



IBM Netfinity X-architecture

Companies demand cutting-edge technologies and extreme performance... at an affordable price. IBM is committed to bringing down the cost of enterprise-class technologies by leveraging its vast portfolio into the industry-standard server market.

A technology vision paper for IT professionals worldwide.

Executive Summary

In an industry transformed by rapid advances, IBM has responded with a new set of initiatives to help you realize more value from your network environment and be better prepared for the challenges of constant change. This fundamental shift has taken the requirement for performance and availability to new levels. **In response, IBM has harnessed decades of expertise and experience that helped build our enterprise systems, and applied that knowledge toward building its new generation of IBM Netfinity® servers.**

IBM Netfinity servers offer the tools and solutions to let your organization realize greater competitive value from its IT investments, while it lets you control your business computing environment more precisely, with less effort than ever before. This reflects IBM's commitment to what we call "the new enterprise"—businesses that must be lean and nimble, at work around the clock, and whose systems must respond in kind—keeping you one step ahead on the technology curve.

IBM Netfinity systems extend today's industry-standard platform beyond the server to include advances in technologies you need to implement business-critical core business, e-business and deep computing applications. **Through advances in migration and integration technologies, reliability, scalability, availability, manageability, security and serviceability, IBM Netfinity X-architecture® helps you realize extreme performance, maintain your existing system investments, and provide an expanded definition of the IT platform of the future.**

In making IBM Netfinity X-architecture the most relevant for your business needs, we designed it to:

- Leverage IBM's vast technology portfolio and server expertise into industry-standard computing
- Bring down the cost of enterprise-class technologies
- Set the standard for enterprise quality through unmatched testing, systems integration and compatibility

Key elements of the Netfinity X-architecture include powerful processors, reliable and available memory systems, scalable I/O, world-class silicon and module technology and advanced caching software. Also included are clustered systems featuring technology derived from IBM's industry-leading S/390® and RS/6000® SP™ product lines, as well as interoperability with existing large and midrange systems. Netfinity Manager™, IBM's system-management middleware for Netfinity servers, lets you focus on running your business rather than maintaining the system. Also, Netfinity servers come with IBM's technology-enabled support and service, including proactive maintenance through advanced warning systems and remote diagnostics.

With IBM Netfinity we promise to deliver **power** to run your business-critical applications, **scalability** to grow your business, **control** to proactively manage the systems on your network and technology-enabled **service**, whenever and wherever you need it.

This paper describes many of the technology innovations IBM is investigating or designing into our line of industry-standard servers. As the world's leading server manufacturer, we are taking the reliability, availability and serviceability features from our large and midrange systems and using them to make Intel processor-based servers ready for your business-critical 24x7x365 applications. Because of the dynamic nature of this industry, changes may occur in IBM's actual implementation. IBM reserves the right to change specifications or other product information without notice.

Designing an Architecture for the Next Century

In designing the IBM Netfinity X-architecture, IBM listened intently to customers. We began by first understanding the trends that are shaping IT system purchases, and then aligning our investments to bring the right combination of technologies together to best meet the demands of your IT needs. We found the most important trends for industry-standard servers are likely to be:

- 1998: Accelerated adoption of industry-standard server solutions based on the popular support and low acquisition cost of Microsoft Windows NT and Intel processor technology
- 1999: Efficient use of IT budgets will drive wide scale server consolidation and superior systems management across multiple platforms.
- 2000: Penalties for system downtime become severe and push requirements for availability toward continuous operations.
- 2001: Companies demand access to a broad array of spectacular applications, development tools and middleware on IA-64 platform.
- 2002: Internet revolutionizes business transactions, information dissemination and drives massive data utilization requirements. (The premise of e-business is fully realized.)
- 2003: IA-64 has widespread commercial appeal.

The result of these trends: Businesses will become increasingly more reliant on their systems and will place ever-greater demands on them. At the same time, businesses will expect the same degree of performance, reliability, scalability and security as provided by traditional high-end systems, as well as systems management and control.

IBM Netfinity X-architecture is designed to address these trends and deliver the most reliable industry-standard server foundation for enterprises large and small.

Your Business Solutions Require Extreme Power and Scalability

Today's computing solutions are synonymous with power. High-function software is fueled by the underlying power of the hardware platform. **Power scales in two dimensions**—vertically (within a system) and horizontally (connecting multiple systems). The key to vertical scalability is a balanced system design incorporating scalable processor, memory, storage and I/O subsystems. Horizontal scalability is enabled by external storage, high-bandwidth I/O, high-speed interconnect technologies and the robust middleware required for connecting multiple systems.

Building servers for e-business starts with designing server nodes that deliver the maximum power embodied in today's microprocessors. **IBM works closely with Intel to keep Netfinity servers on the leading edge of the Intel microprocessor roadmap, while helping Intel to deliver improved performance and reliability.** During development of the Pentium® II Xeon™ chipset, for example, IBM provided assistance to Intel in verification of the chip design. IBM's relationship with Intel will continue through Merced and follow-on IA-64 designs.

As Intel introduces its 64-bit Merced (IA-64) microprocessors, IBM will incorporate them into 4-way and 8-way symmetric multiprocessor (SMP) systems. **To help achieve industry-leading performance and other features central to Netfinity servers, a team of designers with IBM RS/6000 and AS/400® experience is developing the chipset for our IA-64 8-way SMP.** This team has a proven track record on chipsets designed for robust server performance, and they'll use this expertise to incorporate features previously available only in large enterprise-class systems. IBM is committed to delivering Netfinity IA-64 systems with advances in computing power required for a robust architecture.

IBM is also working closely with key application vendors and operating-system providers to help ensure the broadest range of software support for IA-64 architecture. **In June 1998 IBM and Intel announced IBM Applications Solutions Centers (ASCs) to support ISV development, performance tuning and porting of applications to the IBM Netfinity IA-32 offering.** Three ASCs have opened worldwide so far, and more are planned. The ASCs will support IA-64 based products in the future. IBM will leverage our middleware software portfolio for the IA-64 architecture with products such as DB2®, MQSeries® of transaction systems software and VisualAge® application development products.

Even as large investments in money and manpower are being made in IA-64 system development, IBM continues to deliver a full range of Netfinity systems based on the latest and best microprocessors and chipsets from Intel's proven IA-32 architecture. **Further extending the power capabilities of IA-32 network business systems, the IBM Netfinity 7000 M10 4-way SMP server has achieved industry-leading performance across a wide spectrum of benchmarks, including TPC-D, SpecWeb and SAP, to name a few.** This system is based on the latest Pentium II Xeon microprocessor and chipset, featuring up to 4 processors with full speed L2 caches and a 100-MHz front-side bus. In 1999 IBM will introduce a Netfinity 8-way IA-32 SMP server with the Intel Profusion chipset—keeping Netfinity servers in the forefront of performance and technology.

IBM's experience in building SMP systems is demonstrated by the performance capabilities of the Netfinity 7000 M10 4-way SMP, and this experience becomes more important as even larger SMPs become available. **Building an 8-way computer is not**

the difficult challenge. Designing one that scales and delivers useful performance benefits from each of the additional processors requires the talents of an experienced design team—and this will be the IBM Netfinity differentiator in the 8-way SMP arena. In addition to staying on the leading edge with hardware, IBM recognizes opportunities to increase performance using software accelerators. Using advanced caching software developed by IBM Research, **IBM Netfinity Web Server Accelerator**, dramatically enhances Web serving performance. Web Server Accelerator caches static file content on behalf of a Web server. Requests for this content are intercepted at the lowest levels of the protocol stack, avoiding overhead and reducing access times. **In a recent performance test, the Netfinity 7000 M10 used the Web Server Accelerator to provide Web serving performance more than 60% better than competitive Intel processor-based servers.** For more detailed information, see our Web site at www.ibm.com/netfinity.

The history of SMP systems is one of continual growth. Software continues scaling upward to take advantage of added processing elements. IBM will continue to deliver on our promise of scalable power in Netfinity servers as software and hardware elements continue to be enhanced.

Designing Memory Systems to Give Data-Intensive Business Applications Expanded Performance

Although microprocessor power is critical to server performance, it is not the only important parameter. Main memory must expand to complement the power of the processors. Thanks to recent advances in network operating system technology, IBM Netfinity servers now have the ability to scale up to 8 GB of memory. Memory growth is expected to continue as Intel expands its architecture with 64-bit addressing support. Support for memories of 64 GB and larger will be provided on Merced systems, with terabyte (TB) capacity envisioned within the next decade. These increases can significantly improve performance, allow support for a larger number of users and improve scaling characteristics for data-intensive e-business applications. IBM Netfinity will continue to work with **Microsoft, SCO, Novell, Oracle, IBM Software, Lotus, SAP, Tivoli** and other operating-system and application vendors to help your applications scale to their full potential.

Today's business-critical information processing systems require memory subsystems that not only provide the capacity for business applications, but also enterprise-class Reliability, Availability, and Serviceability (RAS). **IBM is exploiting a variety of memory RAS techniques learned from decades of large-system leadership for Netfinity servers.** In addition to Error Checking and Correction (ECC) and soft-error scrubbing, IBM will provide such enhanced functions as redundant memory-bit steering, chip-kill correction, memory mirroring and hot-plugability.

IBM recently announced a new error-correcting memory technology that enhances industry-standard server reliability. Many of today's industry-standard servers provide single-bit error correction, but lack the ability to correct lost data if a DRAM chip fails. IBM has incorporated chip-kill correction, which uses enhanced multi-bit error detection and correction, into memory DIMMs to transparently correct lost data should an entire DRAM chip fail. Providing mainframe-class memory, reliability can be as simple as plugging in one of these new DIMMs.

For memory failures that require physical replacement of memory components, the Netfinity light-path fault-locating system will clearly and quickly identify the failing

component. This will help you perform maintenance with minimum downtime and reduce cost by identifying the correct failing component the first time.

Scalable I/O for Growth and Extreme Performance

Scalable systems require scalable I/O for both vertical growth of a single server and horizontal growth of a cluster of servers. Today, IBM Netfinity servers support scalable I/O using intelligent I/O processors in our Netfinity Fibre Channel RAID Controller Unit, ServeRAID-3 Ultra2 SCSI adapters and Serial Storage Architecture (SSA) PCI RAID adapters. I/O processors are designed to offload the server's main processors by performing I/O-intensive tasks with local processors. All three IBM RAID adapters are driven by high-performance RISC processors. IBM will continue to improve the throughput of RAID adapters and will adopt emerging industry standards by incorporating I₂O[®] device driver support, faster I/O processors, and higher performance versions of PCI.

For solutions requiring extremely high bandwidth, cabling distance and high availability, Netfinity Fibre Channel subsystems will provide scalability to hundreds of terabytes of storage, 100-Mbps throughput and distances of up to 10 kilometers (6 miles) between the server and disk storage to provide additional protection of data. Implementations will include SCSI-over-Fibre solutions to leverage existing investments in SCSI disks and subsystems such as the EXP-15, as well as Ultra3 SCSI and all-Fibre solutions in the future.

The Netfinity Fibre Channel subsystems and Netfinity EXP15 storage expansion units can be used to attach up to 36 TB of disk storage to a *single* PCI expansion slot. Using Netfinity Fibre Channel subsystems, which support single server or N-node cluster configurations, over 432 TB can be attached to a single Netfinity 7000 M10 server or cluster.

The ServeRAID-3 family of adapters supports up to 4.3 TB total external storage when used with the Netfinity 5500, 7000 and Netfinity EXP15 storage expansion units. With SSA technology, the Netfinity 7000 can provide up to 1.7 TB of external storage connected over longer distances than SCSI can support.

All three IBM RAID adapters support Microsoft Cluster Server (MSCS) two-node configurations for high availability. The ServeRAID-3 adapters also support two-node Netware clusters.

The newest IBM Netfinity servers are incorporating even higher levels of performance by implementing industry-standard enhancements to the PCI bus. This includes support for 64-bit PCI, which is used by the ServeRAID-3 and Netfinity Fibre Channel subsystems to support data rates up to 264 MBps.

Industry-standard PCI bus technology has become extremely important in the high-availability server market. The current PCI Local Bus Specification includes a 64-bit 66 MHz definition with peak bandwidth capabilities up to 533 MBps. With higher performance I/O technologies such as Gigabit Ethernet and Fibre Channel on the horizon, even higher bandwidth I/O bus capabilities will be needed in the future.

IBM is working to boost current PCI Bus performance with an extension of current technology called *PCI-X*. This new technology runs at speeds up to 133 MHz and would provide the critical I/O bandwidth needed for industry standard servers running enterprise applications like Gigabit Ethernet, Fibre Channel and Cluster Interconnects. **IBM is driving this new industry-standard technology with a large network of other industry leaders including 3Com, Adaptec, HP and Compaq.** Not only will

customers gain much needed PCI bandwidth with **PCI-X**; it can also provide investment protection by offering backward compatibility to existing systems. **PCI-X based products are expected to be available in enterprise-class products the second half of 1999.** The 64-bit **PCI-X** bus provides burst transfer rates above 1 Gigabyte per second, eight times the rate of the most common PCI implementations today. **PCI-X** also introduces an enhanced protocol to increase the efficiency of data transfer and to simplify electrical timing requirements, an important factor at higher clock frequencies. Allowances for future capabilities are also provided. In addition, **PCI-X** accomplishes this while still remaining backward compatible with the current PCI bus.

Performance Depends on Enterprise-class Storage Technologies

Enterprise storage is an industry-standard term for a set of technologies that IBM has pioneered in large-enterprise computing over the last 30 years. Enterprise storage technology addresses the needs of businesses as they increase their information technology investment and consolidate their information storage platforms. Enterprise storage allows information to be assembled in one place where it can be shared among all computing platforms in a business and managed in a secure, efficient manner. If you have traditionally deployed smaller servers across your network, enterprise storage is a way of moving away from the view of storage as disks on individual servers to a consolidated view of data enterprise storage and management.

As the world's leading IT provider, IBM takes an enterprise-centric approach by providing the total-storage solution for heterogeneous computing environments that span operating systems and system architectures. The IBM Seascope Storage Enterprise Architecture family of products integrates rich software, global enterprise storage management, and a full range of disk, tape and optical storage technologies to create highly reliable, scalable and versatile application-based storage solutions. **The Seascope family of solutions spans the range of servers from Intel processor-based Netfinity servers to the largest IBM S/390 enterprise servers.** Seascope is structured with "snap-in" building blocks that enable rapid integration of new technologies while protecting existing storage investments.

A key element of the Seascope family is the IBM ADSTAR® Distributed Storage Manager (ADSM). This software provides powerful enterprise storage management for distributed environments including backup/restore, hierarchical storage management and a Disaster Recovery Manager feature to help you prepare and execute a disaster recovery plan. **ADSM gives you the basis for managing a complex clustering system and a strong solution for LAN and server consolidation.** Geared toward solving storage management problems for the entire enterprise, ADSM supports a wide variety of IBM and non-IBM platforms for both small and large systems.

Netfinity servers are tested to work in defined configurations with other enterprise-storage vendors' products via the IBM Netfinity ServerProven™ program.¹ This allows customers to use Netfinity servers with existing or new enterprise-class storage subsystems supplied by many vendors.

The creativity that made IBM Storage Systems Division the pioneer in disk storage and RAID technology is being applied to develop adaptive RAID algorithms that intelligently optimize user throughput under variable workload and configuration conditions. Today, RAID subsystem firmware allows RAID-array configuration and, as the workload changes, it automatically adjusts many operating parameters and algorithms to optimize performance.

¹ SP products are warranted and supported by their manufacturers, not IBM.

IBM continues to be the leader in storage and storage-management environments, and Netfinity servers benefit from this experience through advanced technology delivered in our ServeRAID-3 adapters. **Features such as on-line, local and remote management, on-line capacity expansion, migration to RAID, removable battery-backed write caching with memory chip-kill error correction and high-performance processing set the ServeRAID-3 adapter family above the rest.**

The technology that complements the Netfinity enterprise storage strategy is found today in IBM Netfinity servers. The Netfinity 5500 and 7000 M10, for example, feature industry-leading hot-replace and hot-add PCI technology. This allows replacement and addition of I/O adapters, such as ServeRAID™ and network adapters, without bringing the system down. In addition, innovations pioneered by IBM such as predictive failure analysis (PFA) for hard disk drives, fans, power supplies, hard disk and memory enhance IBM Netfinity servers' leadership position in product availability and serviceability.

Cutting Edge Microelectronics Technology Pushes Performance

Netfinity servers not only employ the latest industry-standard Intel microprocessors, but will leverage the best IBM Microelectronics fundamental chip and packaging technologies to provide extra performance.

Copper chip wiring, a technology breakthrough announced by IBM in 1997, enables information to be communicated more quickly throughout the chip. **Copper chip technology will be used to develop future Netfinity servers, enhancing performance of critical system components.**

Another new technology breakthrough by IBM is a commercially viable Silicon On Insulator (SOI) chip. For more than three decades, scientists have been searching for a way to enhance the performance of existing silicon technology. SOI advances chip manufacturing one to two years ahead of conventional bulk silicon. Through reduced capacitance, SOI can be used to improve performance by as much as 25%. **With its superior power and speed characteristics, SOI will enable Netfinity servers to deliver even more performance in the future.**

IBM has long been a leader in electronic packaging, driven by a need for high performance, high reliability and enterprise-class electrical signal integrity in the midrange and mainframe server business. **These packages will be incorporated into IBM Netfinity servers to bring the same performance and reliability benefits to industry-standard computing.**

One recently announced example of high-performance packaging leadership is the new glass ceramic chip carrier, which addresses the three main obstacles to system level integration: chip "real-estate" constraints, thermal and electrical performance and cost. This unique package features the industry's highest off-chip and card-attach densities, with on-package I/O counts of up to 5000, and off-package I/O counts of more than 1600. IBM engineers made this technology accessible by eliminating costly thin films from the glass ceramic process.

Because there is no thermal mismatch between glass ceramic and silicon, die size is not limited, enabling the integration of complex functions onto one chip. IBM's glass ceramic chip carriers will be able to accommodate die sizes of up to 30 millimeters

without chip-to-package reliability concerns, pointing the way to future system-on-a-chip designs.

IBM's glass ceramic chip carriers feature full array flip-chip (C4) die attach, allowing for die shrink and providing improved electrical performance due to shorter interconnect lengths, distributed power and ground. The electrical performance is boosted even higher by the low dielectric of glass ceramic and the presence of copper conductors that can speed the transfer of the signal to the board by up to 28%.

IBM Blue Logic technology combines the performance and cost advantages of custom logic with the time-to-market benefits of designing with thoroughly tested core building blocks. This enables us to use leading-edge technology and provide cost-effective product solutions. Many logic cores are already available, encompassing processors, memory controllers, bus controllers and many other functions supported by the latest application-specific integrated circuit (ASIC) system-level silicon design, test methodology and tools. IBM Netfinity servers will leverage the Blue Logic technologies in our ServeRAID and systems management subsystems.

Clustering Delivers Continuous Performance

The underlying architecture is an important step in determining the characteristics of a system. An SMP system, with a centralized memory, has processor scaling limitations and executes with a single operating system image, therefore creating a single point of failure. Scalability and availability objectives are difficult to meet with a system of this type. Clustering, however, is based on a distributed memory model that improves overall system availability and allows a system to scale to hundreds of processors. IBM cluster solutions focus on three aspects of clustering: availability, scalability, and manageability. IBM will leverage the hardware and software technology of our successful **S/390 Sysplex** and **RS/6000 SP** product lines to deliver these capabilities in our Netfinity cluster solutions.

Continuous Availability

One of the primary motivations for clustering is to achieve continuous availability. There are three steps to accomplishing this:

1. Reduce the probability of a failure
2. Minimize the effects of failures
3. Handle scheduled maintenance and upgrades without downtime

The probability of failure can be reduced by improving the reliability of components and is an ongoing process. This is accomplished through new technology, error correcting codes and other methods.

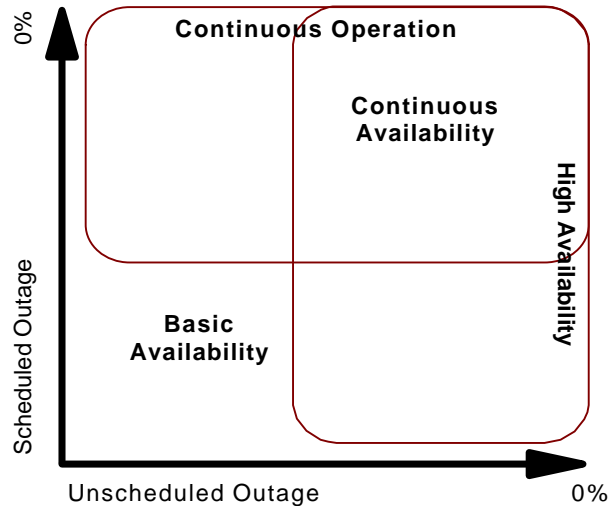
There are numerous approaches to minimizing the effects of failures usually involving the use of redundant components that can take over for failed components until they can be repaired. One approach, used in the RS/6000 SP design, gives particular attention to reduce the probability of total-system failures.

In this approach, the system is designed to be tolerant of single node failures, and the effect of a failing node is isolated to applications using only that node. Independent operating system images on each node make this possible.

IBM Netfinity clusters will allow a node to be powered off and disconnected from the rest of the system while repairs are being made to the node. Later the node can be replaced and powered-up while the system is running. Similarly, hardware and software service and upgrades can be applied to a node or a part of the system without shutting down the entire cluster.

Another approach to minimizing the effects of failures is to anticipate a failure before it actually occurs. This is called Predictive Failure Alerting and is incorporated in Netfinity systems.

The figure to the right illustrates these ideas. *Basic availability* can be achieved by providing a highly reliable system and by servicing as needed. **IBM Netfinity systems are being designed to provide continuous availability by minimizing the effects of a failure or an outage (both scheduled and unscheduled), keeping your business up and running for your end users.**

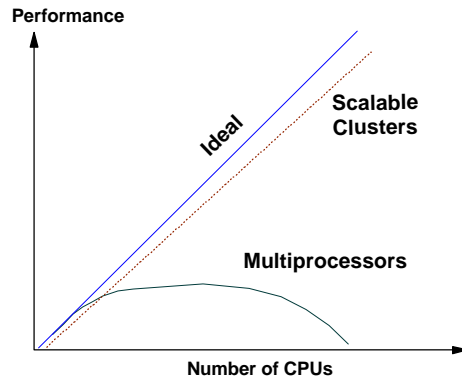


Near-Linear Scalability

While the vertical scalability of an SMP system can satisfy many performance requirements, horizontal scalability through the use of clustering addresses highly scalable performance. Horizontal scalability presents a different set of challenges from vertical scalability, which IBM has already successfully addressed in our large clustered enterprise-class servers. **Reliable, high-speed communications with low latency, shared I/O access and optimal workload balancing are key technologies that need to be in place for horizontal scalability.**

As your business grows and becomes more complex, so do your systems. Scaling lets you keep your systems moving forward while you maintain your current investments. The figure to the right illustrates how clusters provide near-linear scalability as you add more nodes.

If you have an application that has been written for a cluster, the scale factor is much higher. The application can optimize the processing power of the cluster and be more productive.



Clustering Protects IT Investment

Clusters also help protect your investments. The system you bought last year, for example, can be used as a node in your cluster, and you can just add this year's system

alongside it. You don't have to throw away your investment from last year to install a new system; you only need to add another system to create a new node.

Extreme Performance through Scalable Parallel Processing

For many cluster configurations, existing industry-standard technologies such as Ethernet and ATM switching provide ample bandwidth for internode communications and messaging. As the clusters scale to 8, 16, 32 or more nodes per cluster, optimized interconnect technologies called System Area Networks (SANs) will become increasingly important. SANs provide switched connections with extremely high bandwidth and very low latency between nodes. The RS/6000 SP switch and interconnection fabric, a technology which has allowed customers to install 500+ node clusters, will be adapted to industry-standard architecture for Netfinity servers—initially as an 8-port switch (cascadable to more) with simultaneous, bidirectional 150 MBps (300 MBps total) per port capacity, and 1.2 GBps aggregate bandwidth. Port-to-port latency in the switch is less than 300 ns. Future versions will increase bandwidth to 500 MBps each direction and increase the number of ports to 16 and beyond. This switch technology, developed by IBM Research, is capable of scaling to thousands of ports.

Cluster Perspective: Leading Cluster Management Tools

While clusters can provide impressive scalability and availability, efficient management of multiple nodes is often the most important aspect of a cluster. The goals for managing clusters are to offer a single point of control with continuous monitoring for failures and to have the ability to initiate actions if a failure occurs. IBM plans to use the single-point-of-control experiences from the RS/6000 SP Perspectives technology to enhance the manageability of Netfinity clusters. **This technology allows common system management tasks to be performed across all nodes of the cluster using a common graphical user interface (GUI).** This technology will leverage the peer-to-peer capability of the Netfinity Advanced Systems Management processors in each node. This will become increasingly more important as we move toward more complex clusters and server consolidation, especially in the near future when Netfinity and RS/6000 SP systems become more integrated and are capable of being managed from the same control workstation.

Management and middleware technologies are two components critical to horizontal scalability of clusters. **IBM Netfinity servers offer IBM Cluster Systems Management, which builds management and control features on top of Microsoft® Cluster Server (MSCS) feature of Windows NT® Server Enterprise Edition.** IBM Cluster Systems Management gives administrators of MSCS improved control of clustered installations. IBM's offering simplifies cluster administration by providing single-console control of multiple clusters and their respective cluster resources. It also can increase management control by providing resource alerting capabilities to IBM Netfinity Manager, Microsoft SMS and Intel LANDesk management software.

IBM will leverage our clustering experience and technologies to enhance this single-point-of-management capability into the future as clusters become larger and manage ever more complex sets of resources. Key focus areas include:

- Performance monitoring and tuning
- Capacity planning
- Load balancing and scheduling
- Increased alerting and management capabilities

IBM Netfinity servers will provide the tools to balance the processing load among the various nodes in a cluster. Leveraging our scalable Web-server technologies, the network traffic to your Web site will be dynamically distributed and balanced across a cluster. These are just some examples of the capabilities that IBM Netfinity servers will provide to support future requirements for horizontal scalability.

Customized Clusters for Lotus, Oracle and IBM Applications Deliver Advantages for Core, e-business and Deep Computing Solutions

IBM Netfinity clustering solutions offer you the advantages of highly available application clusters. In addition to supporting MSCS, IBM Netfinity servers currently offer solutions for Lotus® Domino™, DB2, IBM e-Network Firewall and Dispatcher and Oracle Parallel Server. These solutions, among others, are designed to handle node, disk or network interface failures in a clustered application environment with little or no downtime.

Scalable Web Server Security. In 1998 Netfinity servers will support IBM's scalable Web server solution, Websphere, secured by the IBM eNetwork™ Firewall. **This powerful combination enables single or distributed Web sites to grow nondisruptively, maintain high availability, and deliver superior end-user response, while minimizing bandwidth costs and security exposures.** The IBM Netfinity solution plans to use the eNetwork Firewall for comprehensive network and Web site security.

WebSphere Performance Pack. IBM components provide bandwidth management, caching, file administration and replication, as well as superior load balancing and availability using IBM's Web Traffic Express, eNetwork Dispatcher and Transarc® AFS® enterprise file system technology. **These capabilities let busy Web sites increase capability by allowing multiple servers to function as one.** Future enhancements are expected to increase the delivery speed to the end user even further.

Lotus Domino. IBM Netfinity servers and Lotus Domino offer you several choices for high availability and scalability in your Domino application and messaging environments. All of these clustering solutions are supported on IBM Netfinity servers. You can utilize the Domino Enterprise Server to have an application cluster of up to six nodes. With the release of Lotus Domino 4.6.1, you can now run Lotus Domino with MSCS under Microsoft Windows NT Enterprise Edition. This provides failover support for both Notes clients and browser clients connected to a Domino server in an active-passive environment. Domino Enterprise Server can also participate in an application MSCS cluster. **With Domino Enterprise Server 4.6.2 running on MSCS-certified IBM Netfinity configurations, you can now have the first active-active configuration in the industry!** Domino Enterprise Server enables failover and load balancing for up to six servers in a production environment in all three configurations. The upcoming Lotus Domino 5.0 will allow for failover of collaborative calendaring and scheduling, failover and load balancing of Web browser clients and support of mail agents.

Thin-Client/Server Solutions. IBM Netfinity servers can run Winframe 1.7a and Wincenter 3.1, and Netfinity servers are vendor-certified for Citrix Metaframe 1.0 running on Windows-Based Terminals. All these operating systems offer the ability to connect to various thin-clients including IBM Netstations. They also include an optional license pack for server load balancing. IBM Netfinity servers have been vendor-certified to run this option. **That means if you are using any of these operating systems with Netfinity servers in a multi-server configuration, you can load balance the server farm.** Users can log onto a particular server farm or a specific set of applications

and a systems administrator can monitor the server farm and allocate and de-allocate resources, depending on load. This is accomplished by viewing the NT performance monitor and adjusting the servers within the farm.

DB2 Universal Database Enterprise Extended Edition. To satisfy the needs of complex decision support and data warehousing applications, IBM has extended the rich feature set of DB2 Universal Database Enterprise Extended Edition to the Windows NT platform. Its “shared nothing” architecture allows parallel query with minimal data transfer across nodes. **Because the number of partitions has little impact on inter-partition traffic, performance can scale in a near-linear fashion as you add Netfinity servers to the cluster.** A unique partition map allows DB2 to manage the distribution and redistribution of data as required.

Oracle Parallel Server. IBM has leveraged the leadership technology of RS/6000 SP Phoenix technology to enable Oracle Parallel Server (OPS) for Netfinity systems. **IBM was the first hardware provider to achieve OPS certification with a Fibre Channel storage solution.** Because OPS relies on “shared disk” technology, Fibre Channel support is a perfect match for the requirements of multi-node OPS solutions. IBM was also an industry leader in certifying a six-node Windows NT-based configuration for Oracle's relational database management technology. For data warehousing and some Online Transaction Processing (OLTP) environments, users will be able to benefit from workload sharing and scaling up to six nodes. Working with Oracle, we can help ensure that more nodes and even greater scale factors are available in the future.

Performance Dependent on Enterprise Interoperability

Many businesses are discovering that expanding their IT operations to include industry-standard computing solutions requires interoperability with existing large and midrange systems. IBM is responding to those needs with hardware and software solutions that maximize the combined capabilities of heterogeneous computing platforms. Two examples of these solutions are the S/390 ESCON® Channel Adapter and the RS/6000 SP Netfinity Server Attachment.

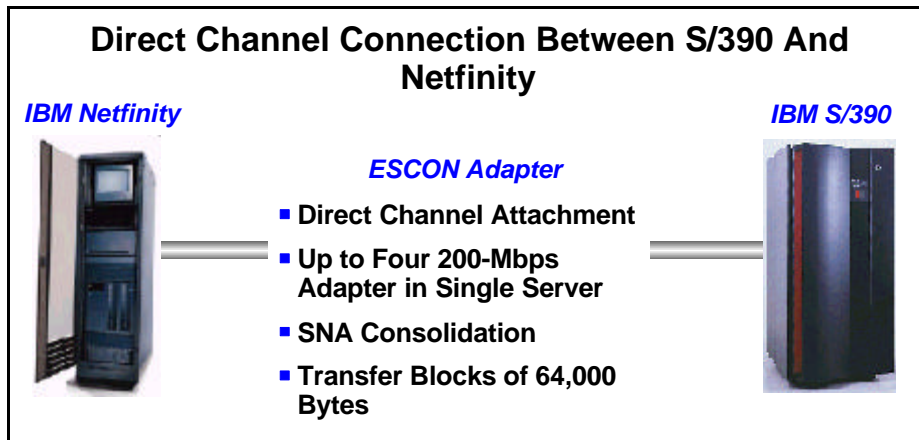
ESCON Channel-Attached Application Servers to S/390 Databases

Business-critical applications, such as enterprise resource planning, e-commerce and data warehousing and data mining for business intelligence, have traditionally been run on mainframe servers. These applications are now being complemented by a new generation of client-server applications centered around powerful, reliable and function-rich industry-standard servers. New e-business applications from companies such as SAP, Oracle, Peoplesoft, Baan and J.D. Edwards are now available for IBM Netfinity servers running Microsoft Windows NT. These servers are also being widely used as Web servers, focal points for intranets, database servers and communications servers.

In many cases, IBM Netfinity servers have to interact with mainframes regularly to retrieve and update databases and files. This is not surprising, given that enterprises with mainframes might have as much as 70% of their data stored on the mainframe. Thus, having the means to conduct efficient, high-bandwidth, quick-response-time, bulk-data transfers between servers and mainframes is essential. **Netfinity servers running Windows NT and acting as application servers can now be directly channel-attached to mainframes using IBM's new, industry-leading, high-throughput PCI bus-based Netfinity ESCON Channel Adapter.**

A Netfinity server running Windows NT can support up to four of the new 200-Mbps IBM Netfinity ESCON Channel Adapters. Therefore, it is possible to have up to four separate channel connections between the server and the mainframe. Depending on the protocol being used for the mainframe communications, a single ESCON connection will be able to support data transfer rates in the range of 32 to 110 Mbps.

Direct channel attachment is a strategic, high-performance, reliable and cost-effective way to realize Windows NT-to-mainframe communications. With aggregate throughput of 440 Mbps for the channel connections, the Netfinity server becomes essentially a co-processor to the mainframe, with both systems achieving rapid data access and bidirectional data interchange.

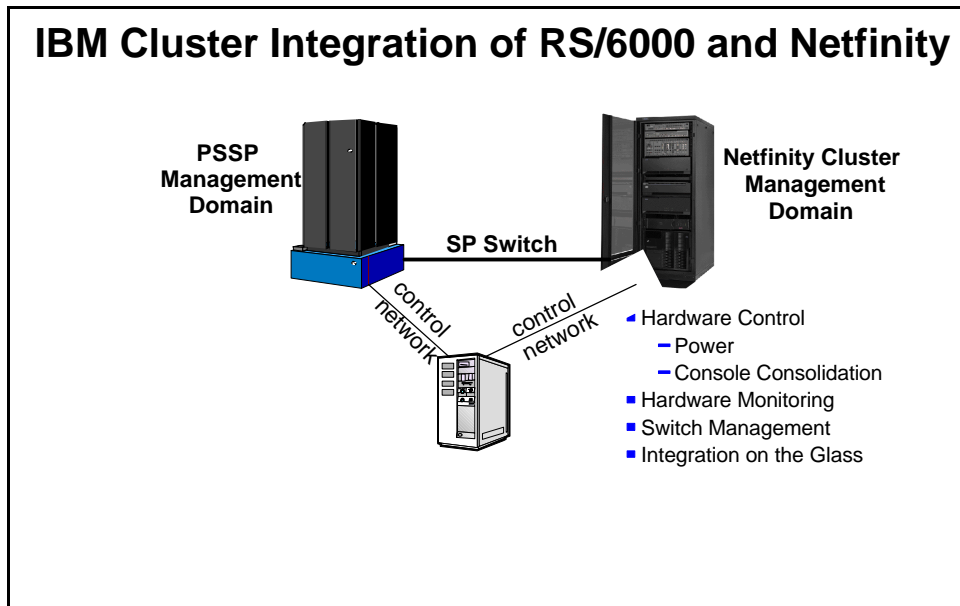


Directly attaching a Netfinity server running Windows NT to a mainframe eliminates delays caused by intermediate LANs and controllers. LANs cannot perform data streaming, which is continuous transfers of very large block sizes, because they do not support such block sizes. For example, the typical maximum block size for FDDI is around 4,096 bytes. For Fast Ethernet, it is 1,518 bytes. **The Netfinity ESCON Channel Adapter, however, can transfer blocks of 64,000 bytes.** In addition, all LAN schemes enforce a mandatory "inter-frame" gap between the transmission of consecutive frames. There can also be a "media access" delay between transmissions while the transmitting station ensures that it has full use of the LAN's actual physical layer. Additionally there are headers, trailers, and preambles that must be appended to each data frame sent over a LAN. The combined effect of the small block sizes, the gaps, delays and header and trailer overhead is such that the actual data transfer rate possible across a high-speed LAN is significantly less than with a direct ESCON channel connection.

Channel attachment also provides unparalleled performance, robust resilience and better overall management. With IBM's Communications Server for Windows NT, a channel-attached server and the applications running on it will have multiple ways to communicate with the mainframe. These communications options include SNA, TCP/IP, APPN®, High Performance Routing (HPR) and Java™. The combination of IBM Communication Server and the Netfinity ESCON adapter provide unique capabilities for Netfinity servers in S/390 installations.

RS/6000 SP Netfinity Server Attachment Integrates UNIX and NT

One increasingly important and frequently encountered mixed environment combines AIX running on RS/6000 SP platforms with Netfinity servers running Windows NT. Many customers have business-critical applications running on RS/6000 servers, but want to add NT applications and platforms. **During the first half of 1999, IBM expects to provide the capability for its RS/6000 SP and Netfinity customers to manage all of their servers from a single point of control through an ethernet connection and the peer-to-peer capabilities of the Netfinity system management processors.** Consolidated system management provided by single point of control can help you reduce the total cost of ownership.



In the future a Netfinity cluster will be able to attach to an RS/6000 SP cluster with the SP switch interconnect. The resulting cluster provides a powerful solution featuring a tiered approach to computing, consisting of Windows NT and AIX® co-operative environments. For example, this solution is very well suited to three-tiered ERP environments where customers may want to utilize Windows NT for their application servers and AIX for their database servers. It is also well suited to LAN consolidation environments.

RS/6000 SP customers who have Netfinity servers today, or who plan to purchase Netfinity servers to complement their AIX computing environment, will benefit from the mixed server environment, simplify system management and reduce operating costs through a single point of control.

Demand Universal Manageability

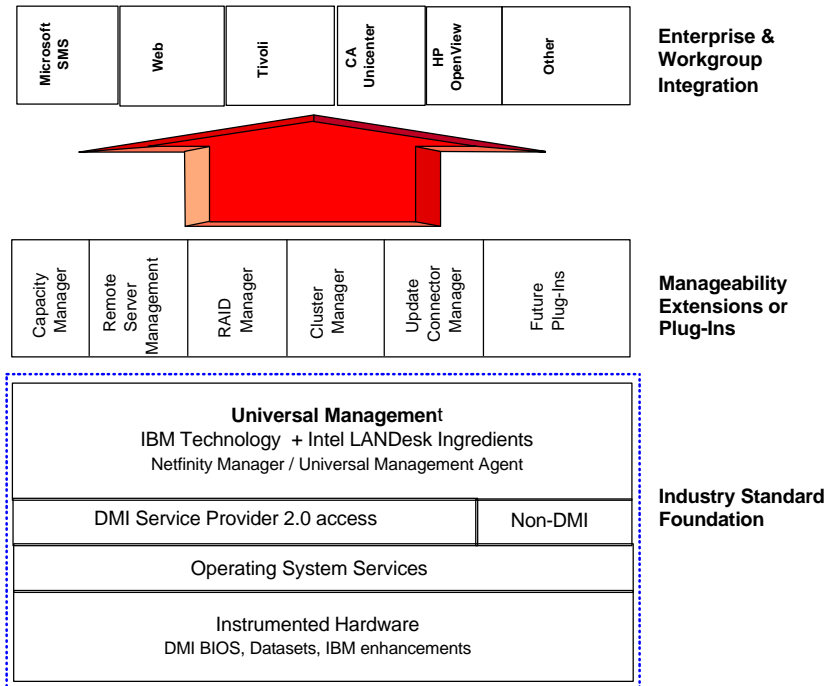
IBM knows that networks are difficult to manage. Additional technology enables more functionality, but often at the cost of increased complexity. IS personnel often struggle with ensuring control over the information systems they manage. The answer is smarter, easier system management. Advanced management can produce a wealth of benefits, many of which go directly to the bottom line—reduced down time, increased productivity,

reduced service and support costs, and end users and IS personnel empowered to focus on running the business rather than running the systems. **IBM is taking the best management capabilities from larger IBM systems and adapting them into a framework that will integrate seamlessly with a wide range of industry-standard, customer-chosen management and operating system environments.**

In addition, we are working with industry leaders such as **Intel** and **Microsoft** to deliver management solutions built on an industry-standard foundation and to drive new and emerging industry standards for manageability. These efforts will result in:

- A standards-based foundation to help users protect investments and to accommodate heterogeneous environments
- A broader set of manageability tools that plugs into the industry-standard framework and delivers even more comprehensive control
- Better integration into higher level management environments at both the workgroup and enterprise levels

The following diagram illustrates the management framework for Netfinity systems.



Being able to access and manage your servers remotely is increasingly critical to your system's availability and, ultimately, to your business. **IBM Netfinity servers are designed with fast problem isolation capabilities, hardware and software architectures that proactively monitor the server to identify and correct problems before they occur, plus system management processors that provide extensive remote management.** IBM Netfinity servers give you comprehensive control of your networked business systems, including:

- Speed and ease of installation and configuration
- Ongoing optimization of system resources
- Recommendation and application of software updates

- Notification and preventive action required to keep the system up and performing
- Rapid recovery in the event of outage
- Access to server management virtually anytime, anywhere

Proven single-point-of-management software, workload-balancing algorithms and comprehensive control over system assets will distinguish IBM Netfinity servers and help you develop enterprise-wide management procedures and policies that integrate into the system-management environment you choose for your business.

IBM Netfinity Manager software uses many protocols, including the use of the Internet, for complete access to your servers wherever they are located. We will continue to enhance Netfinity Manager as we implement IBM's Universal Management system manageability framework. The combination of Netfinity Manager technologies with other key industry-standard technologies, including **Intel LANDesk** and **Microsoft SMS**, allows IBM to provide the most robust framework and most comprehensive set of systems management tools for managing servers, client systems, notebooks and other network computers. It also allows us to provide these tools in a complementary fashion with higher-level system management offerings like **Tivoli® Management Software** and **Microsoft SMS**.

Capacity Management

The Capacity Management tools within Netfinity Manager collect server performance data and display this data over time to help with trend analysis. **With Capacity Management, system administrators can easily determine how every server on their network is performing and can help identify potential bottlenecks before they turn into problems, optimize server resource use, and help plan for future system upgrades to make the best use of your capital spending.** IBM's Capacity Management tools will be enhanced over time with additional features to analyze network resources and optimize server performance.

Remote Server Management

For decades, IBM's experience in system management processor technology has helped customers manage networks and system resources remotely. IBM brings this technology to Netfinity systems via IBM Netfinity Manager software and dedicated system management processors. **This combination allows you to perform remote management of servers that are powered off, thanks to our advanced systems management processor which uses innovative "computer within a computer" technology, as well as to remotely run system diagnostics for more accurate problem determination before initiating onsite repairs.** This technology provides remote management that is independent of the server hardware, allowing more effective systems management. When used in conjunction with IBM's light-path diagnostics for easy local problem identification, the system management processor provides a cost-effective way to troubleshoot and resolve problems that arise with your server.

Available across the IBM Netfinity line, IBM's remote management solutions will be continually enhanced to include powerful reliability, availability and service features designed to diagnose problems, provide maximum recoverability and minimize downtime, including full redundancy for vital hardware subsystems.

Providing for a Secure Environment... the SecureWay

Security is a major concern in the Intel environment. IBM Netfinity hardware security functions are all part of the IBM SecureWay® portfolio of security solutions, services and technologies. IBM SecureWay helps to ensure that appropriate security principles are consistently applied throughout IBM's products and services. IBM, through our SecureWay organization, works with government and standards organizations to raise the level of security and cryptography available to the business world, and to raise the level of security awareness among our customers.

Preventing System Failures Before They Happen

To IT professionals, IBM has always meant world-class service; service to keep systems up and running, and service to bring them back quickly if they fail. **IBM Netfinity systems extend this commitment to world-class service by using technology to help prevent failures before they happen.** Two major elements provide maximum system availability:

- Proactive maintenance triggered by advanced warning systems that allow a customer to schedule service, often without requiring system downtime
- Remote diagnosis and problem isolation for quick application of high-level skills to a problem

Built-in Failure Prevention

Most IBM Netfinity servers are designed with extensive built-in preventive maintenance features. For example:

- The **ECC memory-scrub** feature detects and cleans single-bit errors. This helps prevent the server from being brought down by the accumulation of individually insignificant soft memory errors.
- **System management interrupt (SMI)** handlers allows special software, independent of the operating system, to interface with the Error Detection, Fault Isolation (EDFI) hardware to detect and isolate hardware errors. When combined with dedicated system management processors, the SMI handler also provides a powerful mechanism to report these errors to a system log, which can be accessed through system management software or a dial-up connection.
- **PFA, extensive alerting and real-time diagnostics** are leading a growing list of Netfinity serviceability tools. This technology is resident in the Netfinity 5500 and 7000 series and includes power supplies and cooling fans, as well as in all of the disk drives supported on Netfinity systems. Also, PFA coverage includes processors, system memory, power subsystems and voltage regulator modules. PFA will continue to be enhanced throughout Netfinity X-architecture.
- **Hot Plug & Play Features** allow a system administrator to replace and add I/O adapters without having to take the system down is supported through the addition of hot-replace and hot-add PCI technology.
- **Technically-leading RAID adapters** and subsystems provide the ability to support continuous system availability in the event of one or more hard disk drive failures. Failed disks can be replaced and the lost data recovered automatically onto hot spare disks or after replacement disks have been installed to provide future

protection against additional drive failures. The Netfinity Fibre Channel subsystems support redundant controllers and redundant Fibre Channel host adapter and hub configurations to support configurations with no single points of failure.

Hot Plug and Play Leadership

The PCI Special Interest Group (PCI SIG), an association of microcomputer industry leaders, works to enhance the development of the PCI architecture. As a member of the steering committee, IBM lead the development of the *PCI Hot-Plug Specification*, Revision 1.0. PCI Hot-Plug technology allows you to insert a PCI adapter into an adapter slot while the power is connected to the host system and while other adapters connected to the host-system bus continue operating. PCI Hot-Plug functions can be described in two major categories: Hot Add, where an adapter is added to the system to expand capacity; and Hot Replace, where an adapter in the system is no longer operating correctly and is replaced by a similar adapter. **Current Netfinity systems feature both Hot Replace and Hot Add capabilities for Windows NT 4—an industry first!**

Many growing businesses have experienced the exhaustion of computer resources and determined the need to expand their systems. Traditionally, a system operator will schedule downtime for the system in order to expand its capacity. The IBM Netfinity PCI Hot-Plug solution lets you add I/O adapters while the system is still running. For example, if you need to add users or balance your network traffic, you can add a LAN adapter to your Netfinity system. The device driver will be loaded, the service started and users will continue to work uninterrupted. Today's adapters are well engineered and tested prior to delivery to market. These adapters are rugged and reliable and will generally work well for a long time without failure. With this reliability in mind, IBM is focused on *expanding* the I/O adapter capabilities on the system—without powering down or rebooting the system. **Netfinity PCI Hot-Plug is designed to reduce, if not eliminate, the downtime associated with inadequate or missing resources and the associated loss in productivity.**

IBM and Microsoft have been working on the next generation of servers and desktops to support ACPI on Windows 98 and Windows NT 5. Our desktop systems are already supporting ACPI on Windows 98. Our PCI Hot Plug solution for NT 4 is consistent with ACPI interfaces and is being used by Microsoft as the current reference for proposed NT 5 support for Hot Plug PCI.

OnForever Technology Initiative with Microsoft

Building on the success of its PCI Hot-Plug solution, IBM is working closely with Microsoft to create full, system-level support for other hot-plug technologies. **This initiative, called *OnForever*, will provide uninterrupted computing on IBM Netfinity servers running Windows NT.** OnForever is a system-level solution that will provide extended scalability, higher availability and greater manageability for the entire system, including hardware, operating system, middleware and applications.

OnForever will tie together the technologies of IBM Netfinity servers with the Microsoft Windows NT operating system. For example, the support of processor hot-plug hardware will require new capabilities to be developed, perhaps as an extension to the industry-standard ACPI definition. In addition, the hardware and firmware must be developed in accordance with this new architectural model. **OnForever will extend the benefits currently available with Netfinity PCI Hot-Plug adapters to other major system components such as processors and memories.**

Rapid Fault Recovery—Light Path Shows You the Way

IBM continually strives to produce the most reliable and fault-tolerant servers in the industry, but occasionally systems fail. Should that happen, our goal is to return the server to operation as rapidly and efficiently as possible, with the fault corrected the first time, every time. With the Netfinity 5500, IBM introduced a whole new meaning to diagnostic support. Called *light path*, IBM has introduced a 3-step lighted guidance system to assist technicians in fast identification of failing components.

- Step 1:** Our industry-leading system management processor detects a hardware component that requires replacement. It illuminates an indicator on the **front panel** of the system unit, alerting you that the system requires service. For example, a faulty fan, memory DIMM or a voltage regulator may require replacement.
- Step 2:** The system unit is opened (with no tools access), and a set of lights immediately points to the area of the system where the fault is located.
- Step 3:** An indicator light in close proximity to the faulty component is illuminated.

This system provides a clearly lit path from the front of the system to the faulty component, allowing quick, positive identification of the fault and reducing the time needed to make the repair.

The majority of server failures are software related, and often the server is located in a remote site with no technician standing by. Because of this, IBM has deployed two technologically advanced services: **MoST Connect** and **RemoteConnect**.

- **MoST Connect:** IBM developed the Mobile Service Terminal (MoST) to allow an engineer to connect a failing server to the IBM Global Network. This lets our experts in the development centers *look* into the server, download necessary files and even *take control* of the remote server. This allows high level determination of the root cause of a system failure from experts in remote sites without delay.
- IBM's **RemoteConnect** technology enables a Netfinity server to self-diagnose an alert, call the service organization on its own, and request a replacement part and service technician. The server will also notify you that IBM has received an alert and provide you with a receipt response and tracking number. RemoteConnect is currently available with the Netfinity 5500 and 7000.

Both of these unique IBM service and support features are enabled by Netfinity system technologies, providing the increased system and application availability that is a cornerstone of Netfinity X-architecture.

In the hands of IBM worldwide service personnel and Business Partners who have been trained in the service and repair of midrange and large systems, these tools will minimize the frequency and duration of unplanned outages.

Conclusion

The IBM Netfinity X-architecture drives industry-standard computing to the next level. It combines the latest industry-standard computing hardware and operating system elements with IBM's complete line of tools, middleware and system-management features. All are designed to bring you the highest levels of power, scalability, control and service. Key elements include:

- **Power.** Powerful processors, reliable and available memory systems, storage, Industry-leading PCI-X I/O breakthroughs, world-class silicon and module technology and advanced caching software; Interoperability with existing large and midrange systems
- **Scalability.** Scalable Parallel Clustered systems featuring technology derived from IBM's industry-leading S/390 Sysplex and RS/6000 SP product lines
- **Control.** Universal Management Agent (UMA) and Netfinity Manager system-management middleware
- **Service.** Technology-enabled service and support, including proactive maintenance through advanced warning systems and remote diagnostics; OnForever Technology; Hot Plug & Play Features: Light-path Diagnostics

Together these elements make up a superior platform designed to let you focus your limited corporate resources on running your business rather than running the computer network. The result is a system architecture with the performance, manageability, reliability and availability you need to transform your business—and keep it open 7 days a week, 24 hours a day, 365 days a year anywhere in the world.

Additional Information

For more information on IBM Netfinity directions, products and services, refer to the following white papers, available from our Web site at www.ibm.com/netfinity.

At Your Service...Differentiation beyond technology

IBM Netfinity 8-Way SMP Directions

IBM Netfinity Cluster Directions

IBM Netfinity Fibre Channel Directions

IBM Netfinity PCI Hot-Plug Solutions

IBM Netfinity Predictive Failure Analysis

IBM Netfinity Server Quality

IBM Netfinity Servers and Intel Architecture

IBM Netfinity Storage Management Using Tape Subsystems

IBM Netfinity System Management Processor

IBM Netfinity Ultra2 SCSI Directions

IBM ServerGuide for Netfinity and PC Server Systems

Implementing IBM Netfinity Server Management

Building next-generation performance, scalability, control and service into industry-standard servers

Integrating IBM Netfinity Manager with Intel LANDesk Server Manager

Integrating IBM Netfinity Manager with Microsoft Systems Management Server

Lotus Domino Clusters Installation Primer

Lotus Domino Clusters Overview

System Management for Servers



© International Business Machines Corporation 1998

IBM Personal Systems Group
Department LO6A
3039 Cornwallis Road
Research Triangle Park NC 27709

IBM USA

Printed in the United States of America
12-98
All rights reserved

For terms and conditions or copies of IBM's limited warranty, call 1 800 772-2227 in the U.S. Limited warranty includes International Warranty Service in those countries where this product is sold by IBM or IBM Business Partners (registration required).

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM reserves the right to change specifications or other product information without notice.

IBM Netfinity servers and PC Servers are assembled in the U.S., Great Britain, Japan, Australia and Brazil and are comprised of U.S. and non-U.S. parts.

ADSTAR, AIX, APPN, AS/400, DB2, eNetwork, ESCON, IBM, MQSeries, Netfinity, Netfinity Manager, POWERparallel, RS/6000, SecureWay, Server Proven, ServeRAID, S/390, SP, VisualAge and X-architecture are trademarks of International Business Machines Corporation in the United States and/or other countries.

Domino and Lotus are trademarks of Lotus Development Corporation. Tivoli is a trademark of Tivoli Systems, Inc.

Microsoft, Windows, Windows NT and the Windows 95 logo are registered trademarks of Microsoft Corporation. Java is a trademark of Sun Microsystems, Inc. Intel, LANDesk, Pentium and Pentium II Xeon are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries. Transarc and AFS are registered trademarks of Transarc Corporation.

Other company, product and service names may be trademarks or service marks of other companies.

THIS PUBLICATION MAY INCLUDE
TYPOGRAPHICAL ERRORS AND TECHNICAL
INACCURACIES. THE CONTENT IS PROVIDED "AS
IS," WITHOUT EXPRESS OR IMPLIED WARRANTIES
OF ANY KIND.