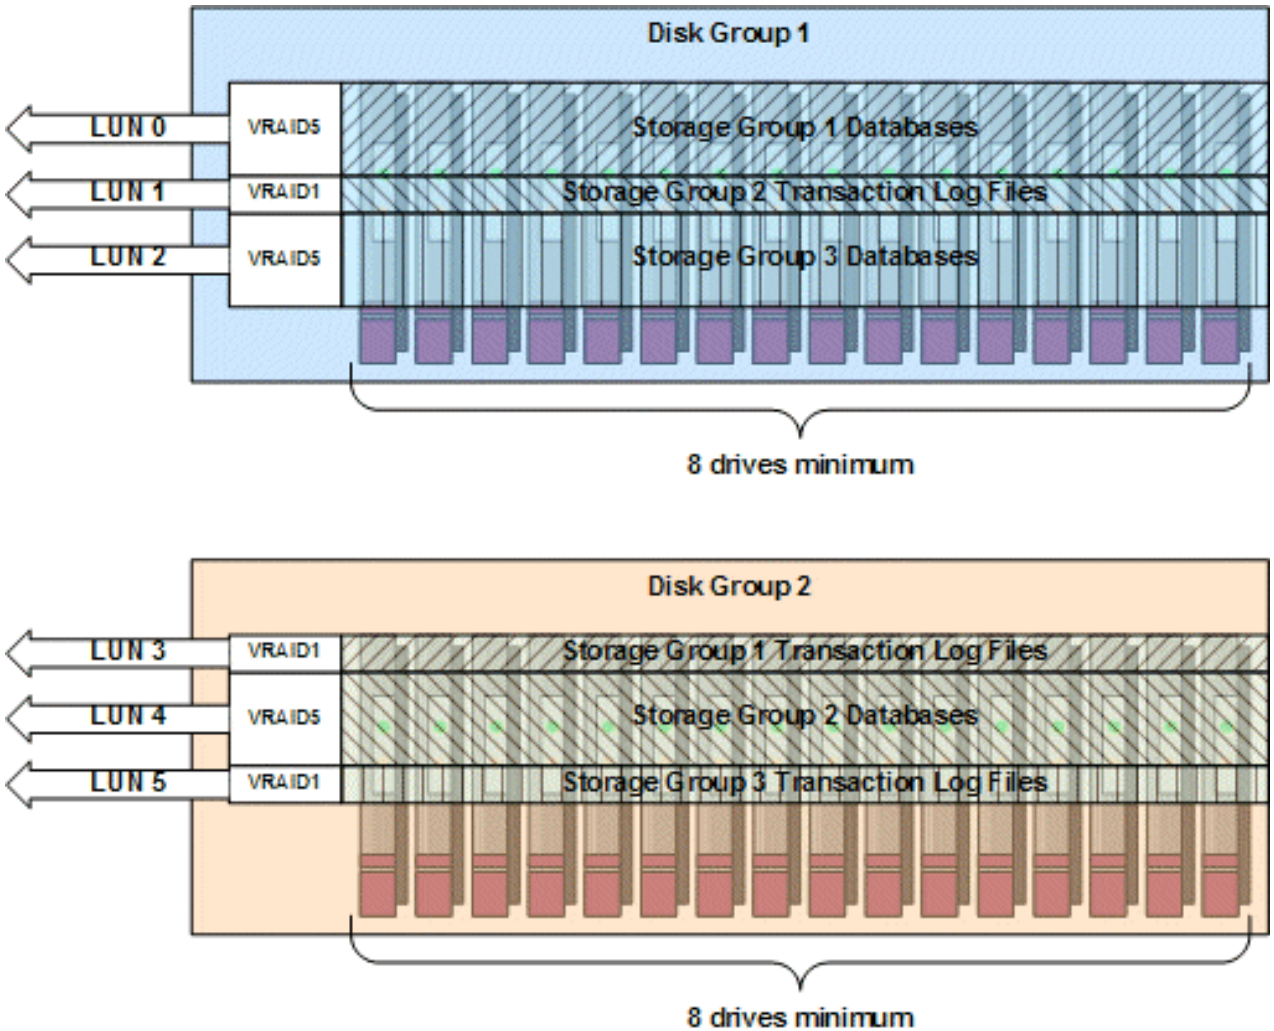

Virtualization: sample

- 8 disks group, including VRAID1 (mirroring) and VRAID5 (distributed) parity
- Adding 4 disks to the group (online)
  - → Capacity increased (raw: 4 x disk size)
  - → Throughput increase (net: 12 x 150 I/O/s = 1,800 I/O/s)

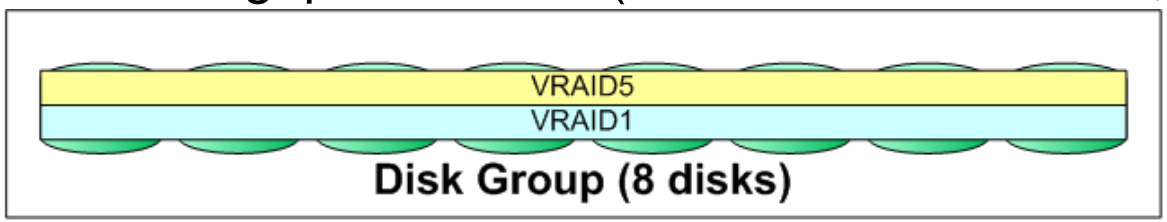

=

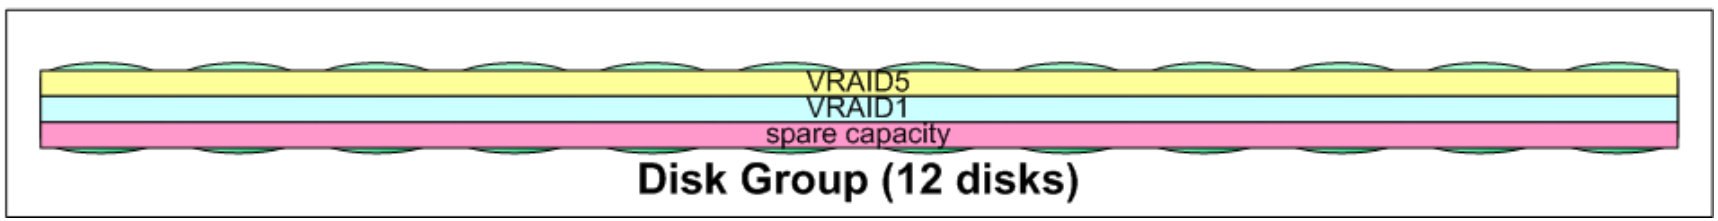

How do Snapshots Work?

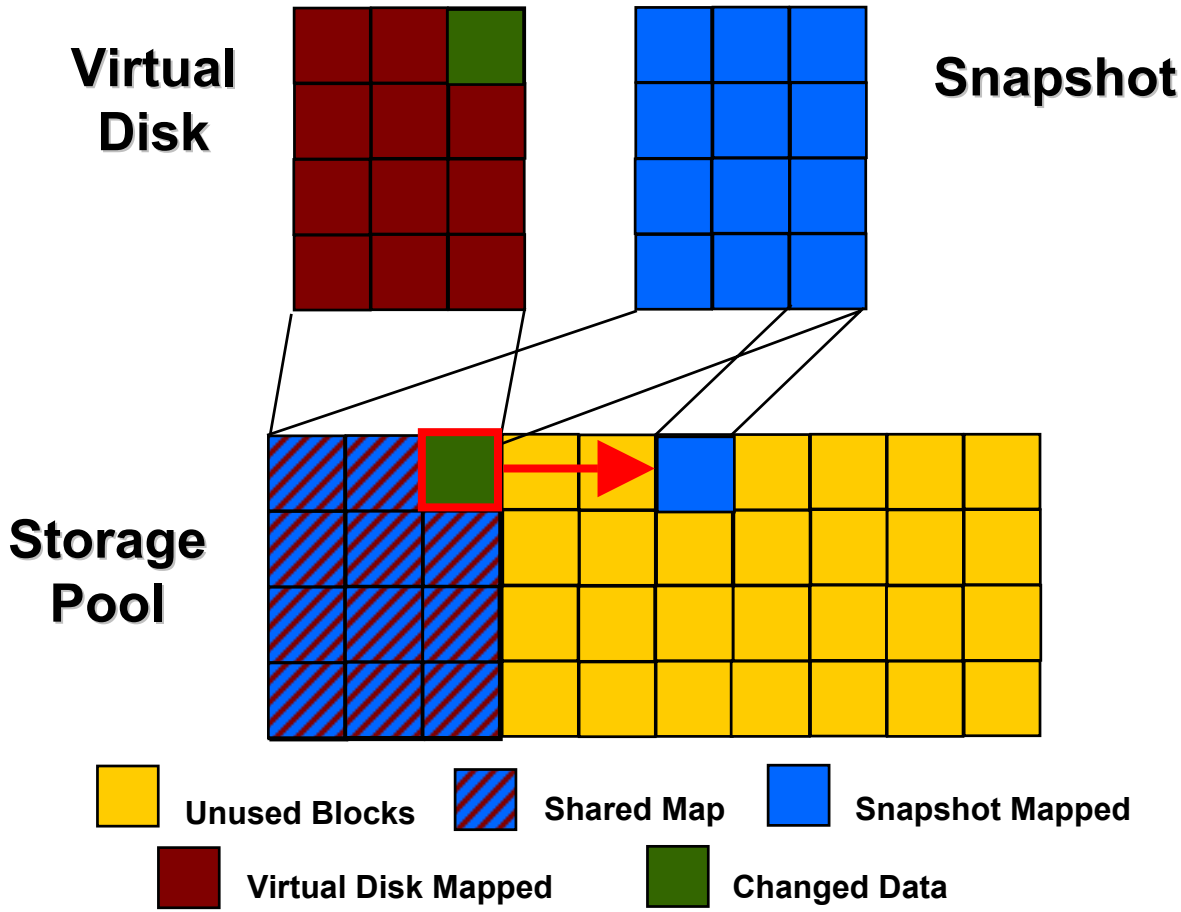

Clone: physical duplication

- Mirror and split

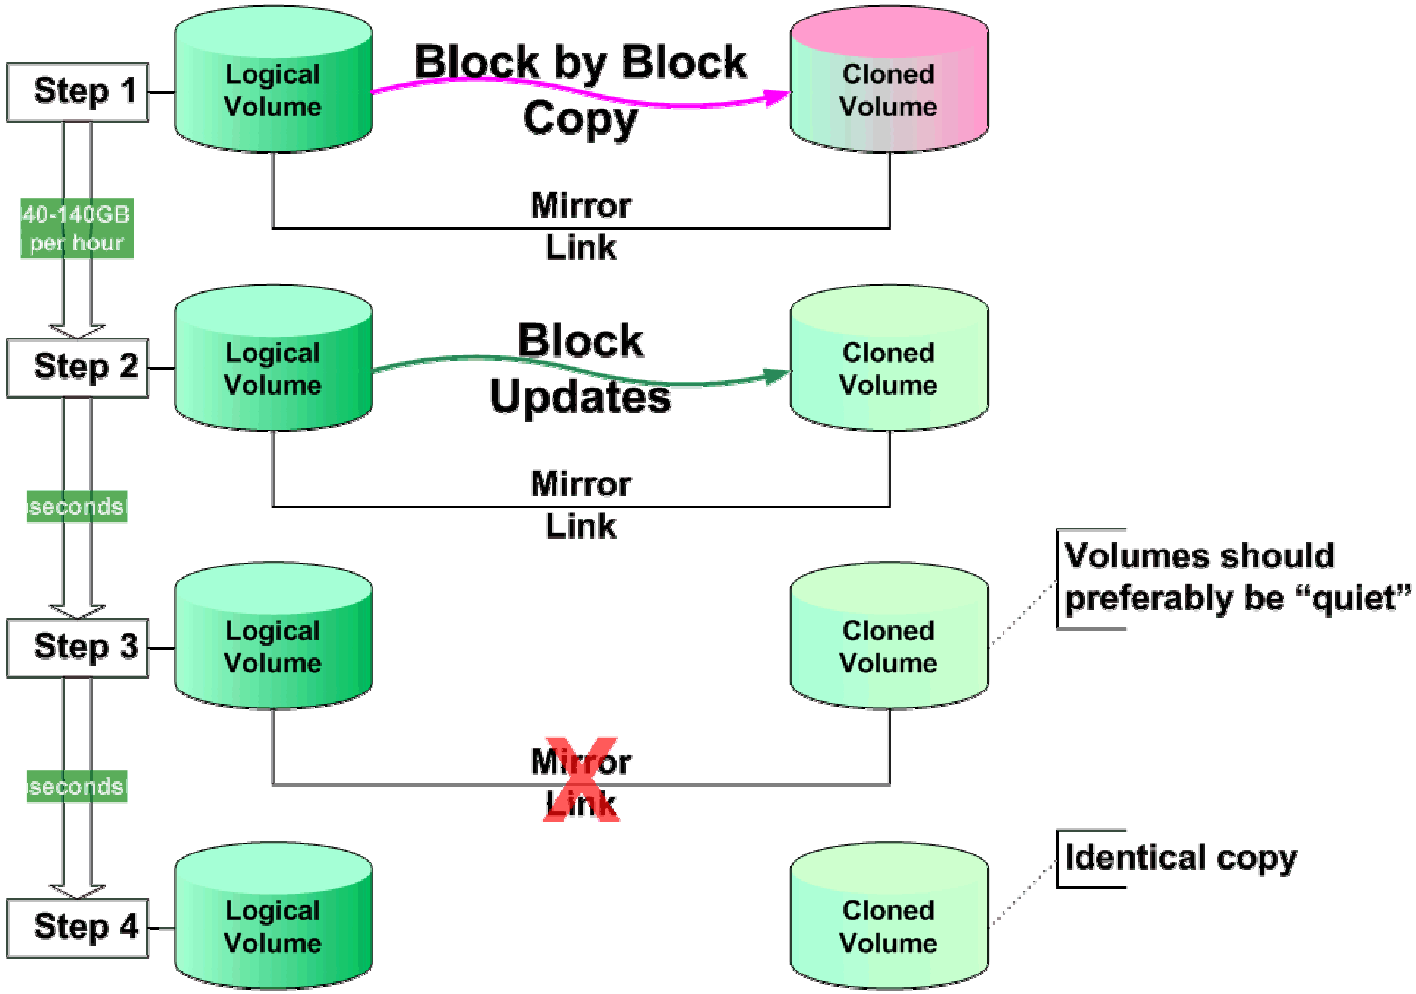

Summary

- .NET ecosystem includes full range of platforms, from mobile to big iron
- Same rules of thumb for server / storage deployment apply as with today's datacenters
  - Constantly improving processing power
  - Balanced architecture more important than fastest processor speed
  - Storage is still key consideration for the back-end tier
  - New clustering and availability technologies also emerging
- Device type will help to drive deployment technologies, which will help to drive infrastructure choices

Interex, Encompass and HP bring you a powerful new HP World.