IBM Netfinity Cluster Directions

Bringing proven leadership in high availability and scalability to industry-standard, business-critical solutions

Executive Summary

In business today, competitive strategies are increasingly dependent on information technology (IT) to meet business-critical demands. These strategies often encounter competitive challenges. Applications and databases outgrow even the largest single systems, high availability is a critical system requirement, and system management costs are skyrocketing. Whether you're launching a business from the Web, or extending your current IT investments, you need to be able to keep your business up and running 7 days a week, 24 hours a day, 365 days a year.

Such challenges require enterprise computing solutions that provide high availability, scalability and manageability. **IBM Netfinity® delivers world-class clustering technology to the industry-standard server marketplace.**

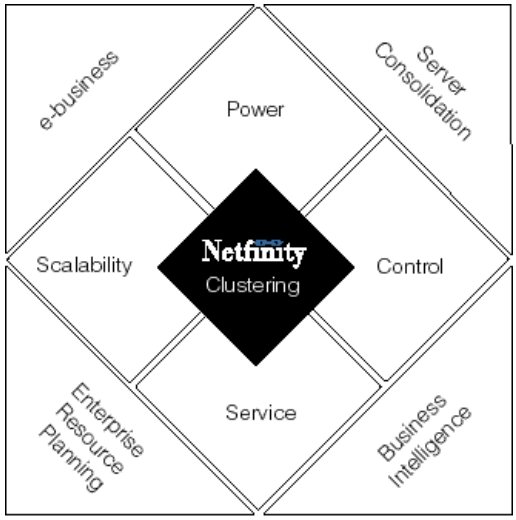

Clustering offers you the power, scalability, control and service you need for your Intel processor-based servers. It can give you:

- High availability by monitoring for failures and initiating automatic recovery mechanisms
- Access to data and shared devices in the cluster
- Improved performance and the ability to manage future growth
- Workload balancing among nodes
- A single point of control and management

This paper describes the required underlying clustering technology for industry-standard server implementation, and it illustrates the current clustering directions for IBM Netfinity systems and their important role in our overall IBM Netfinity X-architecture. Because of the dynamic nature of this industry, changes may occur in IBM's actual implementations.

IBM Netfinity Clustering Architecture

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. In response to growing requirements, many businesses have moved from uniprocessor systems to the vertical scalability of larger systems with symmetric multiprocessors (SMPs).

A key element with an SMP is the microprocessor itself. IBM works closely with Intel to keep Netfinity servers on the leading edge of the Intel microprocessor architecture road map. IBM plans to continue to provide a full range of systems based on the latest and best microprocessors from Intel's proven IA-32 architecture. However, as 64-bit IA-64 microprocessors are introduced, IBM plans to incorporate them into 4-way and 8-way SMP systems. To help achieve industry-leading performance and other features central to Netfinity servers, a team of designers with past IBM RS/6000® and AS/400® servers experience is developing the chipset for our IA-64 eight-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.

As processing and availability needs expand beyond the capabilities of SMP systems, more companies are moving to clusters for the availability, scalability and system management their demanding businesses require. Clustering has proven to be a very effective method for scaling to larger systems. IBM delivered both the Scalable POWERparallel® (SP™) and the S/390® Parallel Sysplex® Systems as large-system parallel families. These systems, built on clustering principles, provide higher levels of availability and scale to nearly unlimited computing power. Another member of IBM's server family, the AS/400 server, has provided not only highly available mid-range platforms, but also industry-leading solution integration. Based on some of the same principles and using our large-system experience, IBM is developing a new clustering architecture with its Netfinity technology.

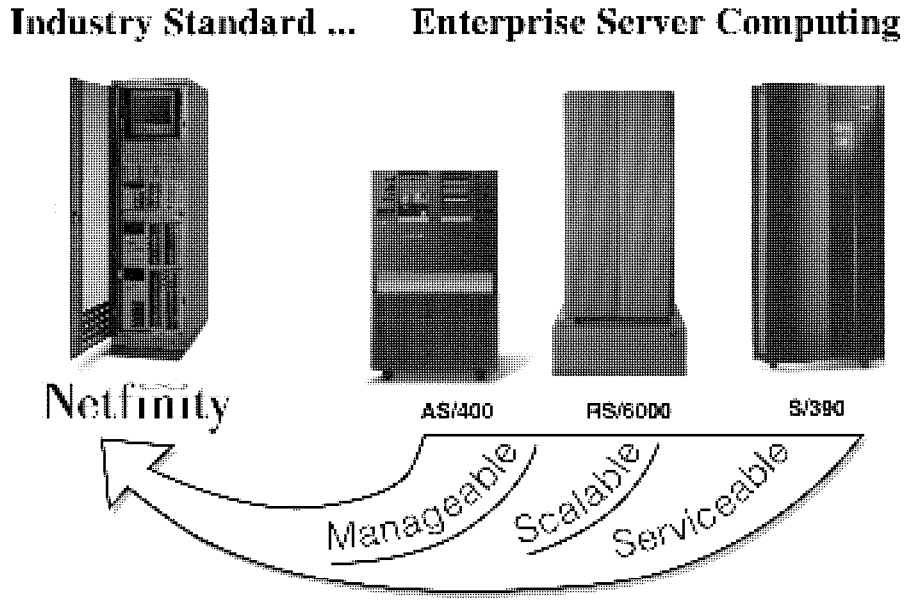

The underlying architecture is an important step in determining the characteristics of a system. An SMP system, with a centralized memory, has processor scaling limitations and executes with a single operating system image, which creates a single point of failure. Scalability and availability objectives are difficult to meet with a system of this type. Clustering, however, is based on a distributed 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.

Continuous Availability

One of the primary motivations for clustering is to provide solutions that approach continuous availability. There are three steps in accomplishing this:

1. Reduce the probability of a failure
2. Minimize the effects of failures
3. Perform 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 reducing 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 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 Analysis™ 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 the system as needed. IBM Netfinity systems are being designed to approach *continuous availability* by minimizing the effects of a failure or an outage (scheduled or unscheduled), keeping your business up and running for your end users. An example of this is IBM's *OnForever* technology initiative with Microsoft.

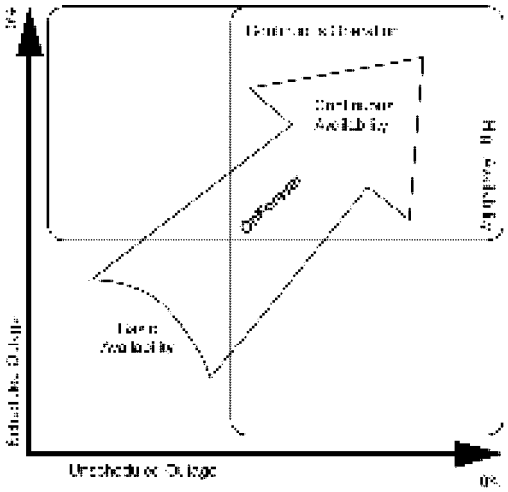

OnForever technology initiative with Microsoft. Building on the success of its PCI Hot-Plug solution, IBM is working closely with Microsoft to create full, system-level support for other hot-plug technologies. This initiative, called *OnForever*, will provide uninterrupted computing on IBM Netfinity servers running Microsoft® Windows NT®. *OnForever* is a system-level solution that will provide 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 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 definition. In addition, the hardware and firmware must be developed in accordance with this new architectural model. This will provide a platform for system diagnostics to be run concurrently with the operating system instead of the current "off-line" approach. *OnForever* will extend the benefits currently available with Netfinity PCI Hot-Plug adapters to other major system components such as processors and memory.

Near-linear Scalability

Although the vertical scalability of an SMP system can satisfy some 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[1], shared I/O access and optimal workload balancing are key technologies that need to be in place to achieve horizontal scalability.

As your business grows and becomes more complex, so can your systems. Scaling allows managed capacity and performance growth on an existing hardware base with "cluster-aware" or optimized applications. The figure to the right illustrates how clusters provide near-linear scalability to harness the processing power of a cluster as you add more nodes.

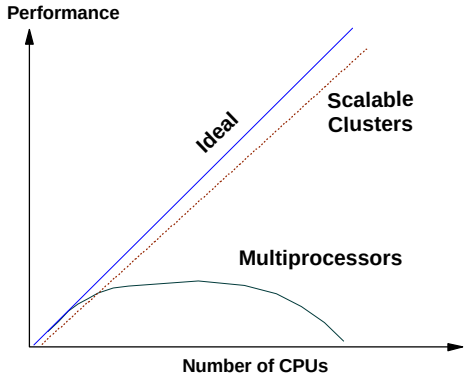

Manageability

Although 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 automated actions if a failure occurs. One of the primary motivations for server consolidation is to improve the manageability and utilization of servers, thus improving total cost of ownership. Consolidating servers and data into a central location can greatly increase efficiency due mostly to better management procedures.

---

[1] Latency is the time it takes for information to be sent from one point in a system to another.

IBM will leverage our years of clustering experience and technologies to enhance this single-point-of-management capability into the future as clusters become larger and consist of 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

Building a Clustered System

To gain the maximum benefit from clustering technology, all of the building blocks must be in place and integrated. These building blocks include:

- Server nodes
- High-speed interconnect
- Scalable I/O
- Enterprise storage
- Robust middleware
- Cluster management

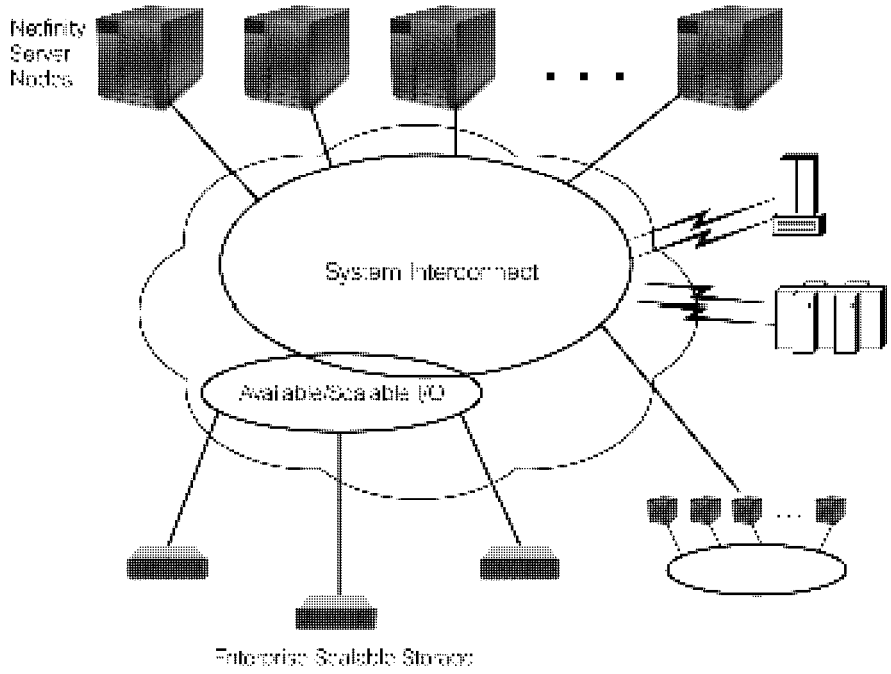

A clustering solution must give you the ability to manage and administer the cluster, and services should be available to assist you in your decision, installation and deployment activities. Also, solutions should be available for common implementations such as databases, collaborative computing, Web serving and server consolidation.

To help you with these activities, IBM has seven ServerProven™[2] Solution Partnership Centers worldwide (with two more planned to open soon). This means that you will find the best solutions from the major ISVs, wherever in the world you call home. And if it's still not possible for an independent software vendor to get to a Partnership Center, IBM can arrange for you to conduct the testing at your own location.

The IBM Netfinity clustering architecture component of IBM Netfinity X-architecture takes all of these building blocks into account. IBM is focused on customer solutions built upon the core technologies of server nodes, scalable I/O, system interconnect, cluster middleware and cluster management—interwoven with services and key system applications and data management.

Server Nodes

As part of the Netfinity clustering architecture, server nodes are interconnected onto a high-speed switch as whole computers. Our strategy is to introduce new hardware features on a timely basis: that is, concurrently with or ahead of the rest of the industry wherever possible. The introduction of new features like the 100-MHz system bus, the addition of 64-bit and 66-MHz PCI and 64-bit IA-64 technology all represent significant milestones in the evolution of server system designs, which IBM will embrace. In addition, IBM will continue to drive standards and innovation in the industry with technologies from our other servers that allow significant new system packaging and I/O expansion capability.

Building servers for e-business starts with designing server nodes that deliver the maximum power embodied in today's microprocessors. As mentioned previously, IBM's work with Intel helps keep Netfinity servers on the leading edge of the Intel microprocessor road map, while helping Intel achieve improved performance and reliability. IBM's relationship with Intel will continue through Merced and follow-on IA-64 designs.

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 servers 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 industry-standard 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 four 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. These advances in server technology will continue to drive the functionality and performance of clusters.

High-speed Interconnect

A core requirement of a cluster is the interconnection between nodes. 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[3] and very low latency

[2] Warranty and support for third-party ServerProven products are provided by the manufacturer, not IBM.

[3] Bandwidth is the peak rate of information flow.

between nodes. The proven RS/6000 SP switch and interconnection fabric will be adapted to industry-standard architecture for Netfinity servers, initially as an 8-port switch (switches are cascadable) 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. Using this same technology, customers have already installed RS/6000 SP systems with more than 500 nodes.

The basic building block of the interconnect is an eight-port switching module. Using multistate interconnection networking (MIN) techniques, this module scales the number of ports as well as aggregate bandwidth. MINs are particularly attractive because of their ability to linearly scale bisection bandwidth, a common and realistic indicator of aggregate network capacity.

In addition, by combining redundant switches within a single enclosure, functionality in the adapters, a switch service processor and software to tie it all together, multiple paths can be set up to provide fault-tolerant and load-balanced communications, thus providing superb scalability and extremely high availability.

Scalable I/O

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. 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 I_2O 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 (TB[4]) of storage, 100-MBps throughput and distances of up to 10 km (6 mi.) 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 EXP15, as well as Ultra3 SCSI and all-Fibre solutions in the future.

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, even higher bandwidth I/O bus capabilities will be needed in the future.

---

[4] When referring to hard drive capacity, TB means one trillion bytes. Total user accessibility may vary depending on operating environments.

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 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.

Enterprise Storage

Clustering technologies drive availability and capacity requirements on storage subsystems. 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.

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.

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, logical drive migration, 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.

Robust Middleware

There are two key aspects to middleware. The first aspect is that middleware provides basic services for clustering, including topology, event and group services. Microsoft Cluster Server (MSCS) offers services for facilitating execution and management of

failover scenarios, where one node in a cluster takes over for another node in the event of a hardware or software failure. To extend cluster function beyond two-node capability, other middleware is required for workload balancing and other, more complicated multi-node scenarios. IBM used clustering technology from our award-winning RS/6000 SP product line in the DB2® Universal Database and Oracle Parallel Server solutions. This sharing of technology from other IBM servers allows Netfinity servers to provide enterprise-ready, industry-tested extensions. IBM fully supports today's two-node MSCS, and we are working closely with Microsoft on its follow-on.

The other aspect is the support of clustering in all the key middleware products, including databases, Web servers, communication servers, transaction-processing software and message-queuing software. In addition to the entire IBM line of standard middleware for Windows NT, IBM is working with vendors such as Oracle to help ensure that their products are fully supported by Netfinity cluster servers. In this way we're bringing our experience in enterprise solutions to the Intel processor-based server market.

Cluster Management

With the complexity of a clustered system, manageability of the cluster becomes a key customer concern. 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 remote control capability of the Netfinity Advanced Systems Managment 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 the 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 Netfinity servers are expected to 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 should provide to support future requirements for horizontal scalability.

Clustering Solutions

Bringing together all of the parts of clustering technology discussed in the previous sections, IBM Netfinity clustering solutions offer you the advantages of highly available and/or scalable application clusters. In addition to supporting MSCS, IBM Netfinity servers currently offer solutions for Lotus® Domino™, DB2, IBM eNetwork™ 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 Lotus
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. Moreover, 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 the 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 Windows 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 an excellent 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 Enterprise Resource Planning (ERP), 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.

Service When You Need It

As clusters are utilized in highly available 7x24 situations, and scalable clusters are used to replace tasks previously provided by larger mainframe and department servers, the importance of fast, high-quality service comes to the forefront. 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 help ensure 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 and Play features:** The Netfinity clustering architecture provides available and scalable I/O as part of its base interconnect that works as well with individual system components. As growing businesses experience performance and capacity constraints of existing computer resources, they expand their systems. Traditionally a system operator schedules downtime for the system in order to expand the

capacity. When a powerful computer—a server—is shared among many employees, productivity can be lost when the server is unavailable. The IBM Netfinity PCI Hot-Plug solution is designed to reduce, if not eliminate, the downtime associated with inadequate or missing resources and the associated loss in productivity.

The IBM Netfinity PCI Hot-Plug solution lets you add I/O adapters while a server is still running. For example, if you need to add users or balance your network traffic, you can Hot-add a LAN adapter to your Netfinity system. The device driver will be loaded, the service started and users will continue to work uninterrupted.

- **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.

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 a field-service representative to connect a failing server to the IBM Global Network. This allows the representative to perform remote problem determination and launch additional resources, including product engineering if required, to solve a problem 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.

In the hands of IBM worldwide service personnel and Business Partners who have been trained in the service and repair of mid-range and large systems, these tools can be used to minimize the frequency and duration of unplanned outages. For more information, refer to the white paper *At Your Service...Differentiation beyond technology*, available from our Web site at **www.ibm.com/netfinity**.

Enterprise Interoperability

Many businesses are discovering that expanding their IT operations to include industry-standard computing solutions requires interoperability with existing large and mid-range systems. IBM is responding to those needs with hardware and software solutions that maximize the combined capabilities of heterogeneous computing platforms.

Netfinity and RS/6000 Cluster Interoperability

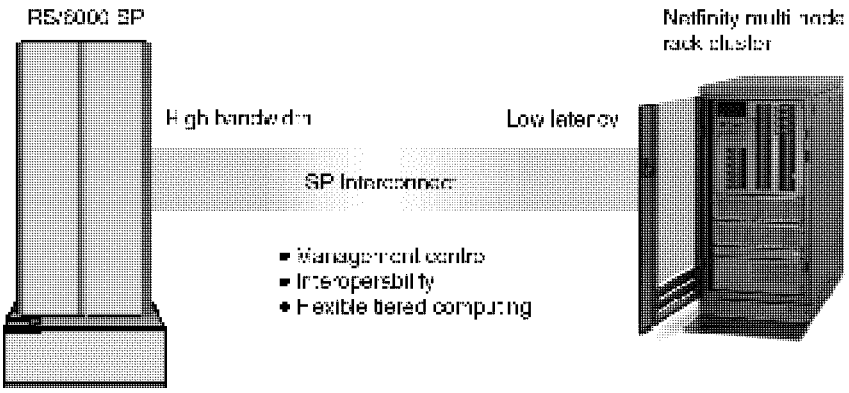

One increasingly important and frequently encountered mixed environment combines AIX® running on RS/6000 SP platforms with Netfinity servers running Windows NT. Many customers have business-critical applications running on RS/6000 servers, but want to add Windows NT applications and platforms. During 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 a 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 cooperative 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 server 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, simplified system management and reduced operating costs through a single point of control.

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 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® Adapter.

A Netfinity server running Windows NT can support up to four of the new 200-Mbps IBM Netfinity ESCON 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 100 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 block size for FDDI is around 4,096 bytes. For Fast Ethernet, it is 1,518 bytes. The Netfinity ESCON 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 can also provide 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 Communications Server and the Netfinity ESCON Adapter provide unique capabilities for Netfinity servers in S/390 installations.

Standards and Industry Relationships

As clustering of Intel processor-based servers grows in the marketplace, it is critical that these technologies be built upon open standards and that major industry providers are working together to provide complementary solutions. Using Netfinity clustering architecture, IBM's Cluster System Management is being built solidly upon open standards. IBM Universal Management, a comprehensive suite of systems management tools designed to lower the cost of ownership throughout the PC life cycle, was developed under the IBM and Intel Advanced Manageability Alliance (AMA). It delivers non-proprietary, standards-based Universal Management tools for compliant desktop, server and mobile systems. IBM Netfinity servers are leading the way by following open systems processes, briefing third-party platform providers and major application developers and sharing API specifications with the industry through appropriate standards bodies.

This approach also includes integrating with other system management products and offerings, such as Microsoft SMS or Tivoli™. IBM fully supports Microsoft's Cluster Manager. IBM was one of the initial vendors providing support for MSCS. IBM is engaged across the middleware components to lead industry standardization efforts to provide open cluster solutions.

IBM is adopting industry-standard cluster interconnect, such as high-speed Ethernet, and introducing innovative high-performance cluster interconnect and switch technology. We intend to support and influence the emerging Virtual Interface architecture. In addition, IBM supports industry-standard SCSI and SSA storage technology and Fibre Channel-attached storage for increased performance and configuration, allowing for flexibility in multi-node clusters.

While customers require complete solutions, IBM realizes that our architecture must accommodate customers who want to add their own elements or replace elements of the overall system. With IBM's participation in and support for open standards, you can be assured that you are acquiring a system that can integrate with the rest of your environment, and that offers added investment protection as you move to the future.

Conclusion

IBM brings years of experience, not only in the technology of clustering, but also in understanding our customers' environments and their business-critical requirements. We know that having the world's best technology alone isn't enough to ensure your success. It requires sharing our knowledge about how to administer and manage different environments while providing the software, and building relationships with other vendors to deliver complete clustering solutions. A complete clustering solution offers enormous benefits in high availability, scalability, and manageability. This translates into higher productivity and reduced cost of ownership for you.

IBM Netfinity cluster solutions offer scalability today with products like DB2 Universal Database Enterprise Extended Edition, the Oracle Parallel server and IBM WebSphere. Availability is highlighted further by our Microsoft Cluster Server and Lotus Notes solutions. Server consolidation with a single point of management control is another area that benefits from IBM's clustering technology. Wherever you are in this arena, IBM has the answer for you. Designed from the customer's point of view, with the solid foundation of the Netfinity clustering architecture, IBM's complete clustering solution will take your business into the future—with not only current world-class technologies, but also with a product and migration plan to let you improve all aspects of your system, including hardware, I/O, interconnect, middleware, system management and solutions.

IBM is committed to continued enhancements of the Netfinity servers by delivering outstanding power, scalability, control and service. Focusing on all aspects of computing enables IBM to drive for reduced cost of ownership for you.

IBM's Netfinity clustering solution is a vital part of our strategic X-architecture, which drives industry-standard computing to the next level. X-architecture 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 and allow us to extend the capabilities of IBM Netfinity, giving you greater levels of performance, higher availability and ease in manageability for your business-critical systems. IBM's heritage with enterprise computing and IBM Netfinity products and services give you the confidence to run your business-critical systems today and in the future.

Additional Information

For more information on clustering, refer to:

Pfister, Gregory F. 1998. In Search of Clusters, 2nd Edition. Upper Saddle River, NJ: Prentice Hall.

IBM's Cluster Solutions and Roadmap, May 1998. (Available at **www.ibm.com/servers/news/clusfnl.htm**)

The following flyer is also available through your IBM representative:

IBM Clusters for High Availability, G221-6109

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**.[5]

Achieving Remote Access Using Microsoft Virtual Private Networking

At Your Service...Differentiation beyond technology

Examples Implementing IBM Netfinity Server Management: Air Conditioning Failure

IBM Netfinity 8-Way SMP Directions

IBM Netfinity ESCON Adapter

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 Netfinity Web Server Accelerator

IBM Netfinity X-architecture

IBM ServerGuide for Netfinity and PC Server Systems

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

---

[5] Are you Year 2000 Ready? Visit **www.ibm.com/pc/year2000** or call 1 800 426-3395 (and request document number 10020 from our faxback database) for the latest information.

IBM Personal Computer Company
Department LO6A
3039 Cornwallis Road
Research Triangle Park NC 27709

Printed in the United States of America
10-98