



## **IBM Netfinity Technology Trends and Directions**

*Applying IBM large-system experience and technology to Intel® processor-based server solutions*

### **Executive Summary**

Today's information technology customers are faced with a paradox. They want industry-standard server solutions, but can't afford to sacrifice the availability and reliability that their business-critical applications demand.

Whether you're launching an e-business on 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. IBM is adapting its large-system experience to the IBM Netfinity® servers, providing performance, reliability and capabilities unavailable in today's industry-standard server marketplace.

Over the years IBM has earned its reputation as a leader in:

- Designing powerful and scalable systems for business-critical, data-intensive environments
- Creating solutions that help ensure application and enterprise availability
- Building technical support infrastructures that include skills, tools and procedures for service and support worldwide

Built on experience gained from decades of large and midrange system leadership, IBM's Netfinity servers allow businesses of all sizes to build a reliable foundation for their networks without sacrificing availability and reliability.

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.

This paper discusses the current technical directions for IBM Netfinity systems. Because of the dynamic nature of this industry, changes may occur in IBM's actual implementations.

## **Power and Scalability**

Today's computing solutions are synonymous with power. Power scales in two dimensions—vertically (within a system) and horizontally (connecting multiple systems). The key to vertical growth is a balanced system design incorporating scaleable processor, memory and I/O subsystems.

The key to horizontal growth is high-bandwidth I/O, high-speed interconnect technologies and the robust middleware required when connecting multiple systems. IBM has vast experience from its large and mid-range systems that it is applying to clustering, memory and cache controller design, as well as scaleable I/O designs.

### **Vertical Scalability**

A key element in vertical scalability is the microprocessor itself. IBM Netfinity servers will remain on the leading edge of the Intel microprocessor architecture road map and will focus on providing you the capability to upgrade systems through more than one generation of microprocessors. IBM's experience in building symmetric multiprocessor (SMP) systems becomes more important as SMPs of more than four processors come into the market. Building an eight-way computer is not the difficult challenge. Designing one that scales and provides performance benefits from the sixth, seventh and eighth processors is the trick and will be the IBM Netfinity difference as we bring our eight-way SMP systems to market (and greater than 8-way systems in the future).

Main memory must expand to balance with 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 4GB<sup>1</sup> of ECC memory. Memory growth is expected to continue in 1998 and beyond as Intel expands its architecture with 64-bit addressing support. These increases will significantly improve performance, allow support for a larger number of users and improve scaling characteristics for data-intensive applications. IBM will continue to work with operating-system and application vendors to help your applications scale to their full potential.

IBM and Intel work very closely on many of the vertical-scalability design issues, including IBM's participation in the IA64 and I<sub>2</sub>O architecture initiatives.

### **Horizontal Scalability**

Horizontal scalability presents a different set of challenges, which IBM has addresses in our large systems through clustering. Communications speed and bandwidth between nodes, shared I/O access and optimal workload balancing are the key hardware technologies that need to be in place for horizontal scalability. IBM will enhance the Netfinity servers' scalability by leveraging the hardware and software technologies from its proven S/390 Sysplex and RS/6000 SP product lines. For two- and four-node connectivity, existing industry-standard technologies, such as Ethernet and ATM switching, will provide ample bandwidth for internode communications and messaging. However, as the clusters scale to 8, 16, 32 or higher nodes per cluster, dedicated switching technologies will become increasingly important. These switches provide

---

<sup>1</sup> When referring to hard disk drive capacity, 1MB equals one-million bytes, 1GB equals one-billion bytes, and 1TB (terabyte) equals one-trillion bytes. Accessible capacity may vary.

point-to-point connections between nodes, with extremely high bandwidth and very low latency. When scalable cluster implementations begin to appear in mid-1998, gigabit ethernet and ATM switching will be the most prevalent. As the clusters grow, they'll require a more robust interconnect.

Management and middleware technologies are two components critical to horizontal scalability of clusters. IBM Netfinity currently offers IBM Cluster Systems Management, which builds management and control features on top of Microsoft's Cluster Server feature of Windows NT Server Enterprise Edition. The key features of resource monitoring and alerting, as well as scheduling and load balancing, will continue to be enhanced for future releases of Microsoft Cluster Server.

IBM Netfinity servers will provide the tools and be able 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 our customers' future requirements for horizontal growth.

### **Scalable I/O**

Scalable processor subsystems require scalable I/O for both vertical and horizontal growth. Today, IBM Netfinity servers support scalable input/output using intelligent I/O processors in our ServeRAID II (Ultra SCSI) and Serial Storage Architecture (SSA) controllers. I/O processors are designed to offload the main processor subsystem by performing I/O-intensive tasks with local processors. Both the ServeRAID II and the SSA adapter are driven by high-performance RISC processors. The ServeRAID II adapter allows Netfinity 7000 and a rack of Netfinity EXP10 storage expansion units to support up to 1.3 terabytes total external storage. With the SSA technology, the Netfinity 7000 can provide up to 1.7 terabytes external storage.

Tomorrow's network environments will demand even more. IBM Netfinity servers will incorporate ever-higher levels of performance by implementing the industry-standard enhancements to the PCI bus including support for 64-bit PCI. IBM will continue to improve the throughput of RAID adapters and embrace industry standards through the incorporation of I<sub>2</sub>O interfaces and faster I/O processors. Additional cable distance and bandwidth will come from support of Low Voltage Differential SCSI (LVDS).

For solutions requiring extremely high bandwidth, cabling distance and high availability, Fibre Channel subsystems will provide scalability to multiple terabytes of storage, at 100 MBps throughput, with distances of up to 10 kilometers between nodes. Implementations will include SCSI-over-Fibre solutions, to leverage existing investments in disk subsystems, as well as all-Fibre solutions. IBM is committed to delivering today's Fibre Channel solutions (primarily SCSI-over-Fibre) and leveraging advances in the supporting Fibre Channel chip set technologies throughout 1998. The improvements in SSA and Fibre Channel technologies would enable the more-than two-node clustering solutions and provide additional flexibility in more dispersed physical layouts.

## **Control**

IBM knows that your network can quickly become unmanageable. So, we're taking the best management capabilities from larger IBM systems and adapting them into a framework that will integrate seamlessly with a wide range of customer-chosen management and operating system environments.

Being able to access and manage your servers remotely is increasingly critical to your system's availability. IBM Netfinity systems will focus on comprehensive control through:

- Speed and ease of installation and configuration
- Ongoing optimization of system resources
- Recommendation and application of software updates
- Proactive 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 seamlessly 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 system manageability framework. The combination of Netfinity Manager technologies with other key industry-standard technologies, including Intel's LANDesk® and Microsoft's SMS management products allows IBM to provide the most robust framework for managing servers, client systems, notebooks, and other network computers seamlessly within the context of higher-level system management offerings like Tivoli Management Software, CA Unicenter and Microsoft System Management Server (SMS).

### ***Cluster Management***

IBM Cluster Systems Management gives administrators of Microsoft NT Cluster Server 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 load balancing/scheduling and increased alerting and management capabilities.

### ***RAID Technology***

IBM has long been a leader in the storage and storage-management environments. Netfinity servers benefit from this experience through advanced technology delivered in the ServeRAID II controller. Features such as on-line, local and remote management and high-performance processing set the ServeRAID II above the rest.

The Logical Drive Migration (LDM) feature supported by the ServeRAID adapter family lets you add up to three disks at a time to an existing RAID array without shutting the server down or performing a backup/restore operation. LDM also lets you change the RAID level of your existing array or rearrange the disks in the array “on the fly” to increase performance or add fault tolerance. It is the fault tolerance of the ServeRAID adapters which ensures that even if the LDM operation is interrupted, it will automatically complete when the interruption is cleared, without loss of data.

The IBM ServeRAID II adapter has a significantly enhanced initialization routine, which dramatically cuts the time required to install a new disk array. This technology allows the adapter to initialize disk storage at the moment new space is required, allowing you to begin using a new disk array almost immediately.

The creativity that made IBM's System Storage business 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 you to tune your RAID arrays to meet your needs, but in the future it will *automatically* adjust operating parameters and algorithms to optimize performance as the workload changes.

### ***Remote Server Management***

For decades, IBM's experience in systems 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 the 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.

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 down time, including full redundancy for vital hardware subsystems. For more information, refer to the white paper *System Management for Servers*, available from our Web site at [www.ibm.com/netfinity](http://www.ibm.com/netfinity).

## **Service**

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 solutions now extend this commitment to world-class service by developing capabilities that will help prevent failures before they happen. Two major factors will make this a reality:

- 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 Prevention***

Most IBM Netfinity servers will have built-in preventive maintenance features. For example:

- The ECC memory-scrub feature will detect and clean single-bit errors. This will prevent the server from being brought down by the accumulation of insignificant errors.
- The ability to replace and add I/O adapters without having to take the system down will be supported through the addition of hot-swap and hot-add PCI technology.
- System management interrupt (SMI) handlers will allow 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 will also provide a powerful mechanism to report these errors to a system log, which can be accessed through system management software or a dial-up connection.
- Predictive Failure Analysis, extensive proactive alerting and real-time diagnostics will lead a growing list of Netfinity serviceability tools. This technology can be found today in the IBM Netfinity 7000's power supplies and cooling fans, as well as in all of the disk drives supported on Netfinity systems.

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 will 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/efinity](http://www.ibm.com/efinity).

## ***Fault Recovery***

IBM's service capabilities aren't limited to the covers of the server. Being able to quickly diagnose, fix and recover from major problems is one of the reasons IBM is frequently given the moniker "IBM Means Service." Taking service and support capabilities that were previously only available on large systems is key to the Netfinity advantage.

Two key service and support offerings are **MoST Connect** and **RemoteConnect**. IBM developed the Mobile Service Terminal (MoST) to allow an engineer to connect a troubled 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. IBM's RemoteConnect technology enables a Netfinity 7000 server to send a Netfinity Manager alert directly to IBM's service organization, allowing fast and accurate maintenance. It can also notify you that IBM has received an alert and provide you with a receipt response and tracking number. Both of these unique IBM service and support technologies are enabled by Netfinity system technologies, and they provide increased system and application availability. Remote Connect is available for the Netfinity 7000 and MoSTConnect is available for all Netfinity servers.

## **Conclusion**

IBM has built an enviable reputation as a provider of large and midrange solutions. And now, IBM Netfinity solutions provide that same skill and experience in the industry-standard server market. IBM will continue to enhance the Netfinity servers using the skills, experience and technology that only IBM can offer to help you build the reliable foundation for your network—your way.

### **Additional Information**

For more information on IBM Netfinity directions, products and services, refer to the following papers and briefs, available from our Web site at [www.ibm.com/netfinity](http://www.ibm.com/netfinity).

*IBM Netfinity Servers and Intel Architecture*

*IBM Netfinity Cluster Directions*

*IBM Netfinity 8-Way SMP Directions*

*IBM Netfinity Fibre Channel Directions*

*IBM Netfinity Ultra2 SCSI Directions*

*System Management for Servers*

*At Your Service...Differentiation beyond technology*

*IBM Netfinity Server Quality*



© International Business Machines Corporation 1998

IBM Personal Computer Company  
Department T7VA  
3039 Cornwallis Road  
Research Triangle Park NC 27709

IBM USA

Printed in the United States of America  
5-98  
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.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Business Partner, IBM, Netfinity and SP2.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation. Intel is a trademark of Intel Corporation. Tivoli and TME 10 are trademarks of Tivoli Systems, Inc. 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.