


Implementing Windows 2000 in the DataCentre



Ian Bromehead
Global Initiative Manager Windows
MSO
Hewlett Packard

Joe Brown
Microsoft OEM Manager
Microsoft



Agenda

- Introduction
- DataCentre Platform Components
 - Servers
 - SAN & Storage
 - Windows 2000
- Application Scenarios
- Recommended Services
- Q & A



What's a DataCentre ?

- Consolidated IT Infrastructure
- Infocentre
- Performance Provider
 - OnLine & Batch transaction processing
 - Data Mining
 - Information Management & Protection
 - Scalable Connectivity through server farms
- Hosted Services DataStore



What's required in a DataCentre ?

- Scalable transaction capacity
- Very Large DB capacity
- High speed expandable connectivity
- Redundancy in IT architecture
- IT Management
- Dynamic load balancing
- Service Level Management processes




Reliable Scalability for the Enterprise Datacenter

**HP
NetServers
for the
DataCentre**



HP NetServers


DataCentres using 6-way and 8-Way




- Collaboration
- Messaging Consolidation




- Decision Support Systems
- Data Marts
- Databases



- Enterprise Resource Planning
- Online Transaction Processing
- E-Business



Blending Intel-Based Costs with RISC-Based Performance



Scalable Power for Growing Datacenters


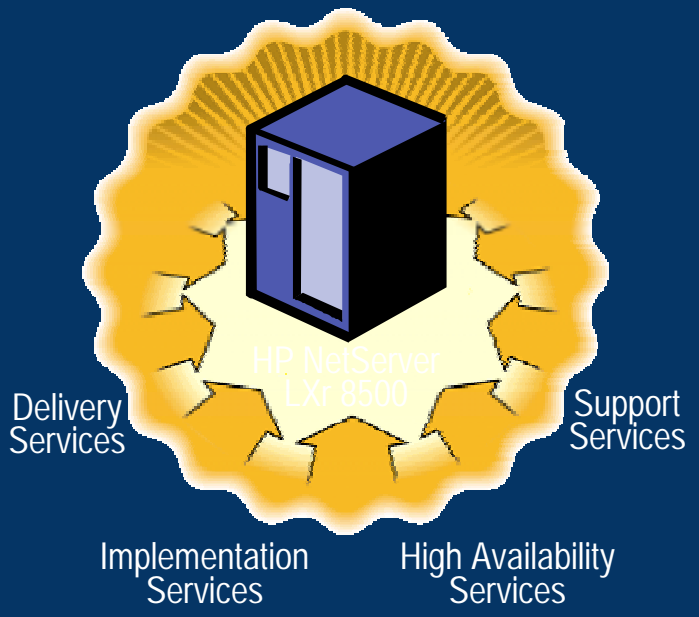


Efficient Control of Messaging Consolidation



HP NetServer

Why 6 & 8-Way from HP?

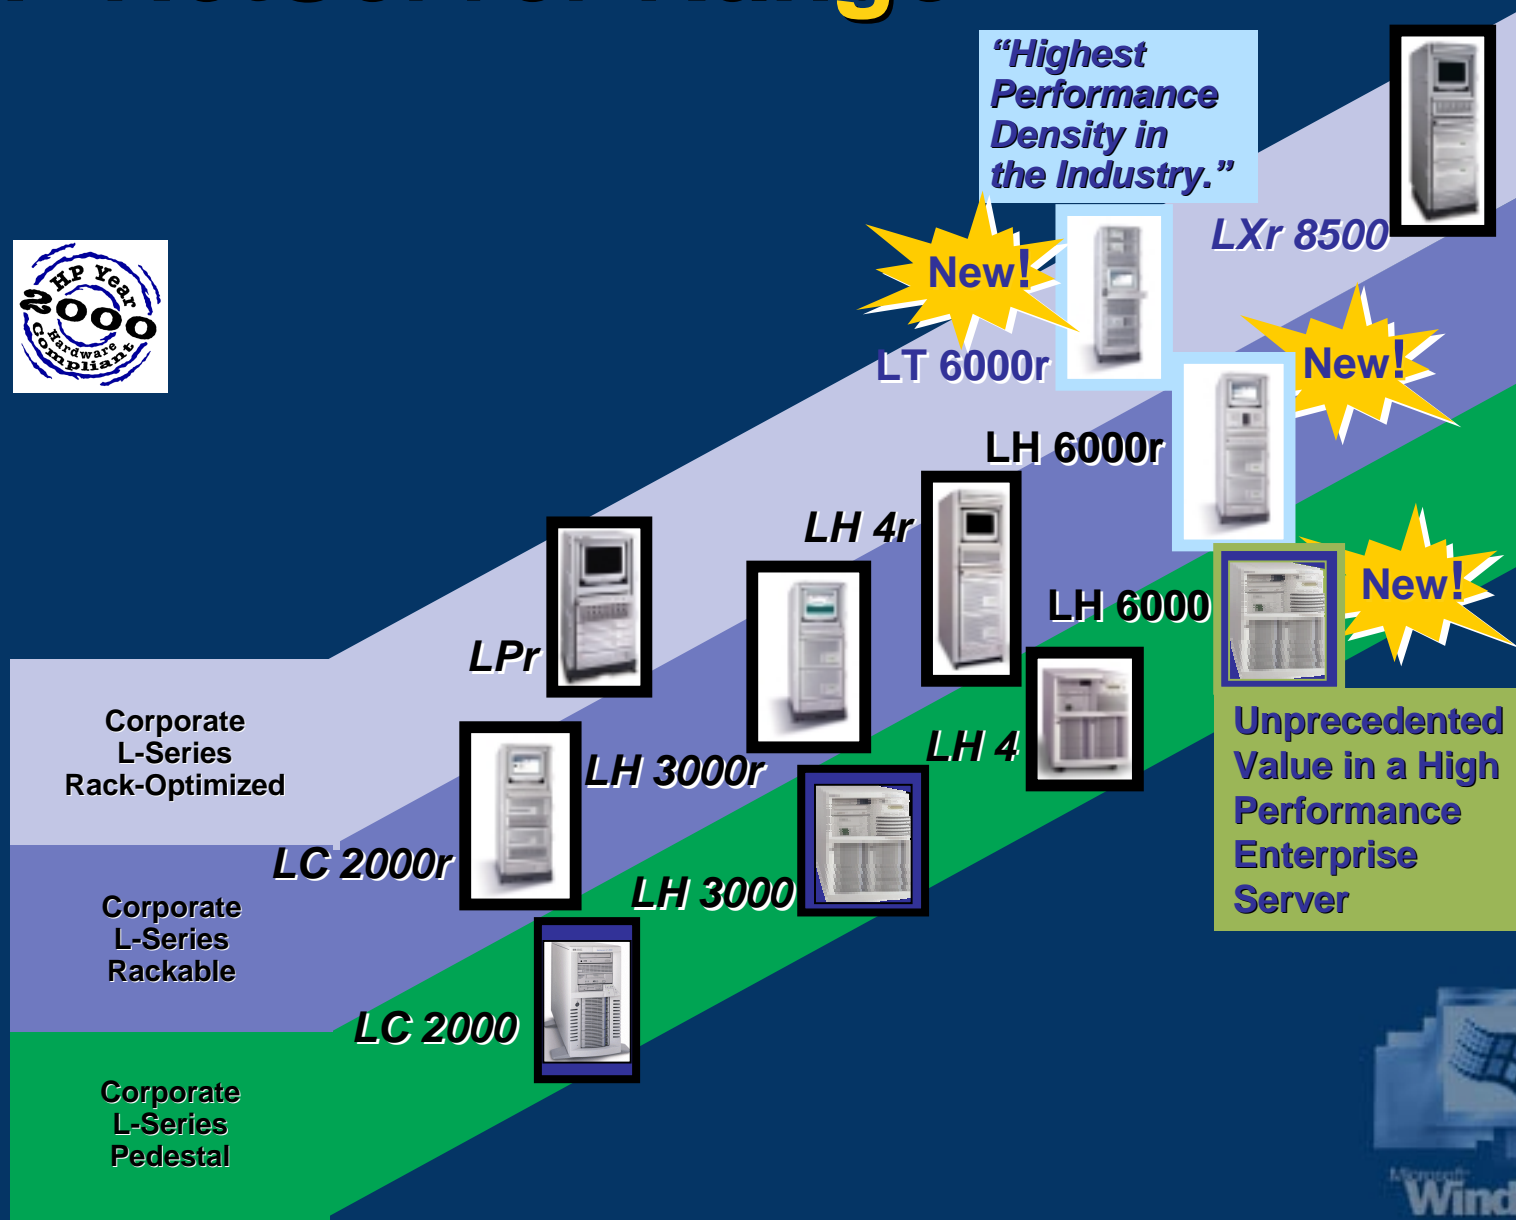


HP's Intel-Based 6 & 8-Way Solutions
Designed for Reliability through World-Class Services

HP and Intel's Relationship
Standards-Based Co-Development of Enterprise-Class Technology



HP NetServer Range



HP NetServers

User Value Proposition

- Providing scalable, fast solutions for enterprise resource planning, data marts or decision support systems
- Leveraging enterprise expertise to deploy and support mission critical computing environments
 - HP Select Express and Racks Express
 - HP Implementation Services
 - HP High Availability Observatory
 - HP Fault Notifier
 - HP SureStore One Button Disaster Recovery
 - HP Mission Critical Services & Support

hp netserver = performance



hp netserver = enterprise

HP NetServers

Scaling Enterprise-Class Reliability for Datacenters



HP Select Express & Racks Express

Mission Critical Services and Support Offerings

HP Implementation Services

HP Flexible Warranties

Camp Four: Services

Integration with Enterprise Management Solutions

HP OpenView ManageX/SE

HP TopTools for Servers

Camp Three: Software

HP TopTools Remote Control Card

Portfolio of Clustering Solutions

HP One Button Disaster Recovery

Pre-Failure Alerting for Components

Hot-Swap, Redundant Components

Camp Two: Hardware

Rack-Optimized, In-Box Upgrade Kit

HP Balanced Performance Architecture

Base Camp One: System Design



HP NetServers

“One Button Disaster Recovery” Tape Backup Solution

- The world's first hardware-based One Button Disaster Recovery tape backup solution
- “The most reliable way to recover from a disaster”
 - Easy implementation at remote sites
 - Full hard disk restore at the touch of a button
 - Available on the HP NetServer LXr 8500 with purchase of an HP SureStore DAT24 tape backup accessory



HP TopTools 4.5

Providing Intelligent Manageability for Flexible Datacenter Control

- Manage Servers, Desktops, Mobiles, Hub & Switches and Printers from One Application
- Built-In Device Discovery and Alerting
- HP TopTools 4.5
 - HP 9000 Auto Discovery and Alert Integration
 - Pre-Failure Alerting on Hard Disk Drives and Memory
 - Email Alerts
 - Optimized for Microsoft Internet Explorer 5.0
 - WFM 2.0 Compliance

The screenshot displays the HP TopTools 4.5 web interface. The top navigation bar includes 'Device Types', 'Topology', 'Custom Groups', and 'Search'. Below this is a table of devices with columns for Name, Type, Polling, Alerts, SNMP, DMI, WBEM, and Web. The table lists several HP NetServer devices with their respective status indicators (green, yellow, red) and icons for actions like 'Home', 'Refresh', and 'Delete'.

Name	Type	Polling	Alerts	SNMP	DMI	WBEM	Web
cadbury.cup.hp.com	Server, HP 9000	●	●	✓			
cranberries.rgv.hp.com	Server, HP NetServer	●	●	✓			
hpdmsx1.rgv.hp.com	Server, HP NetServer	●	●	✓			
hpdpc814.rgv.hp.com	Server	●	●	✓			✓
hpdpt043.rgv.hp.com	Server, HP NetServer	●	●	✓			
hpdsvr4.rgv.hp.com	Server	●	●	✓			
hpdweb.rgv.hp.com	Server, HP NetServer	●	●	✓			

Below the table, an 'Alerts' window is open, showing a list of alerts. The most prominent alert is a 'Critical' 'EMS Change Event: Value "2" for "/system/humUsers" (Threshold: change)'. The alert details include a description: 'The Event Monitoring Service (EMS) has generated an event for a monitored resource. A Problem Event occurs when the threshold condition defined for the resource is true. A Change Event occurs when the monitored value changes. Monitoring requests configured for notification when the value changes may generate this event.'

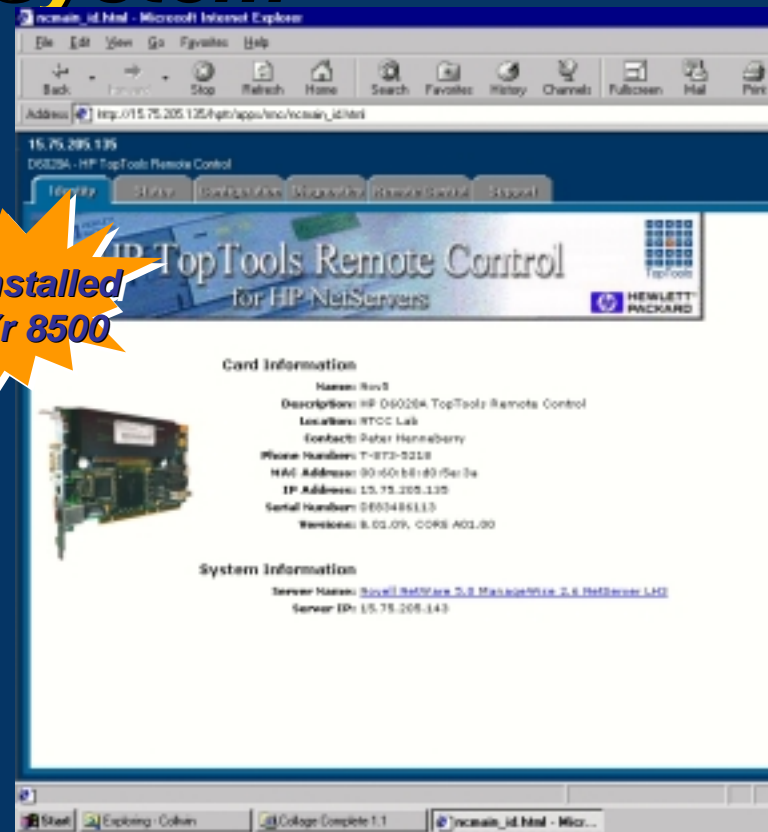
TopTools Remote Control Card

Remote Management of HP NetServers Regardless of System State

- Easy Access Anytime, Anywhere
- Faster Troubleshooting
- Powerful Remote Control
- Advanced Security

- Remote Plug and Play
- Group Actions
- Support for DHCP
- All New Features Require Only Firmware Update!






**Pre-Installed
in LXR 8500**



HP NetServer Support

Enterprise Expertise for Datacenters

New Services!

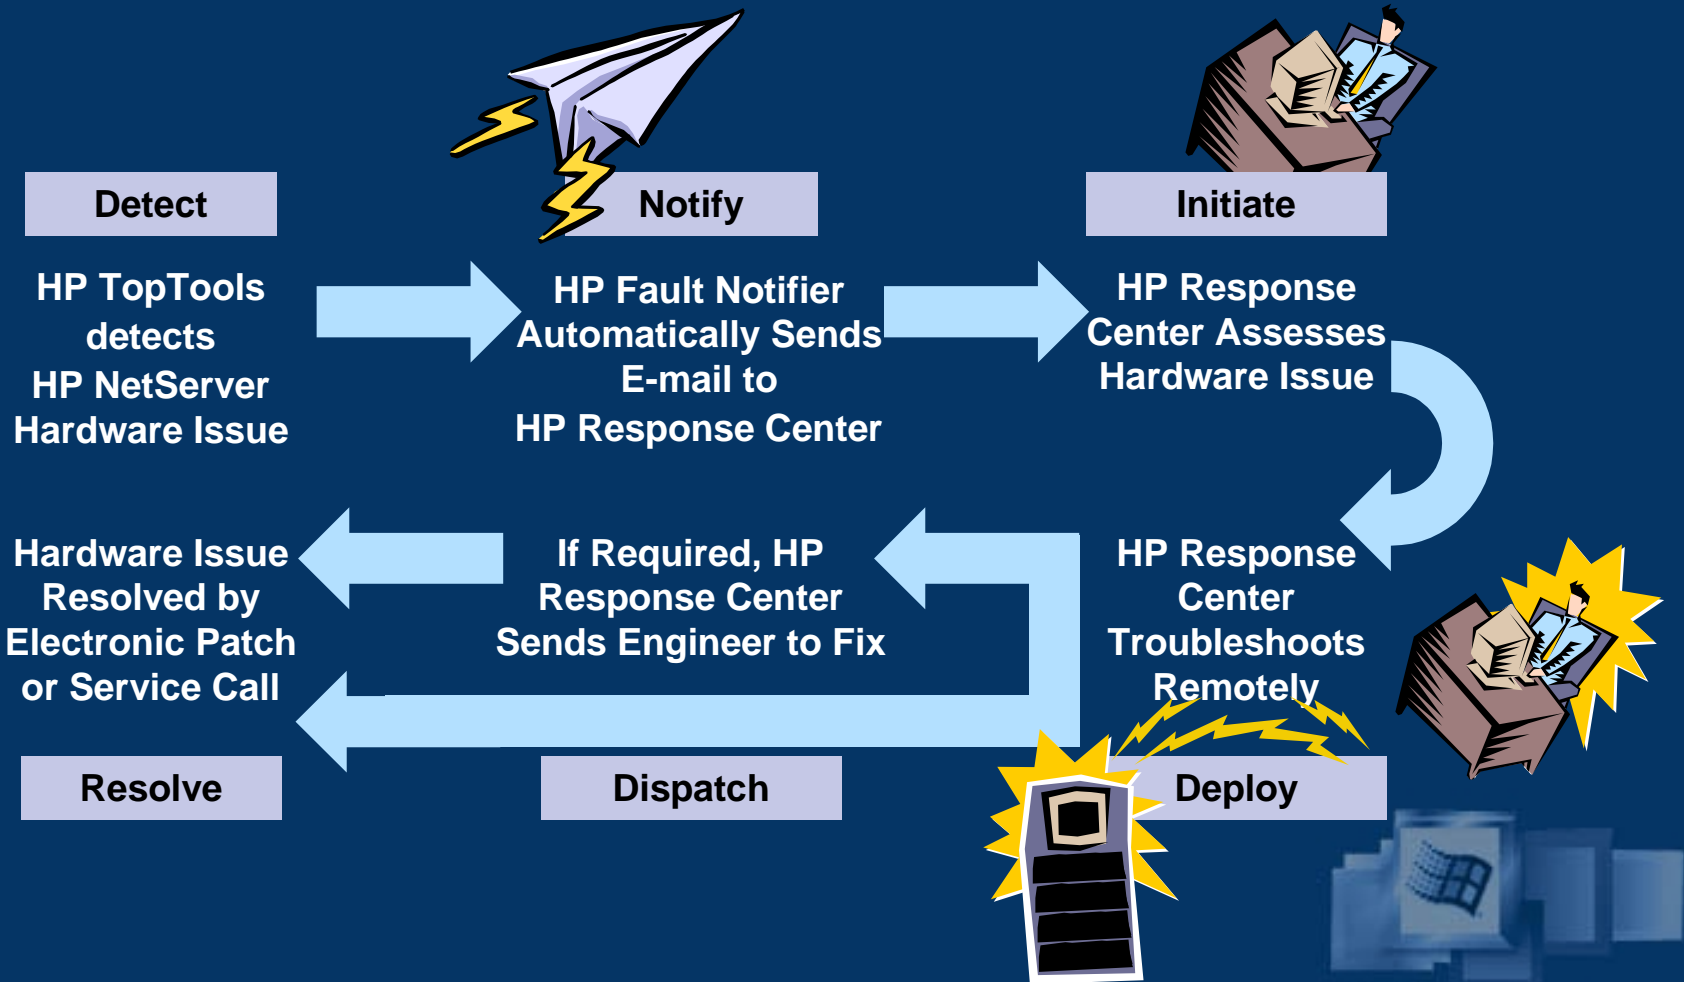
Hardware SupportPacks	Software SupportPacks	Implementation Services
<p>Installation and Network Configuration Services</p> 	<p>HAO Enhanced!</p> <p>Critical System Support for NT</p> 	
<p>Hardware Call-to-Repair (3 & 5 years) Immediate Response</p> 	<p>Personalized System Support</p>	
<p>24 x 7 (3 & 5 years) 4 Hour Response</p>	<p>Comprehensive Technical Support</p> 	
<p>Same Business Day (3 & 5 years) 4 Hour Response</p>		
<p>Next Day On-Site (3 & 5 years)</p> 		
<p>Flexible Warranties (1 & 3 years parts or 1, 2 & 3 yrs parts & labor)</p>		

Available Now!



HP NetServers

HP Fault Notifier Maximizes Uptime with Reactive Support Capability



HP NetServers

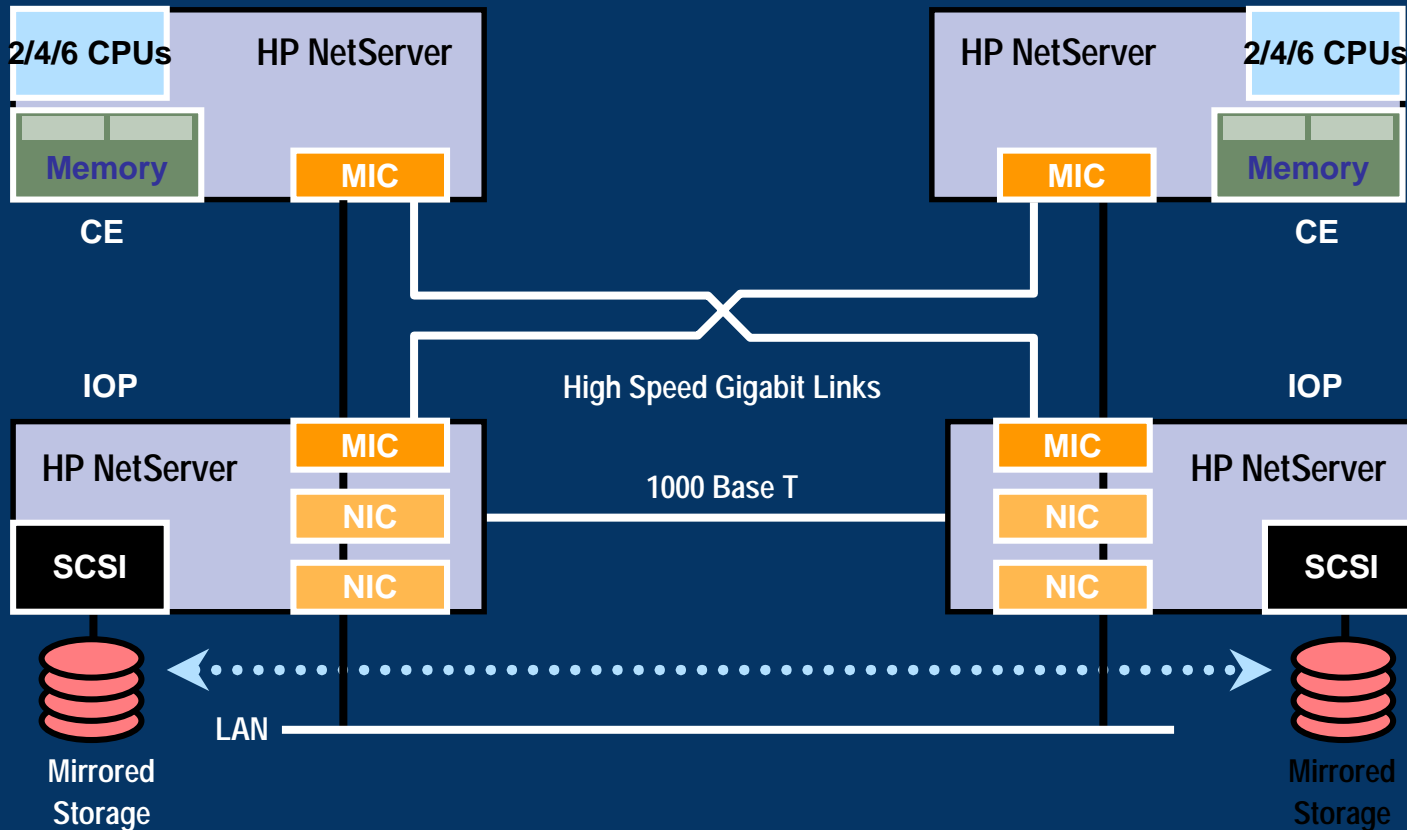
Reliable Datacenter Solutions Delivery

HP Select Express	HP Racks Express
<p>Featuring</p> <ul style="list-style-type: none">• Integration of HP NetServer LXr 8500 Select Express models, external storage, accessories and other HP NetServer Select Express models into a fully configured HP Racks Express Solution• System cabling and labeling for easy port and slot identification• Hardware RAID configuration• Inside delivery to datacenters• On-site Customer Engineer installation <p>Benefits</p> <ul style="list-style-type: none">• Easy customization of HP NetServer LXr 8500 solution• HP quality rack-optimized solutions• Reduces installation, configuration and integration issues• Solution receipt to power-on in hours, not days	
Standard Availability (14 Days)	Standard Availability (21-28 Days)



Assured Availability Fault Tolerance

Best In Class



Strategy	Multiple servers act as a single virtual server. End users are “masked from failures without interruption of service.
Strengths	Protects against hardware and many software failures.
Weaknesses	High hardware redundancy. Some performance overhead.
Examples	HP NetServer AA 6200 (2 CPU) and 4000 Solutions

Assured Availability Fault Tolerant DataCentre

- 2 way MP Support for Compute Elements
- Faster PCI Cards & Communication Protocols
- No SSDs in tuple link & now Gigabit Ethernet
- ManageX Policies and OpenView Integration
- 500- 20km Meters Between Tuples
- Fiber Connectivity

AA 6200 Example

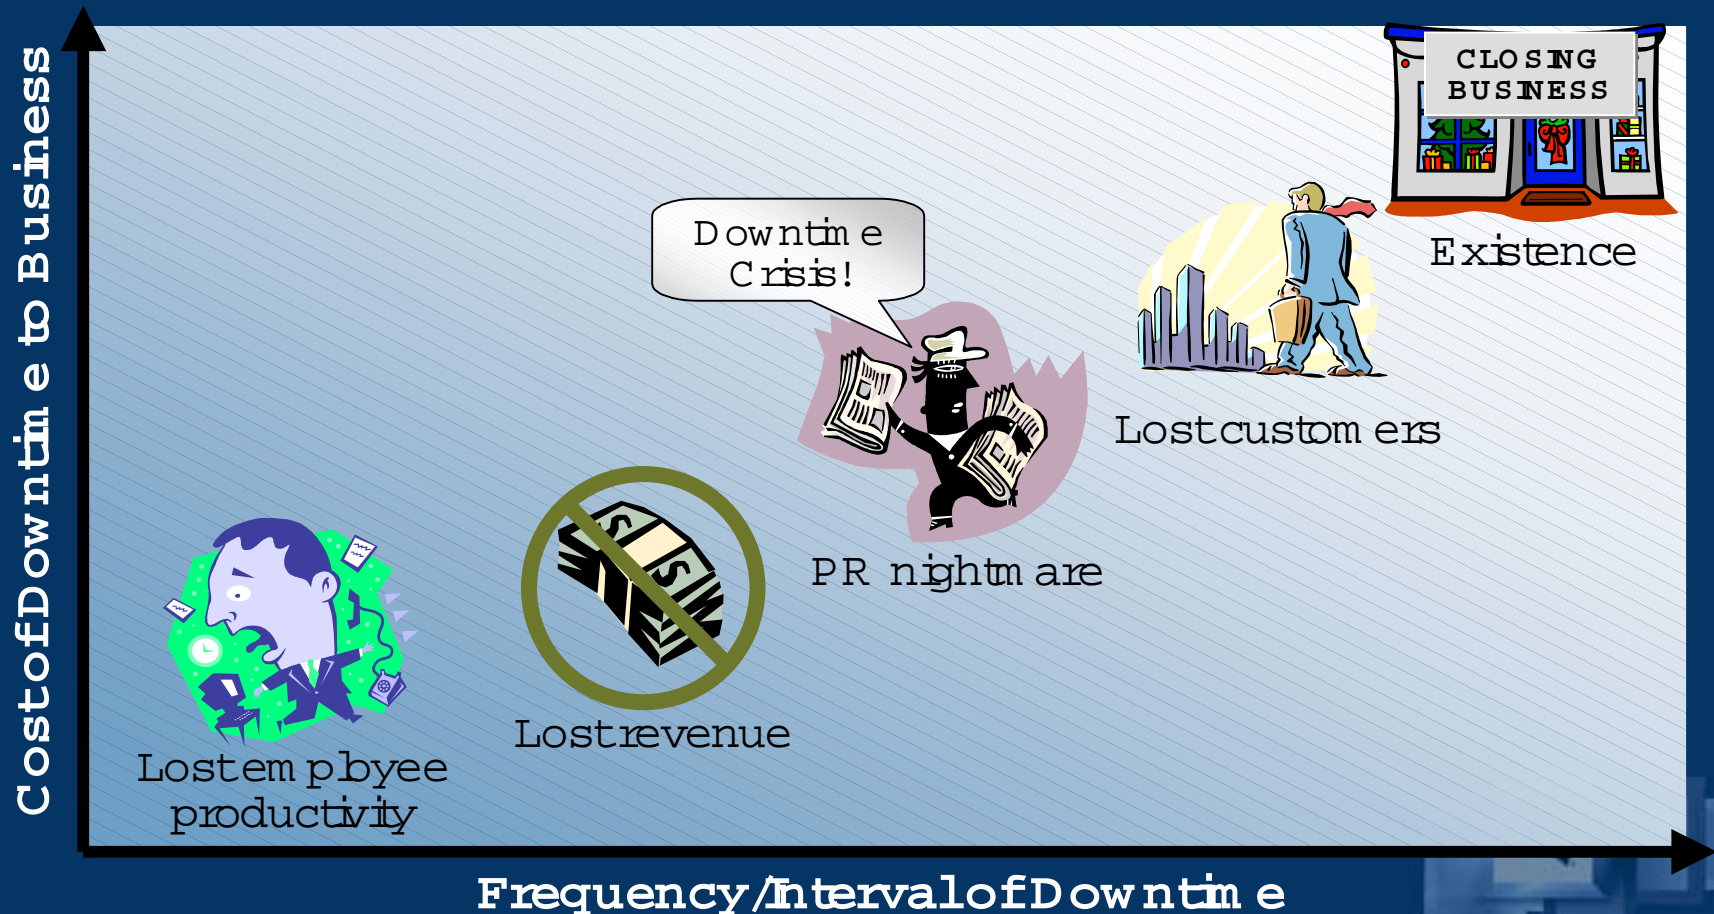
- 2 HP NetServer LC2000 Systems (CE) and 2 HP NetServer LH3000 Systems (IOP)
- External Storage
- Tape Backup Devices
- Advanced System Management
 - Endurance Manager
 - ManageX Event Manager Integration
 - TopTools (IOP Only)
 - Remote Control Card
- Mission Critical Services
 - Hardware Call-to-Repair Support Service
 - Priority Recovery for HP NetServer AA 6000 Software and Windows NT Support

AA 6200 rack



High Availability of Data

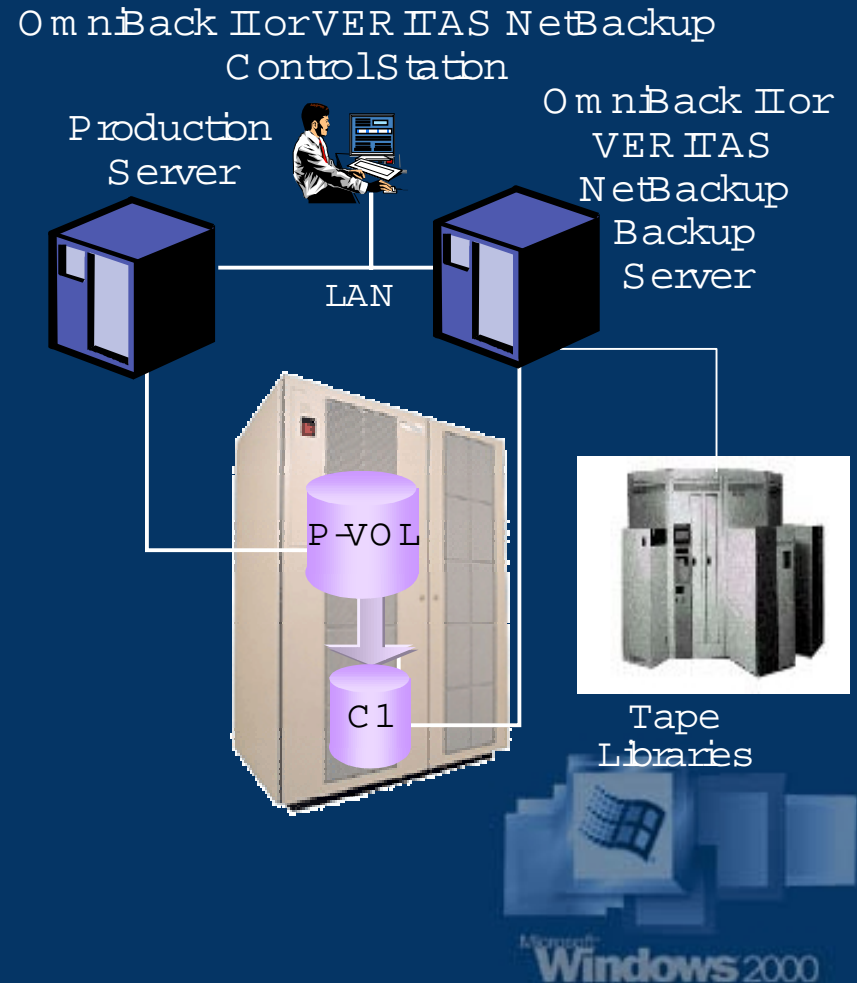
Downtime is not an option in the internet age



Data Management Solutions Ensuring Business Continuance

Zero Downtime Backup and Recovery is required


- Data protection at all times
- Continuous access to business-critical data resources
- No application performance degradation
- Fast recoverability in case of data loss or data corruption
- Full automation and integration resulting in lower TCO



Data Management Solutions Ensuring Business Continuance

Load Balancing and Automatic Path Failover

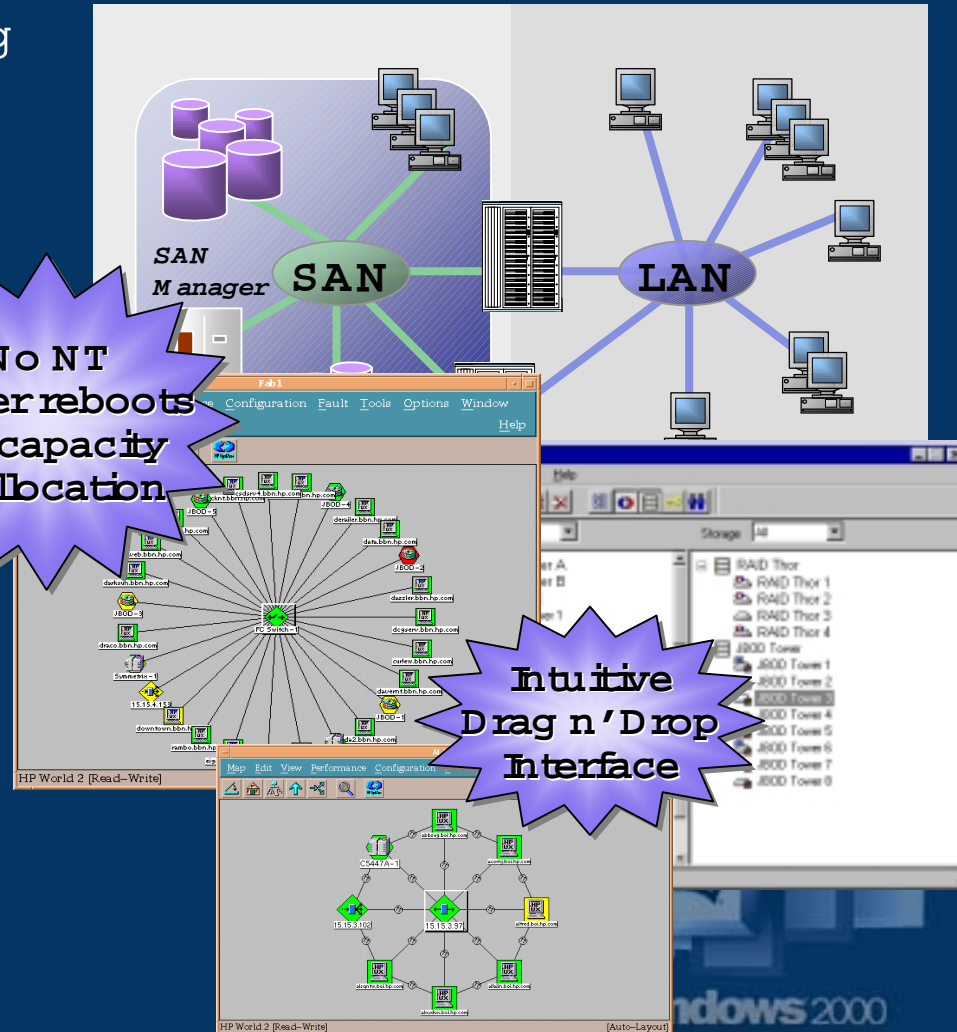
- Eliminates HBAs, cables, LAN components and host interface as single points of failure
- Distributes I/O across multiple paths to avoid overloading a single path to improve I/O performance
- Enables MSCS failover if an HBA, cable or SAN component fails



Data Management Solutions Ensuring Business Continuity

Operations & Change Control through efficient Management


- Single-view administration reducing administration time and costs
- Storage is assigned dynamically; network remains 100% available during SAN administration
- Improves LAN performance by offloading storage & backup traffic
- Provides an extremely scalable storage pool
- Integration with Enterprise Management Tools



Data Management Solutions Ensuring Business Continuance

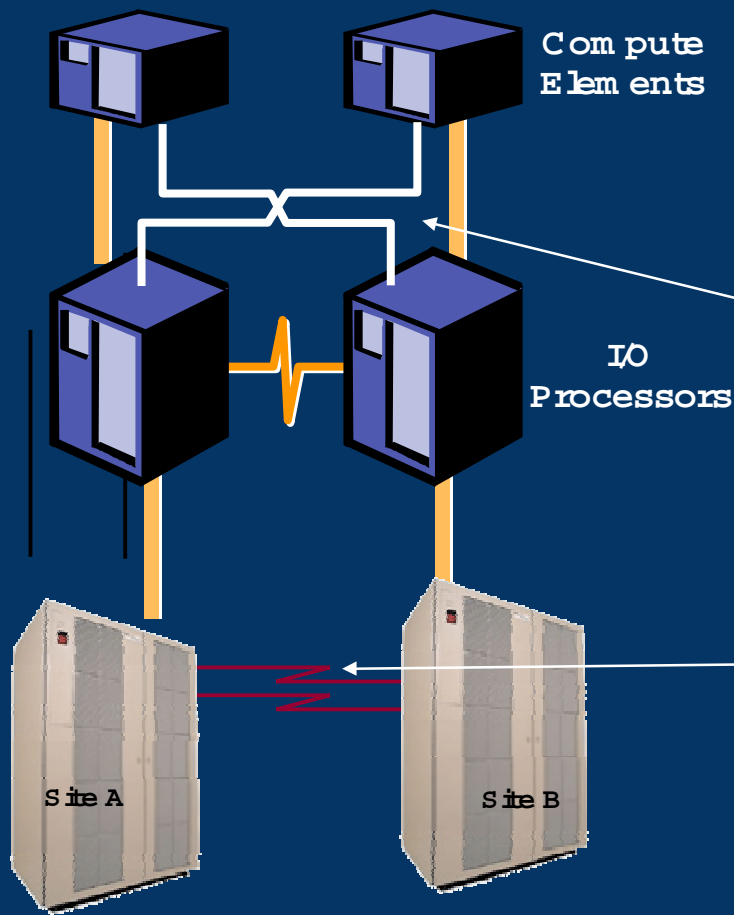
99.999% means the most reliable data protection

- **Designed to be ALWAYS UP!
ALWAYS AVAILABLE!**
 - No single-points-of-failure in Storage bays or SAN
 - Non-disruptive online cache, firmware upgrades
 - 24-hour phone-home capabilities
- Local and remote data mirroring for disaster recovery



Data Management Solutions Ensuring Business Continuance

Full Application Disaster Recovery Solutions



- Logical devices mirrored between two separate arrays in widely separated locations
- Assured Availability provides kilometers server separation
- A higher-availability alternative to MSCS
 - Looks like single logical node
 - Eliminates fail-over and fail-back
 - Does not require cluster-aware applications
- Full application, host, and storage fail-over over distances up to 43km with Microsoft Cluster Server or other clustering solutions
- Fulltime processing and continuous data access
- Disaster tolerance and constant performance
- Reduced downtime



Windows 2000 Is Doing 24x7

99.999+% under load of
~30 million hits per day

24x7 since Jan 1
5x9 on standard hardware
with no web server failure

NASDAQ

Uptime: 99.95 %
Service Levels: 24x7x365

Source: Aberdeen Group interviews with dotComs running Windows 2000



Buy.com was experiencing several connection and reliability related problems daily. Today most of them are gone! Servers which previously could handle only **25** concurrent connections are now handling **800**

Windows 2000 Server Family Positioning



Windows 2000 Advanced Server
The operating system for Line of Business & .com

Target Usage

- Mid-scale Line-of-Business & .com (backend)
- Distributed deployment of messaging and file & print

Scalability

- 1-8 processors
- 8 GB

Reliability & Availability

- 2-node clustering
- Component testing

Serviceability

- Premier / Alliance recommended
- Uptime Guarantees available from OEMs
- Configuration Flexibility



Windows 2000 Server Family Positioning



Windows 2000 Datacenter Server

The operating system for the most demanding levels of availability and scale

Target Usage

- Large scale Line-of-Business & .com (backend)
- Server Consolidation

Scalability


- 1-32 processors
- 64 GB

Reliability & Availability

- 4-node clustering
- System testing (Datacenter HCL)
- Coordinated (HW & SW) change control


Serviceability

- MCSC for Datacenter
- Limited Configurations
- Other services offered by OEMs



Windows 2000 Datacenter

Product Overview



Joe Brown
OEM Manager
Microsoft Corp



Applicable Laws of “Physics”

■ Moore’s Law

- Capability of silicon doubles every 18 months
- Applies to chips: (not directly to the interconnections)
 - Processors, memory controllers, etc

■ Maxwell’s Equations

- *“You can’t fool an electron”*
- Applies to connections between the chips (among other things)

■ Gilder’s Law

- Networking bandwidth increases 3X / year for the next 25 years

■ Amdahl’s Law

- Total execution time = $\frac{\text{parallel portion}}{N} + \text{serial portion}$

And data grows faster than all of the above




IO Performance



Windows 2000 I/O Scaling

Core and Network

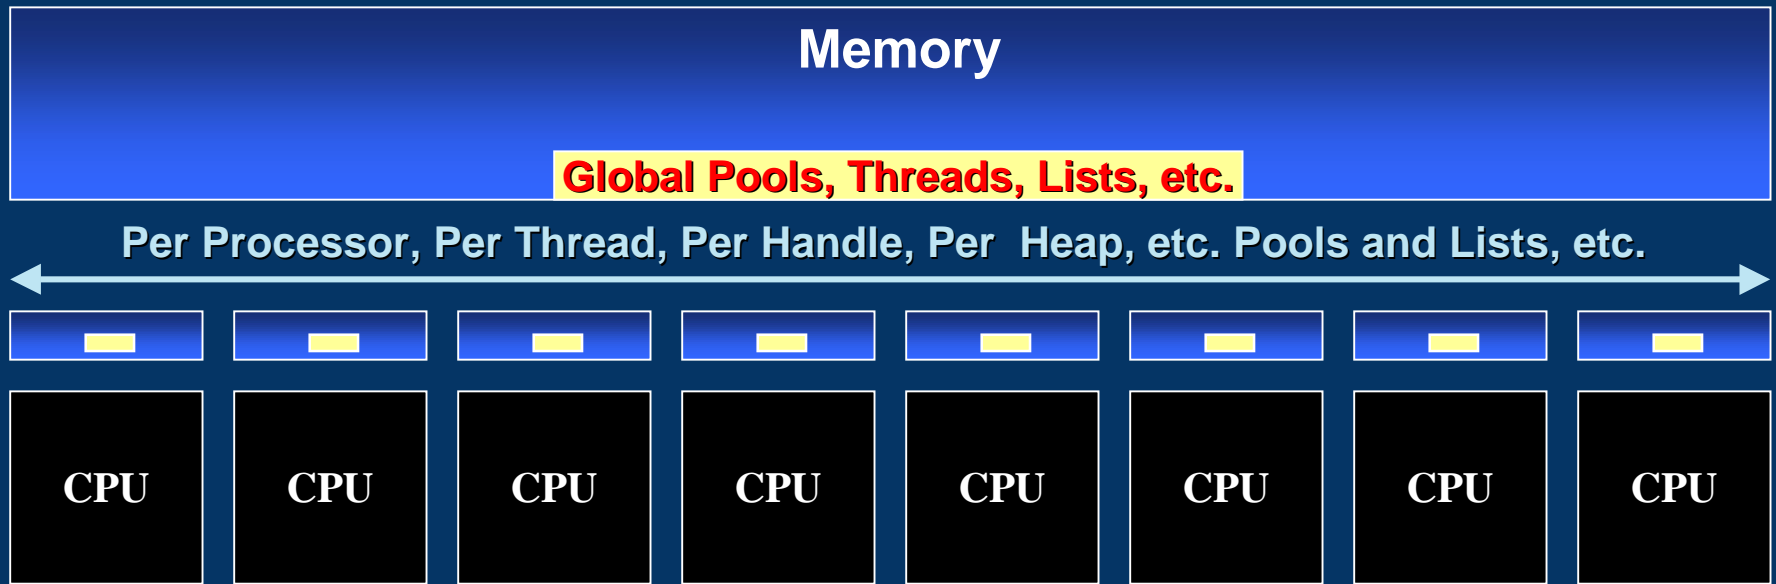
- I/O drivers - path lengths reduced, asynchronous design *
- NDIS - deserialization on SMP**, batch sends, jumbo frames
- Interrupt I/O affinity - improves locality of memory references
- NTFS improvements - reduced context switches
- SCSI - reduced contention on spinlocks
- High performance I/O - scatter/gather – reduced map registers
- TCP/IP stack performance improvements
 - Offload to HW
 - Code path length optimizations
- Large frame support
 - More efficient (64KB vs. 1.5KB packets) on the wire
 - Processing overhead reduction
- Selective acknowledgments
 - Allows the receiver to request from the sender only those packets that were missing or corrupted during initial delivery
- Benefits – less CPU time spent ‘pushing bits’ for IO, more cycles used for applications
- Benefits – network bandwidth more efficiently utilized, less ‘chatter’, more pipe left for apps and data



Symmetric Multi-Processing



Windows 2000 SMP Scaling



- Per processor:
 - Look-aside lists for memory allocation (hot lists)
 - Paged + non-paged pool look-aside lists
 - Thread pools
 - I/O completion ports
- Benefit – ‘sweet spot’ for NT increases beyond 6-8 CPUs



Windows 2000 SMP Scaling

- **PFN lock**
 - AFD improvements
 - File system cache VA space increased by 50%
 - LRU algorithm reduces PFN and dispatch lock use
- **Dispatch lock**
 - Hold time of the dispatch lock reduced
 - Fibers reduces dispatch lock use >2x per IO
- **Queued Locks**
 - Anywhere lock could not be made more granular
- **Handle improvements**
 - Global handle executive replaced w/ per handle sync
 - Less contention, fewer context swaps
- **Fibers**
 - Less memory cost than threads
 - Less context 'switch' cost than threads
- **SMP Design throughout***



Memory



Windows 2000 and PAE

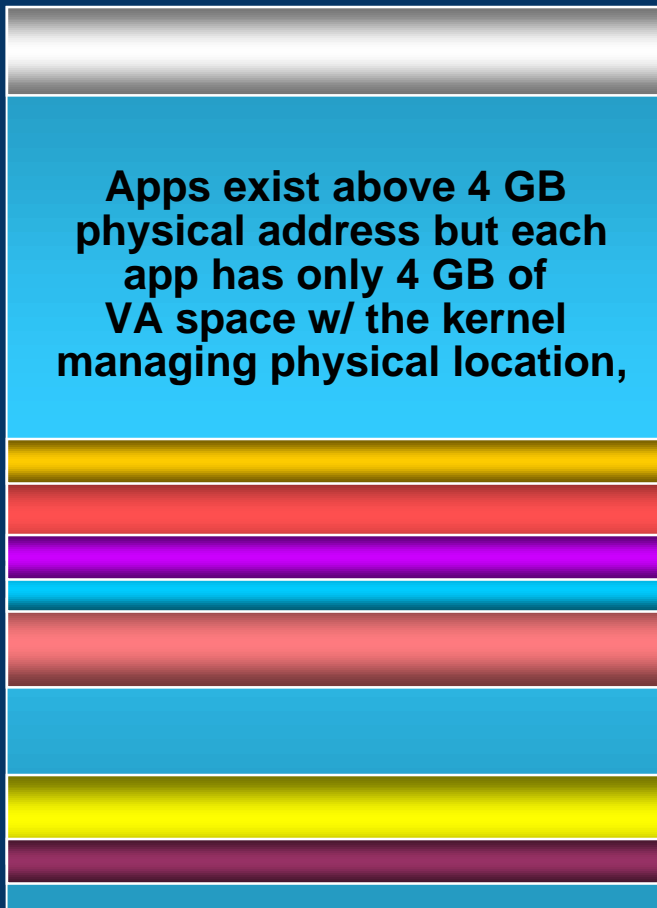
Memory Scaling

- **Address windowing extensions (AWE)**
 - Allow Apps with 32-bit VA limits to address >32-bit PA
 - Non-paged memory and window views
 - Extremely fast map/remap capability to be implemented
- **AWE API**
 - VirtualAlloc w/ MEM_PHYSICAL flag
 - AllocateUserPhysicalPages
 - MapUserPhysicalPages
 - FreeUserPhysicalPages
- Retains many of the NT MM architecture benefits*



Windows 2000 Large Memory

Consolidation Server



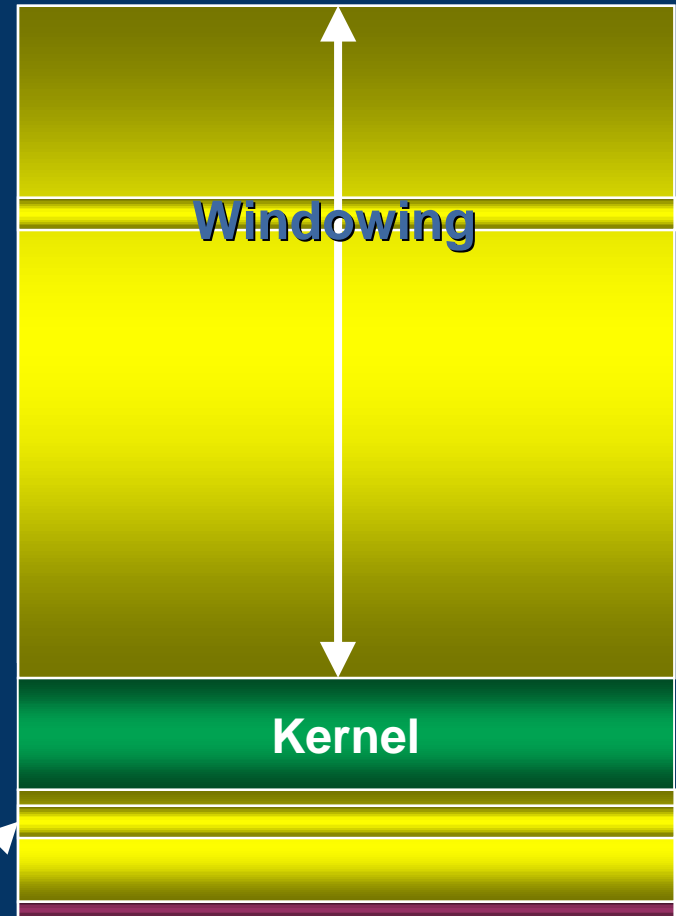
No change in applications

Single AWE enabled app

Physical Space

MapUserPhysicalPages

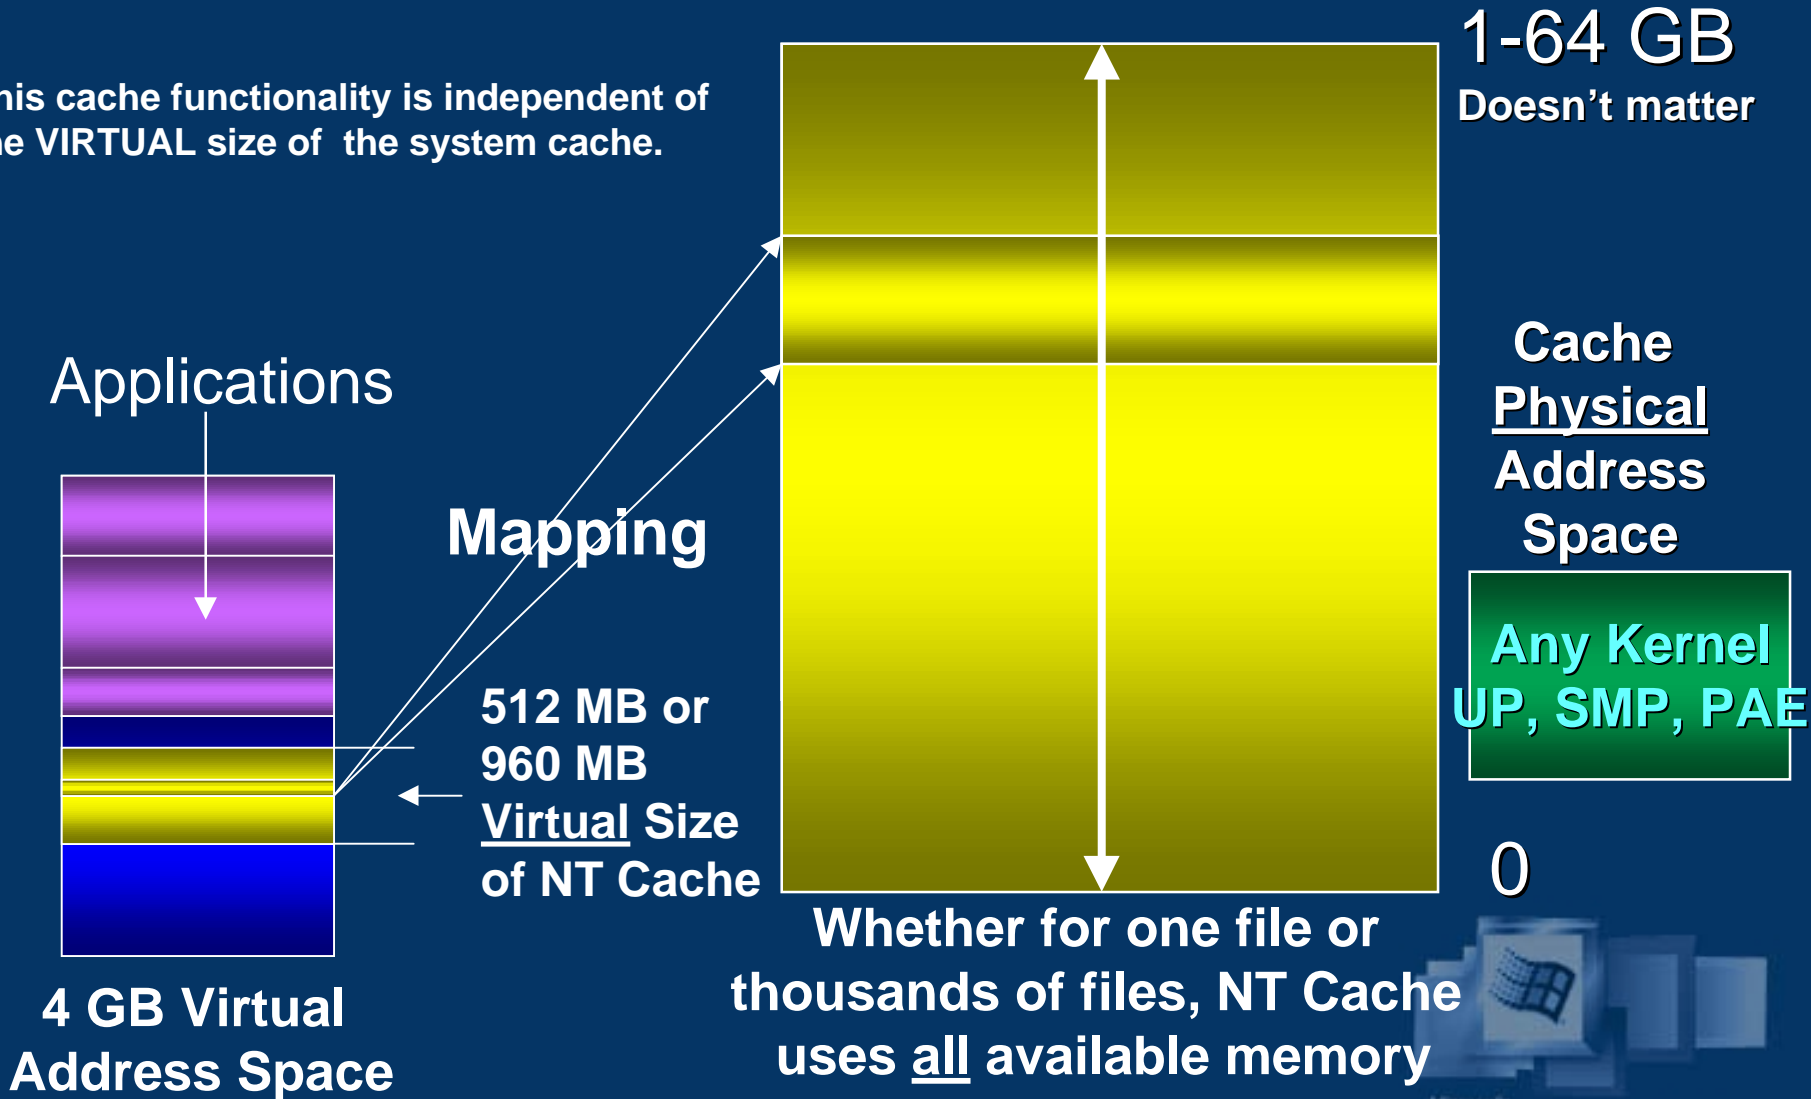
AWE Region



Application re-coded for AWE

Windows 2000 Caching

This cache functionality is independent of the VIRTUAL size of the system cache.

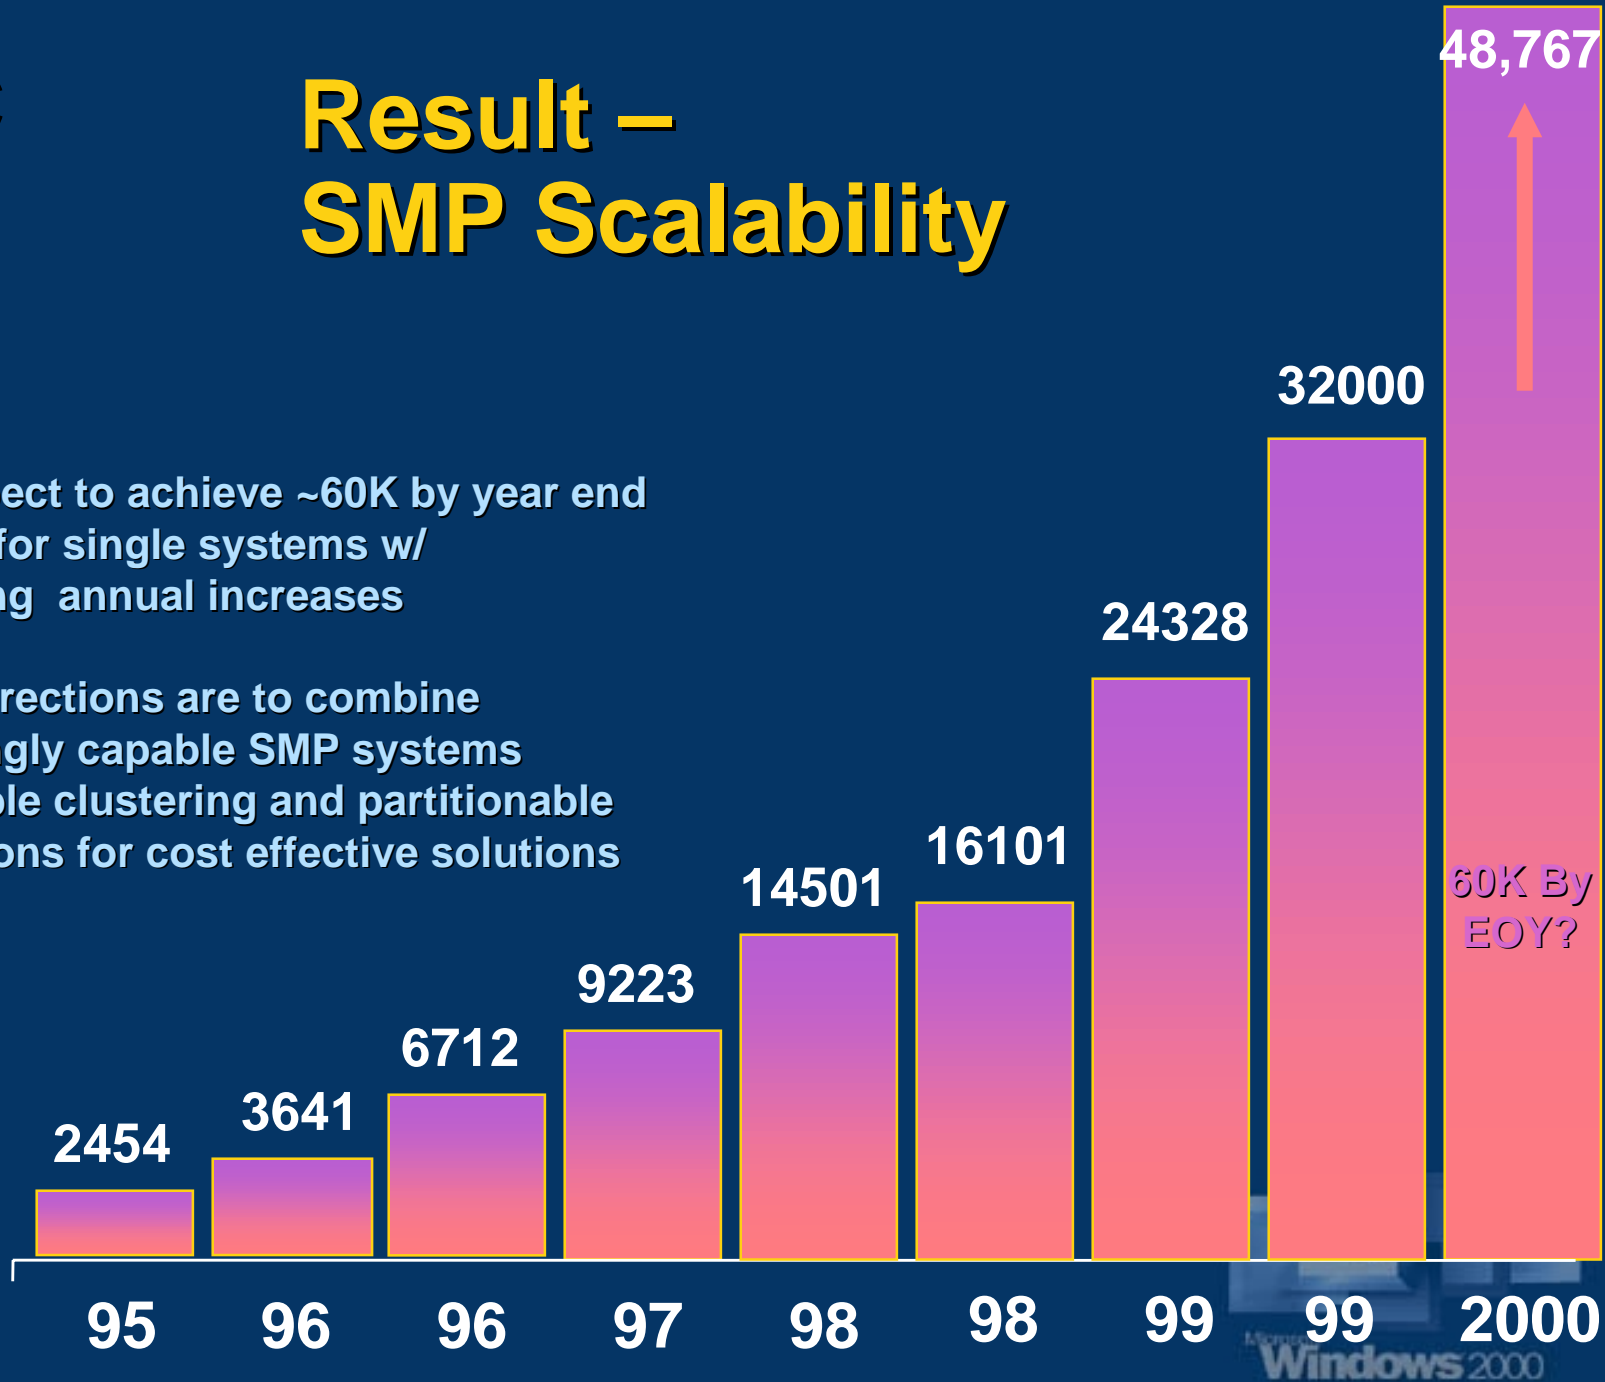


tpmC

Result – SMP Scalability

Fully expect to achieve ~60K by year end
CY 2000 for single systems w/
continuing annual increases


Future directions are to combine
increasingly capable SMP systems
w/ scalable clustering and partitionable
applications for cost effective solutions



Managing Large Systems



Job Objects



- **Job object**
 - Provides a namable, securable, inheritable, sharable object that controls associated processes
 - Limit possible adverse impacts (leaks)
 - Manage groups of processes as a unit
 - Enforce limits on each process associated with job



Job Object

■ Limits and Constraints

- Working set size
- Time limits [various]
- Maximum number of active processes
- Processor Affinity
- Priority class

■ Security and Access

- ACLs and tokens
- Handle access
- Clipboard access
- System changes
- System or Process

exit



Process Control - Using Job Objects for Mgnt

- Partitioning of large system resources;
 - Department, Application, etc.
- Improved Server Management
- Complement taskmgr & sysmon
- Persistent between boots
- Easy to use from GUI, CLI or WSH
- Control of runaway applications
- Scale apps to the appropriate number of CPUs
- More manageable cluster fail-over
- Applicable to Service Level Agreements



Process Control

The screenshot shows the 'Process Control' console window in Windows 2000. The window title is 'Process Control' and it has a menu bar with 'Console', 'Window', and 'Help'. Below the menu bar is a toolbar with icons for 'Action', 'View', 'Favorites', and navigation. The main area is divided into a 'Tree' view on the left and a list of processes on the right. The 'Tree' view shows a hierarchy: 'Process Control (Local)' > 'Rules' > 'Process Alias Rules' > 'Process Execution Rules' > 'Process Group Execution Rules' > 'Processes' (selected). Under 'Processes', there are sub-items: 'Process Groups' > 'IExplorerJob', 'IISJob', and 'SqlJob'. The list of processes has the following columns: Process Alias, Image Name, PID, Status, Affinity, and Priority. The processes listed are: Dfssvc (Dfssvc.exe, PID 836, Normal), Explorer (Explorer.exe, PID 1056, Normal), iexplore (iexplore.exe, PID 532, Normal), Iexplorer (Iexplorer.exe, PID 532, Manag...), Iis (Iis.exe, PID 532, Manag...), inetinfo (inetinfo.exe, PID 876, Normal), llssrv (llssrv.exe, PID 672, Normal), lsass (lsass.exe, PID 244, High), mmc (mmc.exe, PID 864, Normal), msdtc (msdtc.exe, PID 520, Normal), mspaint (mspaint.exe, PID 328, Normal), MSTask (MSTask.exe, PID 784, Normal), ProcConMDB (ProcConMDB.exe, PID 792, Normal), proconsv (proconsv.exe, PID 728, Normal), regsv (regsv.exe, PID 772, Normal), services (services.exe, PID 232, Normal), smss (smss.exe, PID 156, Normal), spoolsv (spoolsv.exe, PID 460, Normal), SQLServer (SQLServer.exe, PID 460, Manag...), svchost (svchost.exe, PID 428, Normal), svchost (svchost.exe, PID 652, Normal), and System (System.exe, PID 8, Normal).

Process Alias	Image Name	PID	Status	Affinity	Priority
Dfssvc	Dfssvc.exe	836		0x1	Normal
Explorer	Explorer.exe	1056		0x1	Normal
iexplore	iexplore.exe	532		0x1	Normal
Iexplorer	Iexplorer.exe	532	Manag...		
Iis	Iis.exe	532	Manag...		
inetinfo	inetinfo.exe	876		0x1	Normal
llssrv	llssrv.exe	672		0x1	Normal
lsass	lsass.exe	244		0x1	High
mmc	mmc.exe	864		0x1	Normal
msdtc	msdtc.exe	520		0x1	Normal
mspaint	mspaint.exe	328		0x1	Normal
MSTask	MSTask.exe	784		0x1	Normal
ProcConMDB	ProcConMDB.exe	792		0x1	Normal
proconsv	proconsv.exe	728		0x1	Normal
regsv	regsv.exe	772		0x1	Normal
services	services.exe	232		0x1	Normal
smss	smss.exe	156		0x1	Normal
spoolsv	spoolsv.exe	460		0x1	Normal
SQLServer	SQLServer.exe	460	Manag...		
svchost	svchost.exe	428		0x1	Normal
svchost	svchost.exe	652		0x1	Normal
System	System.exe	8		0x1	Normal

System Area Networking

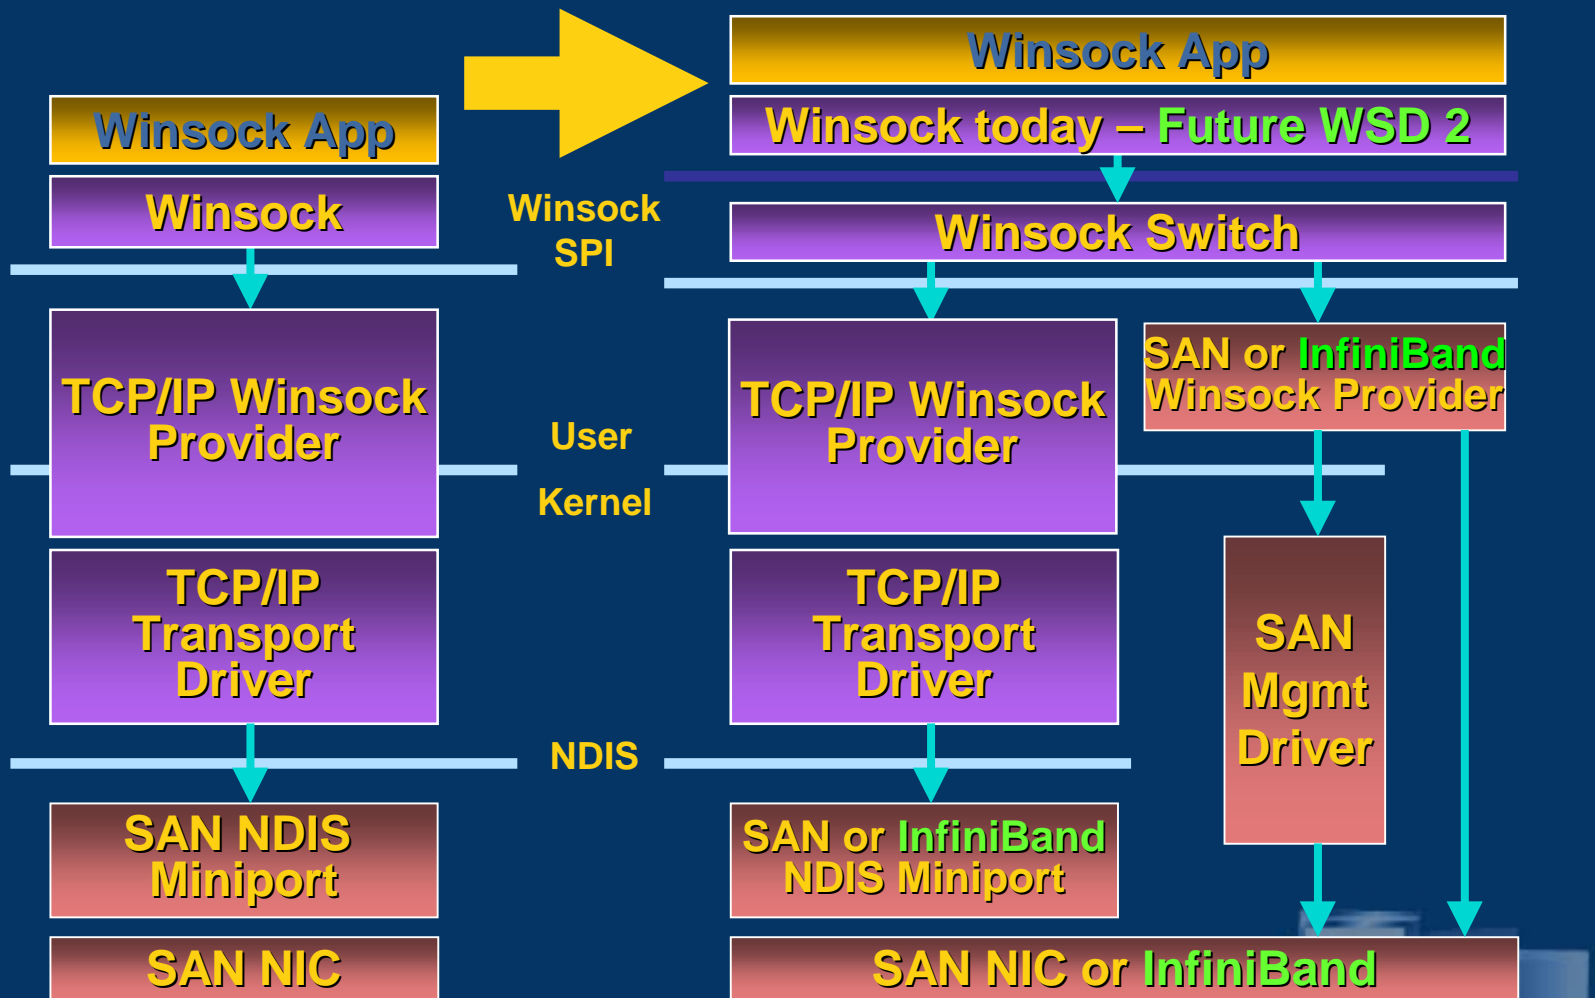


Windows 2000 Data Center Winsock Direct

- Microsoft's approach to Systems Area Networking
 - Windows Sockets Direct Path (WSD)
- Benefits
 - Application compatibility is ~100%
 - High performance w/o app rewrite
 - Agnostic to technology – perf gains from newer HW are possible w/o App churn




WSD Architecture




Early Results

WSD vs Alteon (Synchronous)




WSD vs Alteon (Asynchronous)




Throughput for SAN levels off at 430 Mbps because the system ran out of resources, only 64 MB on system

Percentage Improvement of WSD over TCP/IP (Asynchronous)



Another way of visualizing that the cycles per byte cost for Ethernet approaches WSD since the cost decreases for large messages.

However, 80% of all traffic in a SAN are small messages & data xfers




Scaling 'Out'

- **Distributed and Parallel Applications**
 - **DB/2, Oracle and SQL 2000 are parallel today**
 - **SAP, Baan, PeopleSoft, JD Edwards**
 - **IIS 5 Server farms**
 - **COM+ Load Balancing Server farms**



NLB - Microsoft Dogfood



Note: this is only the Virginia site, not all mirrors

Available Download Bandwidth
Download bandwidth utilized (Peak)
Number of download servers
Number of WWW servers
Peak Bandwidth from Microsoft
Page Views on Microsoft.com


IE4.0 Launch
400Mb/Sec
400Mb/Sec
16
18
300Mb/Sec
32Million

IE5.0 Launch
2.5Gb/Sec
1.2Gb/Sec
40
30
1.03Gb/Sec
60.1Million



SMP and Clusters

The Final Scalability Answer



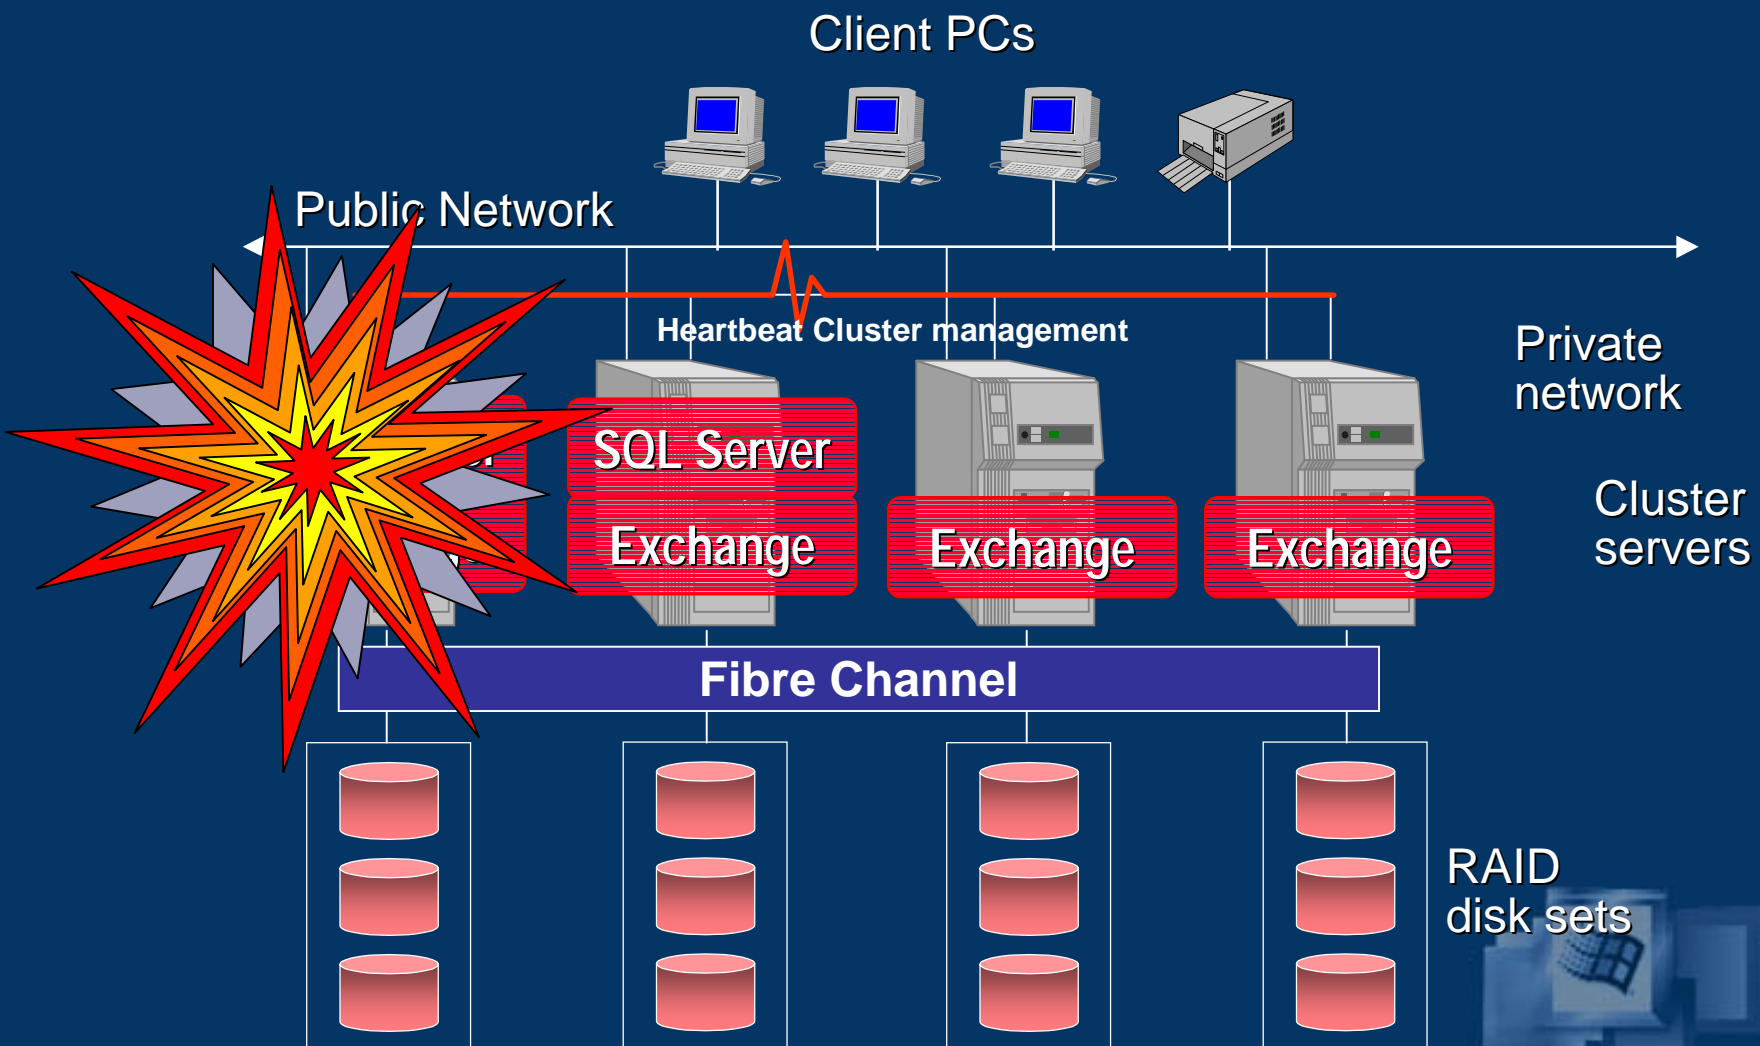
TPC-C Performance: Feb 2000

Rank	Operating System	Database	Hardware System Model	tpmC	Price/ tpmC	Total Price
1	Microsoft Windows 2000	SQL Server 2000	8-way Profusion (12 nodes)	227,079	\$19.12	\$4.3M
2	Microsoft Windows 2000	SQL Server 2000	8-way Profusion (8 nodes)	152,207	\$18.93	\$2.9M
3	IBM AIX	Oracle 8i	IBM RS 6000 Enterprise Srv S80	135,815	\$52.70	\$7.2M
4	IBM AIX	Oracle 8i	Bull Escala EPC 2400 c/s	135,815	\$54.94	\$7.5M
5	Sun Solaris	Oracle 8i	Sun Enterprise 6500 Cluster	135,461	\$97.10	\$13.2M
6	Sun Solaris	Oracle 8i	Sun Enterprise 10000 S70	115,395	\$105.63	\$12.2M
7	IBM AIX	Oracle 8.0	IBM RS 6000 Enterprise Server	110,434	\$122.44	\$13.5M
8	Digital UNIX	Oracle 8.0	Compaq AlphaServer 8400	102,541	\$139.49	\$14.3M
9	HP HP-UX	Sybase ASE 12	HP 9000 Model V2500	102,023	\$63.21	\$6.4M
10	Microsoft Windows NT	Oracle 8i	Intel SHV based PDC/02000	101,657	\$35.68	\$3.6M

Source: Transaction Processing Performance Council

4-Node Back-end Server Cluster

Windows 2000 Datacenter Server



RAID disk sets

4-Node Server Clusters

- **Up to 4 nodes in cluster**
 - Lower cost of excess capacity
 - Increased management flexibility
- **Multiple “possible” and “preferred” owners**
 - Control failover of mixed workloads
 - Easily re-assign server roles in failover scenarios
 - “Cascading” failover to multiple alternates
- **New “Server Group” certification rules**
 - Lower cost & effort to certify configurations
 - More choices in mixing servers within clusters



Multi-Node Cluster Certification

	2-node certification	Multi-node certification
Size & OS certification of server	<ul style="list-style-type: none">■ Advanced Server<ul style="list-style-type: none">■ ≤8P & 8GB■ If ≤2P & 2GB eligible for “low cost clustering”■ Datacenter Server<ul style="list-style-type: none">■ 8P capable or larger	Datacenter Server only. 8P capable or larger.
Storage connection	<ul style="list-style-type: none">■ Advanced Server<ul style="list-style-type: none">■ SCSI or Fibre■ Datacenter Server<ul style="list-style-type: none">■ Fibre only	Fibre only
Stress test time	4 days	7 days
Server “mix and match”	No	Yes, “Server Group” rules: <ul style="list-style-type: none">■ Each server in a 4-node homogeneous configuration■ Each server in a 4-node mixed configuration with every other server in group

Cluster Service

- **Simplified Setup**
- **Rolling Upgrade Support**
- **4-node support in Datacenter**
 - **Complete 4x4 'package' from OEMs only**
- **Windows 2000 Infrastructure Support**
 - **Active Directory - Class Store for Resource DLL's**
 - **Plug and Play Support - Network and Disk**
 - **Better Network and SMB Services**
 - **DHCP, WINS, DFS, and IIS V5 Support**
 - **MMC Integration**
 - **CluAdmin invoked from MMC Snapin**
 - **OCM Support**
 - **Cluster installation during Windows 2000 Setup**
 - **Configuration Completed from Configure Your Server**



Cluster Service

- **COM Interfaces**
 - **MSCLUS COM Automation Interface to CLUSAPI**
- **Client Network Recovery**
 - **NIC Failure detection**
- **UI Updates**
 - **Wizard for Application Virtual Server Setup**
 - **Network simplification**
- **Unattended Setup**
- **Backup API's for Hive and Checkpoint Files**




Infrastructure Support

- File Share
 - Print Spooler
 - Generic Application
 - Generic Service
 - DFS 
 - IIS
 - NNTP 
 - SMTP 
- MSMQ
 - MSDTC
 - DHCP 
 - WINS 
 - SQL Server
 - Exchange

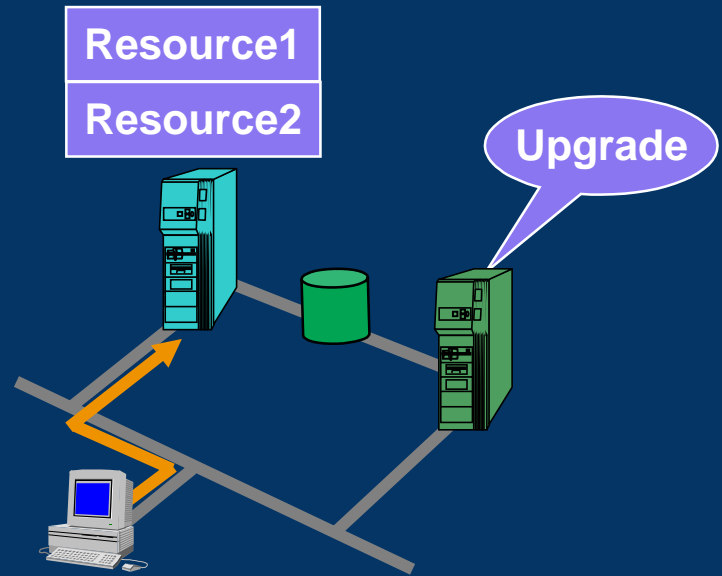


Rolling Upgrades


1



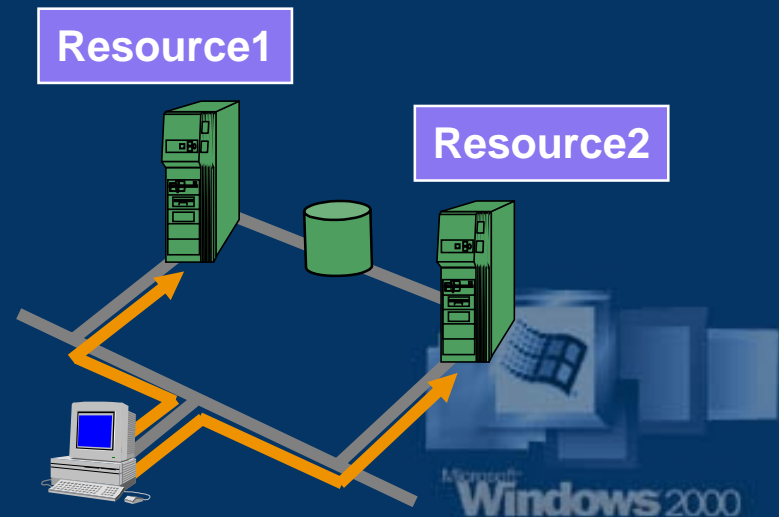
2



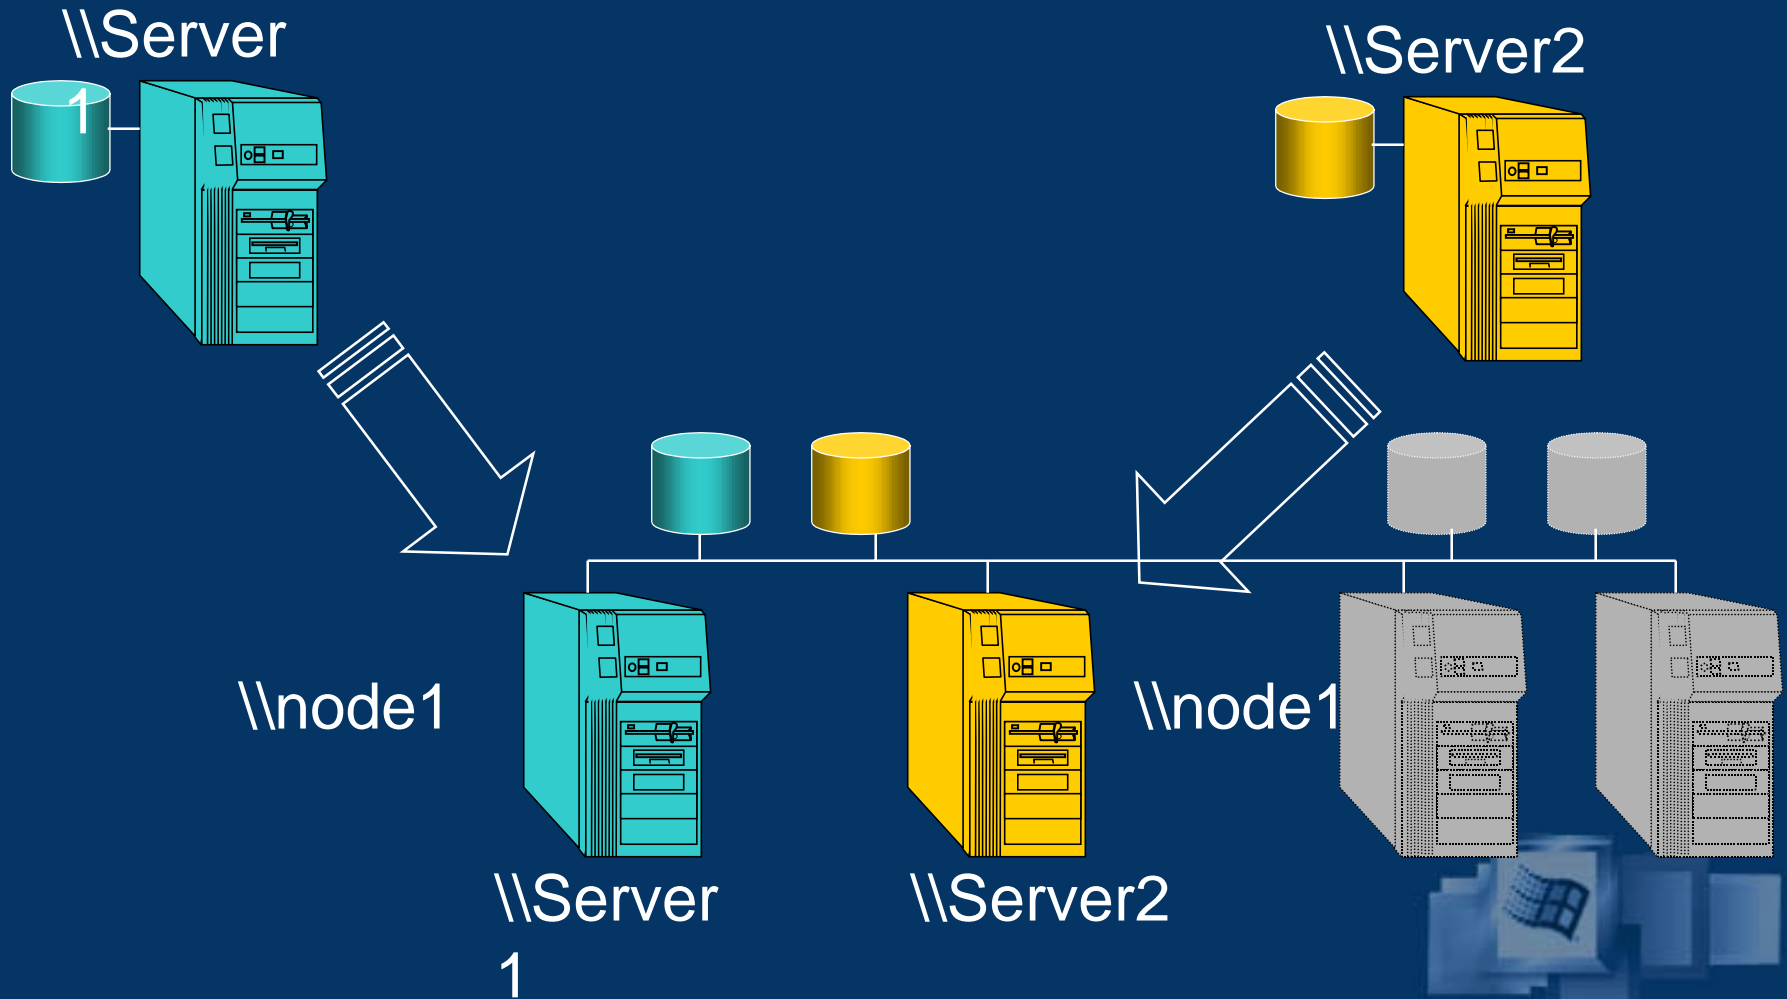
3



4



Server Consolidation

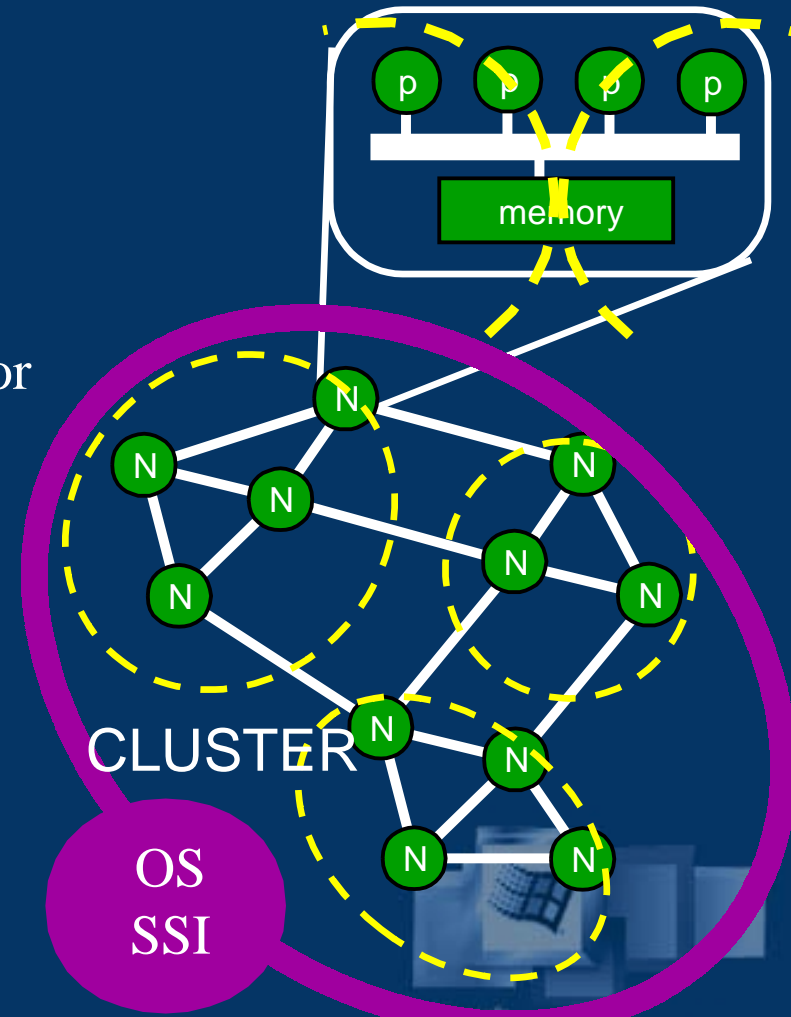


Tomorrow's Compute Complex

■ Y2000+ Midrange/High-End System Structure

Partitioned SMP/Cluster

- Nodes of M-way SMP
 - $N * M >$ anybody needs/wants
 - size of N & M will vary with vendor
- Cluster partitioning of N nodes
 - Availability, workload boundaries
 - Interconnect hardware for memory and application communications
- Partition within large-N nodes
- Dynamic partition adjustments w/o reboot
- Single-System management



IA 64 Market Opportunity

- **With PCs Average Selling Prices at ~\$1K, IA64 not relevant to very high volume**
 - PC Market ~140 Million units growing at 15-16%
 - Probably years before \$1K PCs can be built on 64-bit
- **For the next 3 years relevant to “Workstation” and “>\$10K Server” markets**
 - Total available market for workstations: ~3 Million units growing at 12-15%
 - Total available market for servers >\$10K: ~2.5 Million units growing at 12-15%
- **Memory will be driving factor for system cost (today 4GB=\$14-18k)**



Benefits of Windows on IA-64

- **64-Bit Chip and API allow**
 - Greater than 4GB of physical memory
 - 64-bit arithmetic in hardware
- **64-Bit filesystem allows**
 - File sizes greater than 2 gigabytes



IA-64 Server Usage Scenarios

- **Large Databases**
 - SQL Server, other databases
 - Exchange & other messaging servers
- **Demanding LOB & .com**
 - ERP
 - E-Commerce
 - Encryption
- **In-Memory Data Stores**
- **Complex Consolidation**
 - Heterogeneous workloads
 - Terminal Services



Win64 Project Update

- Targeting first release of 64-bit Windows with the availability of Itanium – 2H CY00
- 64-bit Design Review - March '00
 - Approx 200 ISVs, IHVs & OEMs attend.
- Shipped Windows 2000 – Feb '00
 - Much of the 64-bit porting work was done in Windows 2000.
- Microsoft & Intel Announce 64-bit SDK - Feb '00.
- Windows boots on Itanium - August '99
- 64-bit Windows Demo via Merced Simulator - April '99
- First 64-bit Windows DDK, Second SDK - April '98
- First 64-bit Windows SDK - Jan '98
- First 64-bit Windows Design Preview - Jan '98
 - 170+ partners attend.



Datacenter Server Scenarios



Scenarios

- **Large DB**
 - 32P support
 - AWE API & PAE memory
 - Multiple IO buses
 - *SQL 2000, IBM DB/2, Oracle
- **Large Groupware/Email**
 - Multiple IO buses
 - PAE memory [NT Cache effect]
 - *Exchange 2000, Notes



Scenarios 2

- **Application Consolidation**
 - **Heterogeneous – lots of apps**
 - **PAE**
 - **32P support**
 - **Process Control**
 - **Lone Wolf**
 - **Push for cluster awareness**
 - **App interactions mngd w/ Process Control**
 - **Homogeneous – multi-instance apps**
 - **PAE**
 - **32P support**
 - **Lone Wolf**
 - **SQL 2000**
 - **Exchange 2000**



Scenarios 3

- **Infrastructure Consolidation**
 - **Classic**
 - File, Print, etc.
 - Multiple IO buses
 - **Network Services**
 - DHCP, WINS, etc.
 - **Directory and Domain Services**
 - DNS, etc. [no cluster support]
 - 32P support, PAE [NT Cache], SCM
 - **Gateways**
 - SFM, GSNW, SNA, SFU, etc.
 - Use SCM



Scenarios 4

- **Distributed and Parallel Applications**
 - SQL 2000, IBM DB/2 Enterprise, Oracle
 - SAP, Baan, PeopleSoft, JD Edwards
 - IIS 5 Server 'farms' to back-end DB
 - AppCenter Server CLB 'farms'
 - WinSock Direct System Area Networking
 - Terminal Server and NLB
 - Scale Out the front-end w/ NLB and TS
 - Availability in back-end data systems w/ clusters
 - Highest perf, lowest latency w/ WSD



Why Datacenter?

- Reliability – Certification, SP1 inclusion, etc.
- Availability – 4 Node clusters, Lone Wolf
- Scalability – 64GB, 32P, Winsock Direct, etc.
- Manageability – Process Control, etc.
- Service and Support
 - Datacenter Program w/ OEMs




Datacenter Status



Scaling “Up”

- **Scalability**
 - Up to 32x SMP
 - Up to 64GB memory
 - Kernel and NTFS optimizations
 - Scalable system services for authentication, components and transaction processing
- **Partitioning Support**
 - Dynamic Resource Partitioning
 - Static system partitioning [feature is hardware platform dependent]
- **Investment Protection**
 - Source-code ready for WIN64



Scale-Up Scenarios

- **Large databases for:**
 - ERP
 - Customer relationship management
 - Business intelligence (e.g. data mining)
 - E-commerce
- **Legacy application servers**
 - High number of concurrent transactions
 - Complex business logic
 - Complex / large in-memory databases
- **Server consolidation**
 - Reduce the number of servers
 - Physical consolidation of n-tier / application servers on fewer machines
 - Simplify management and operations




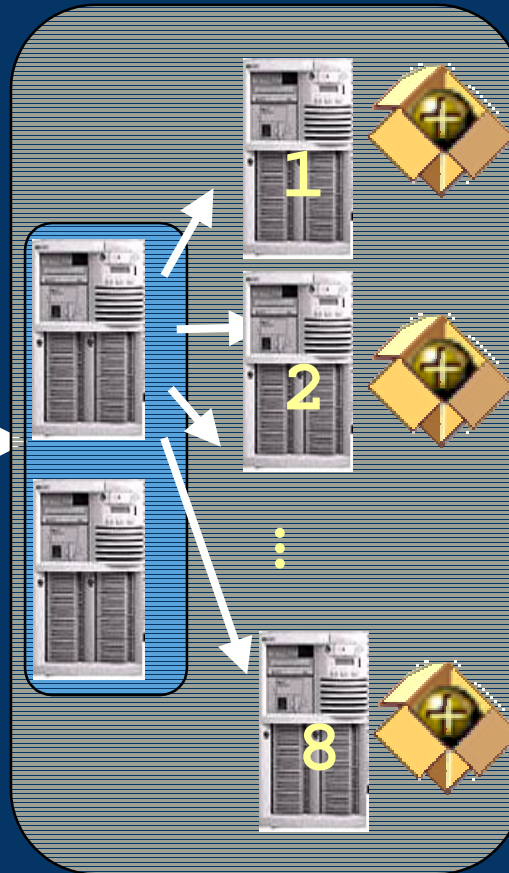
Scaling "out"

Clients

Network Load Balancing

Component Load Balancing (COM+)

Cluster & Storage Service



IS Web Server
or other IP based services

COM+ Components
Application Servers

Data Servers
SQL, Exchange, File


Server-based Computing with



- **Savings**
 - Ease of support and administration - Save 20-40%
 - Control maintenance and License fees
- **Manageability**
 - Centralized administration
 - Convenient software upgrade
- **Speed**
 - Rapid application deployment and upgrades
 - Rapid hardware deployment
- **Scalability**
 - Modular infrastructure
 - User “roaming” possible
 - Consistent bandwidth requirements
 - Flexible performance
- **Simplicity**
 - Broad application and data access from a single client
 - Easy new technology deployment
 - Complexity on the IT manager, not on the users
- **Security**
 - Data security
 - Access control
 - Virus proof



Windows 2000 &
Citrix Metaframe
Connectivity Farm



File Edit Processing Maintenance Reports Options Windows Macros Help

Company Code: Swico Voucher No. 1000000000 [Enter] [F5] [F6] [F7] [F8] [F9] [F10] [F11] [F12] [Print] [F4] [F3] [F2] [F1] [Home] [End] [Page Up] [Page Down] [Find] [Help]

Demo Consulting Company Invoice No. 94529

Vendor: ACR001 Resub To: [Tax] [Payment] [Address] [Aging]

Allen Crafton Engineering

Currency Code: USD Posting Code: ADMIN Voucher Type: STANDARD Cycle Code: INTAX Tax Code: [Tax] [Payment] [Address] [Aging]

Approval Code: Purchase Order No. Order No. Ticket No.

Print Method: Terms Code: Discount Date: Due Date: Class Code: Branch Code: Classification

CHECK NET 30 07/05/1998 07/05/1998 SPECIAL BLSMRY STOR0

Invoice Date: Apply Date: Aging Date: Shipping Tax Request Date: Receipt Date

06/05/1998 06/30/1998 06/05/1998 06/30/1998 06/05/1998


Voucher Comment: 1099 Amounts reported annually on 1099 Internal Comment:

Reason: Apply To No. [One Voucher per C]

Location	Item Code	Received	U/M	Unit Price	Extended Price
Description	Expense Account	Ordered	Tax Code	Tax	Discount
Company	Expense Account	Reference	1099 Coc	Misc.	Freight
		1.		3,696.00	\$3,696.00
design work R. Boogs		1.	NOTAX	\$ 0.00	\$ 0.00
Swico 6030 CH Mto 100			CONTRA	\$ 0.00	\$ 0.00
		1.		5,263.00	\$5,263.00
install cable on Boston job		1.	NOTAX	\$ 0.00	\$ 0.00
Swico 6030 00 000 300			CONTRA	\$ 0.00	\$ 0.00

Enter the voucher control number 01/31/2000 NUM

CLBS App
Servers



Storage



(This CE Device is accessing an ERP application over a wireless connection)



Server-Based Computing

The application on demand

	Light Users	Typical Users	Heavy Users
--	-------------	---------------	-------------




HP NetServer: LPr 2*450MHz, 768MB RAM	110	65	32
------------------------------------------	-----	----	----

HP NetServer: LH3 2* 450MHz, 768MB RAM	110	65	32
-------------------------------------------	-----	----	----

HP NetServer: LXr8500 4*Xeon 400MHz, 1GB RAM	200	120	60
-------------------------------------------------	-----	-----	----

Outsourced SBC
ASP model
+ 90% / year till 2003

ASP



Internal SBC
Hosting model
+ 60% / year till 2003




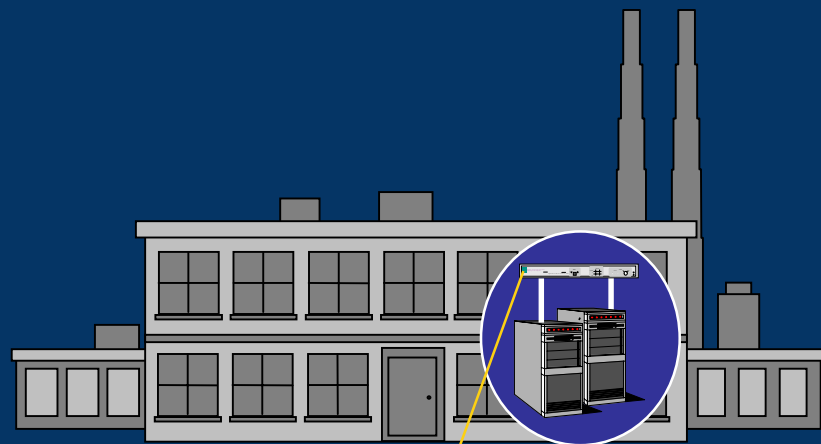
Scale-Out Scenarios

- **ASP (Application Hosting)**
 - **Web Servers**
 - **Enterprise Apps**
- **.COM**
 - **Front end and Mid-tier structures**
 - **Back-end for partitionable data stores**
- **Enterprise Infrastructure**
 - **Exchange 2000**
 - **Replicated Front Ends and Partitioned Data Stores**
 - **File and Print**




SplitSite™ Disaster Tolerance

- *Disaster Tolerance*
 - *Local - 20 meters*
 - *Campus - 2.0 Km*
 - *GeoCluster – 20 km*
- *Real-Time “Hot-Site”*
- *True User Application Session Transparency*
- *Full Data Replication*
- *System Failover if required*



Disaster Recovery : Backup

Theoretical Maximum Transfer Rates
LAN 10 Mbit/s -- 0.8 MByte/sec or 2.9 GB/hour
LAN 100 Mbit/s -- 8.0 MByte/sec or 29 GB/hour



Dedicated Backup LAN -- 100 Mbit/s

Single DLT 7000 (w/ 1.5x compression)
7.5 MB/s or 27 GB/hour



Disaster Recovery : Backup

Shared Disk, Direct to Tape



**Direct
Disk to Tape**
One-way data flow
from disk to tape

Configuration

- Disk space is managed as a conglomerate for multiple servers
- Sequential backups are completed from Fibre disk to tape without server intervention.

Value - Needs

- Backup window: The server is never down for backup.
- Backup data is essential to resume business if data is lost due to hardware or application SW error
- Easy to scale up by adding more more disk or libraries.






Manageability : Why ?

Getting Operations Wrong Bites




Development

Where's my project gone?



Infrastructure

Where's my money gone?



Operations

Where's my business gone?

“Once the CEO sees the corporate dependence on distributed systems, he or she will issue a mandate to the CIO -- Keep these systems running or.... “

Forrester Group, Computing Strategies, February 1998



Microsoft Operations Framework

Achieving Excellence in Service Management

- **Vision:** Create comprehensive operations guidance for achieving mission-critical production system reliability, availability and manageability on the Microsoft platform
- **Goals:**
 - Provide rich intellectual property (IP) to help the field, partners and customers address *process, organization, and technical operations* on our platform
 - Reduce the *time and complexity* involved in running MS products in an enterprise operations center
 - Provide guidance for the interoperation of the MS platform with other technology platforms
 - Extend the *reach / accessibility* of the operations knowledge developed
 - Deliver Operations Kits eg for DataCentre Operation, and Characterized Environments eg hosted IT configurations for ASP's
 - Drive Microsoft and Partner Service Offerings




MOF Service Mgt Functions

Capacity Management
Availability Management
Cost Management
Workforce Mgt
Contingency Planning

Release Readiness

Change Management
Configuration Management
Release Management

SLA Review



Implementation Review

Help Desk
Problem Management
Failover & Recovery




Operations Review

System Administration
Monitor/Measurement
Directory Services Administration
Security Administration
Network Administration



VantagePoint for Windows Highlights

- Sophisticated Windows, Solaris and HP-UX management from Windows
- Most advanced event browser and performance subsystem Windows has ever seen
- Most advanced object technology agent Windows has ever seen
- Best-in class reporting exposed via IIS
- Extensive out of the box management for: Windows core OS, SQL Server, SNA, MSMQ, DNS, WTS, RAS, ...
- Advanced SPIs for Exchange, Internet Servers, SAP, Oracle
- Leverage legacy OpenView knowledge
- 100% developed on



VantagePoint Provides:



- **Business-driven Intelligence**

(See business impact immediately)
Business transaction insight
Context-driven
Business-centric views



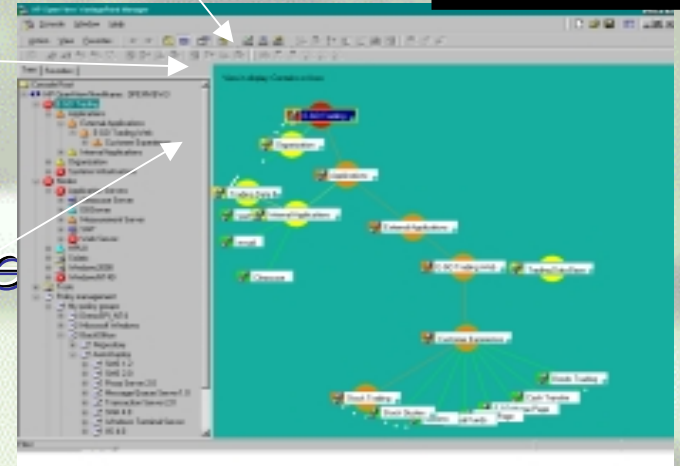
- **Instant Intelligence**

(Fast Results)
Service discovery
Policy deployment
Application knowledge
Business/IT knowledge captured



- **Active Intelligence**

(Respond to the unexpected)
Dynamic measurement
Guided root-cause analysis
Automatic correction



MOF Example Implementation

Exchange Management Reports

- Reporting enables IT administrators and experts to very rapidly access up to date information wherever they are and whenever they want.
- Automated generation and web publication of these decreases TCO considerations

[Exchange SPI Mailbox Usage Trends](#)

Click on the hyperlink
for example reports

- Reporting is possible on all standard Exchange performance counters AND results of specific calculations performed by the SPI (such as round trip time)
- A powerful means to build SLA management solution for the Exchange backbone

[Exchange SPI Messaging Availability](#)

Click on the hyperlink
for example reports




**HP services help you
unleash the power of your
investment in Microsoft®
technologies.**



A Full Range of Service Solutions

Consulting



Plan for the optimum solution

- IT Infrastructure Assessment and Design
- Internet Architecture Planning
- E-services Architecture
- MS Exchange and Knowledge Management

Integration Center



Ensure a smooth fit into your existing environment

- Hardware, Software and Network Integration, Installation and Configuration
- QuickStart Installation services
- Relocation Services

Operations & Management



Ensure trouble-free transition

- Microsoft Software Licensing services
- Network Management
- Desktop Management
- Business Recovery

Education

Develop Windows NT/2000 expertise


- MS Official Curriculum for Windows NT/2000 & BackOffice
- UNIX-Windows NT Integration
- Customized end user training



Financing & Leasing

Multivendor financing from a single source


- Project planning and deployment
- Funding with flexible terms
- Leasing
- Asset Management



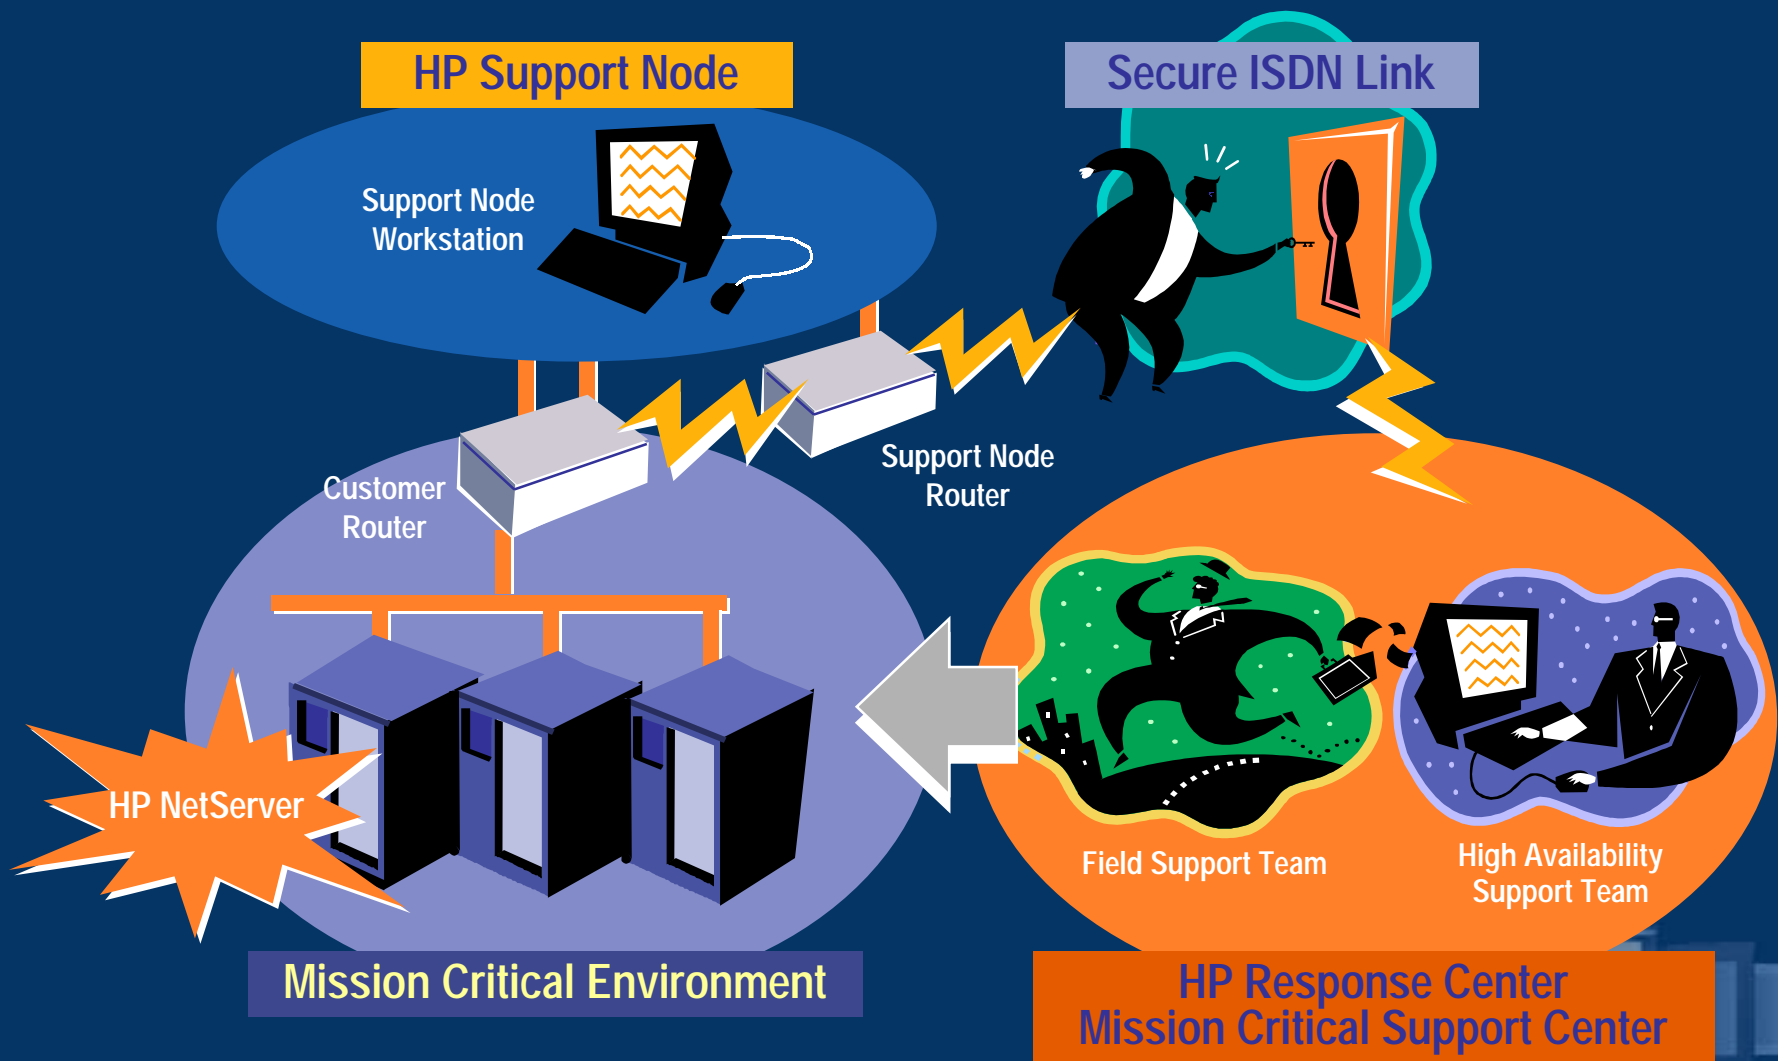
Support

Keep it running & fix it fast

- Mission critical services and support
- Remote and onsite technical assistance
- Multivendor hardware, software, and network support



HP High Availability Observatory



Windows 2000 Data Center Solutions Alignment

Built from HP's Leadership!

- ✓ Installation and startup services for software and hardware
- ✓ Availability Assessment
- ✓ 24x7 Hardware & Software Support
- ✓ One-stop shop, no finger pointing
- ✓ On-site parts and services - 4 hour response time commitment
- ✓ Direct Access to CPR/QFE
- ✓ Ability to deliver 99.9% uptime solutions
- ✓ Robust Solution Testing and Certification Programs



Data Center Solutions

Focus

Windows 2000 Datacenter Server

The operating system for the most demanding levels of availability and scale

- **Support MS Scale up and Scale Out Strategies in n-tier architectures**
- **Drive availability with 4-node fail-over clustering, advanced clustering & Fault Tolerant Solutions**
- **Industry leading customer value in complete Server and Storage Consolidation Solutions**
- **“Killer-App” solution sets (e.g. Database, Exch 2K...)**
- **Industry leading Service and Support**



What is HP Doing Today with Windows 2000 DataCenter

- OEM Leadership in working with MS on DataCenter Program
- Developing and testing DataCenter configurations and solutions
- Corporate ISV partner management
- Leading support and development of > 4 way 4 node clusters with Microsoft
- Finalizing service and support offerings for DataCenter Solutions
- MCSC Problem Management with JointQ
- DataCentre HCL Configuration Certification



MCSC for Datacenter

- MCSC = Microsoft Certified Support Center
- New service program but based on existing MCSC program
- As opposed to MCSC only covers the Windows 2000 Datacenter Server OS
- Includes a joint support queue staffed by Authorized OEMs and Microsoft
- Consistent core set of services across Authorized Datacenter OEMs
- Lots of room for service differentiation




Joint Support Queue

- Queue owns Datacenter OS & HW problem all-up
 - Problem isolation and reproduction
 - Driving problem resolution into either Microsoft or OEM engineering team
 - Providing relief and testing and delivery of Quick Fix to customer
 - Joint Queue is only mechanism for distribution of QFEs for Datacenter
- Accessed via a support contract sold by OEM
 - Directly to customer
 - Or sold to Microsoft on behalf of customer in Alliance and Alliance HA situations
- OEM or Microsoft escalation process may occur before call is handed to joint queue
- Staffed by Microsoft & Authorized Datacenter OEM personnel



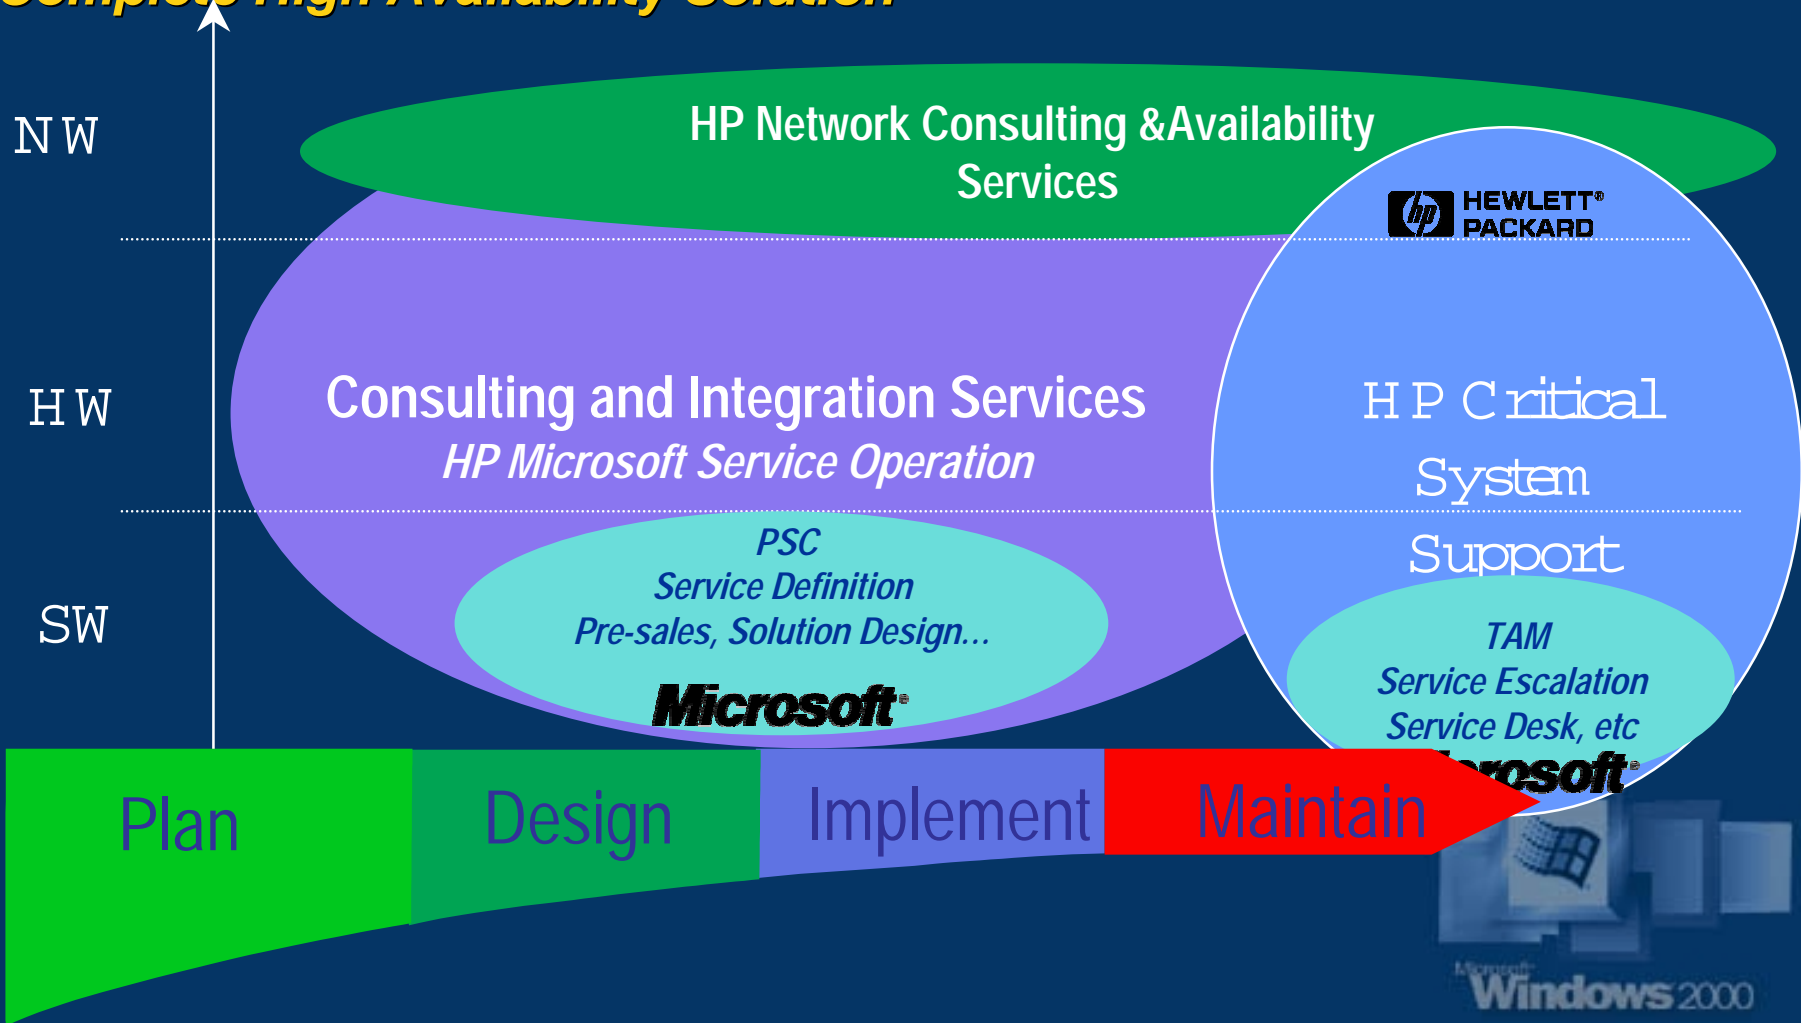
Datacenter HCL

- Stringent system validation – New client/server test
- 2-week initial configuration test
- 1-week re-test for “minor” configuration changes
 - Firmware
 - New adapters
 - Service Packs
- All “kernel” mode components must be loaded
- Validation driven by 8 “load generators systems” stressing the server under test
- Validation suite creates a system configuration file documenting the tested configuration
- Configuration tool enables customers to check if their system is still validated
- Signed system configuration files downloadable from Datacenter HCL site



What Microsoft and HP can offer together

Complete High Availability Solution



HP Services and Support Network

- Over 27,000 employees
- 250 HP Technical Education Centers

- 9 HP Operations Service Centers
- 35 HP Response Center network locations in 34 countries



- Over 600 Support Offices located in more than 120 countries
- 3 Support Material Distribution Centers

- 7 Microsoft Competency Centers
- >40 years support experience



- **Joint HP and Microsoft initiatives**
- **7 centers around the world, connected by a worldwide knowledge management system**
- **Provide benchmarking and proof-of-concept services for Windows 2000 based Infrastructures**
- **Focused on:**
 - **Windows 2000 and BackOffice architecture design, integration, management and deployment**
 - **Microsoft Exchange, Microsoft Commerce Internet System (MCIS), Windows DNA architecture, and high availability solutions**





HEWLETT®
PACKARD

A *We take you there...* ALLIANCE E

Microsoft®

