IBM Netfinity Predictive Failure Analysis

Preventing unscheduled outages before they happen

Executive Summary

Predictive failure analysis (PFA) for IBM Netfinity® servers is a collection of techniques that help you run your business with less unscheduled downtime. PFA allows the server to monitor the status of critical subsystems and to notify the system administrator when components appear to be degrading. Whether you're launching a business from the Web, or expanding 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 stands behind its server products with a 3-year, on-site limited warranty. Labor and IBM parts are covered for the full duration of the warranty period, including parts identified during predictive-failure analysis.

IBM is adapting its large-system experience to the IBM Netfinity server product line, providing complete, highly available application solutions that incorporate performance, reliability and scalability unparalleled in today's industry-standard server marketplace.

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

- Designing powerful and scalable enterprise systems for data-intensive, high-volume business-critical applications
- Creating solutions that enhance application and enterprise availability
- Building technology-enabled support infrastructures that include skills, tools and procedures for service and support worldwide

Focusing on all aspects of computing enables IBM to drive for reduced total cost of ownership for you. Built on experience gained from decades of large and midrange system leadership, IBM Netfinity servers allow businesses of all sizes to build an industry-standard 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**.

Combined with IBM service and warranty support, PFA can increase the availability of your IBM Netfinity server while helping you reduce your total cost of ownership. This paper provides a brief overview of PFA for IBM Netfinity servers. Because of the dynamic nature of this industry, changes may occur in IBM's actual implementations.

The IBM Netfinity PFA Solution

IBM Netfinity PFA monitors the status of your server's critical subsystems. In most cases, replacements of failing parts can be performed as part of planned maintenance activity. As a result, unscheduled outages can be prevented and your system stays up and running for your users.

IBM pioneered PFA on DASD. This technology is currently available on IBM Netfinity servers. System fan PFA was first offered on the Netfinity 7000. Now the IBM Netfinity 5500, our newest Netfinity mainstream server, is fulfilling our promise to introduce a Netfinity server across the three lines—entry-level (Netfinity 3500), mainstream (Netfinity 5500) and high-end (Netfinity 7000).

In the past, customers had to buy proprietary hardware and custom operating systems to reach the levels of high availability now offered in the industry-standard Netfinity 5500. Capitalizing on IBM's knowledge in server technology and our heritage in enterprise computing, the Netfinity 5500 introduces a new level of price, performance and availability. It sets the standard for IBM Netfinity high-availability server PFA offerings.

IBM Netfinity PFA uses a dedicated system-management processor that runs independently of the system processor and provides the intelligence for remote management, system monitoring, alert notification, error logging and environmental monitoring. Integrating processor intelligence into the motherboard allows the Netfinity 5500 to monitor system characteristics more accurately than a plug-in adapter. Couple this processor with the Netfinity Manager software (included in all our servers at no additional charge), and the result is the best overall integrated system-management function in the industry. PFA features on the Netfinity 5500 include:

- System memory
- System fans
- Processors
- Power subsystems
- Voltage Regulator Modules
- DASD

Each is designed to provide local and remote warnings of impending failures that might cause unscheduled downtime. The potential benefit is higher server availability and reduced total cost of ownership.

You never know where or when business-critical problems will occur. IBM Netfinity PFA provides an additional level of protection against unplanned downtime. Plus IBM's 3-year, on-site warranty covers hardware problem-determination, performed on-site as well as remotely, with our latest technology and tools. Labor and IBM parts are provided for the full duration of the warranty period, including service for PFA notifications.

IBM Netfinity Manager

IBM Netfinity Manager is a powerful suite of tools and utilities designed to manage PC-based servers, desktop and notebook systems. Among its functions is the ability to process alerts, including alerts for PFA events. Actions in response to alerts include:

- Logging an alert
- Notifying the user
- Forwarding the alert to another system
- Executing a program
- Playing a WAV file[1]
- Generating an SNMP alert message
- Dialing out to a digital pager service[2]
- Taking an application-defined action

Actions are user-definable, and a detailed log tracks all alerts received. Information from the log includes the date and time the alert was received, type and severity of the alert, identification of the application that generated the alert, any text that was generated and any action taken. The screen capture below shows an example alert.

[1] Available only on multimedia systems.

[2] Available only on systems that have a modem.

System Memory PFA

The need to increase data integrity within the memory subsystem has received much attention recently. As more business-critical applications are moved to Intel processor-based servers, it is imperative that the data be protected from soft errors that can occur in memory systems. To combat this problem, IBM Netfinity servers use schemes to detect and correct memory errors. These schemes are called *error-correcting code* (ECC), or sometimes *error checking and correction*. ECC can detect and correct single-bit errors. It can also detect double-bit errors and some triple-bit errors. ECC works similar to parity by generating extra check bits from the data and storing these extra bits with the data in memory. However, while parity uses only 1 check bit per byte of data, ECC uses 7 check bits for a 32-bit word and eight check bits for a 64-bit word. IBM Netfinity servers use these extra check bits, along with a special hardware algorithm, to detect and correct single-bit errors real-time as the data is read from memory.

Semiconductor memory in today's computers is subject to both hard and soft faults. Erroneous data in memory due to a soft fault can be corrected by rewriting the corrected bit to that location. Erroneous data due to a hard fault cannot be corrected. It becomes a latent fault, whereby another bit fault in the same ECC word results in a double-bit, uncorrectable error. Because ECC corrects the data prior to use, there is little potential for uncorrectable errors to occur, even if a small number of randomly distributed, permanently faulty cells exist. IBM Netfinity memory PFA allows the hardware to remove soft errors from the system memory and focuses on identifying conditions where cells with permanent failures create the potential for an uncorrectable memory error. When one of these conditions is detected, the event is logged and an alert is generated by the service processor.

Memory PFA Alerts

When Netfinity Manager receives an alert from the system management processor that a memory PFA event has occurred, it displays the following message:

> **Machine check architecture error. Please check the system management processor error log for more information.**

This message instructs the user to examine the system management processor error log and identify the failing part. A failing memory part identified by PFA is not catastrophic . The server is still correcting errors without compromising data integrity. However, the memory should be replaced at the first opportunity.

System Fan PFA

System fan PFA is designed to give early warning that a fan could fail. Used in conjunction with redundant fans, system fan PFA provides the best protection against server outages caused by overheating.

System fan PFA uses tachometers to measure fan speed. Monitoring is provided by the system management processor. If a fan is running either too fast or too slow, it must be replaced. The IBM Netfinity 5500 has three hot-swap fans: one fan that provides cooling for I/O cards and media bays, and two fans that provide critical cooling to two CPUs and hot-swap DASD bays. The redundant fans have hub-mounted speed controllers that receive data from two external

thermistors located near the system processors. The speed of these fans will change from 1800 RPM to 2400 RPM if the internal temperature exceeds its threshold value.

Note: A failing fan might still be running and visually be indistinguishable from the other fans.

IBM Netfinity 5500 light-path diagnostics identify the failing fan. The system-error LED on the control panel and the diagnostic panel LED illuminate, identifying the fan that needs to be replaced. An error is logged and can be viewed through the Netfinity system setup or the system diagnostics. In addition, the thermistor associated with the non-failing fan senses that the temperature is rising and increases speed of that fan, compensating for the failing fan.

System Fan PFA Alerts

When Netfinity Manager receives an alert from the systems management processor that a fan is running too fast or too slow, it displays the following message:

A fan has failed.

If you receive this message, check the error log or the LEDs and identify the failing fan. Servers with redundant cooling fans, such as the Netfinity 5500, can continue operating in a safe temperature range with one failed fan. It is best to replace a failing fan as soon as possible, however, because a backup fan could fail and cause excessive internal temperatures in the server. This could lead to a server shutdown.

Processor PFA

Servers with multiple CPU processors are becoming commonplace. Typically these servers are designed for high-performance applications that require more CPU power than is currently available from a single chip. CPU faults are usually unrecoverable errors, meaning that neither the operating system nor the hardware can manage the fault conditions. If a server encounters a CPU fault, the operating system halts and users lose their network connections. When this happens, the CPU needs to be taken offline for reconfiguration, after which the server can be restarted and users can reconnect to the server. In some cases, the IBM Netfinity 5500 will do this automatically.

Some processor problems, such as L2 Cache error corrections, can be monitored by the system. If these errors become too frequent, the likelihood of a disruptive failure increases. The processor PFA feature monitors corrected errors reported by the processors. If the frequency of these errors exceeds the preset threshold, the system management processor generates an alert. Processor replacement can then be scheduled as part of normal maintenance.

Processor PFA Alerts

When Netfinity Manager receives a processor PFA alert, it displays the following message:

Machine check architecture error. Please check the system management processor error log for more information.

This message instructs the user to check the system management processor error log and identify the failing part. A failing processor identified by PFA is still processing normally, and data integrity is maintained. The processor should be replaced at the next scheduled server maintenance shutdown.

Power Subsystem and Voltage Regulator Module PFAs

An uninterruptible power supply (UPS) is the most cost-effective first step in increasing overall availability. Once you have protected and conditioned the power reaching the system itself, the next component to consider is the power supply. Because power supplies have a lower meantime between failure (MTBF) than digital electronic circuits, it's a good idea to have redundancy built in.

PFA for the power supply subsystem provides an additional measure of protection. The IBM Netfinity system-management processor monitors the power supply voltages (+5v, +12v, -12v and 3.3v). Additionally, each processor on the Netfinity 5500 is powered by Voltage Regulator Modules (VRMs). The system-management processor also monitors the voltage of each VRM.

Each voltage has an upper and lower threshold for which a PFA alert is issued. If any voltage is out of tolerance, the system-management processor sends an alert to Netfinity Manager and the error is logged to the system-management error log.

For extreme deviations of power supply voltage from the specification, the system management processor issues shutdown commands to prevent hardware damage and protect customer data from unstable circuitry that can result from fluctuating voltage levels.

Power and VRM PFA Alerts

When Netfinity Manager receives an alert from the service processor indicating an out-of-tolerance voltage, it displays the following message:

System voltage is out of specification. Please check the system management processor error log for more information.

This message tells the user to check the system management processor error log and verify the source of the problem. For a VRM problem, the Netfinity light-path diagnostics will indicate which VRM has the out-of-spec voltage.

DASD PFA

DASD PFA is based on the Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.) specification to provide early warning of some hard-disk drive failures. This allows critical data to be protected. S.M.A.R.T. is the industry-standard reliability prediction indicator for hard disk drives. IBM paved the way for S.M.A.R.T. by marketing the industry's first failure prediction capability for SCSI hard disk drives. Regular backups, combined with S.M.A.R.T.-capable hard disk drives, help safeguard against loss of data.

There are two kinds of hard-disk drive failures—unpredictable and predictable. As you might expect, unpredictable failures happen quickly, without advance warning. These failures can be caused by static electricity, handling damage or thermal-related solder problems. Predictable failures, on the other hand, are the types of failures that S.M.A.R.T. attempts to detect. These failures result from the gradual degradation of the drive's performance. In fact, 60% of drive failures are mechanical—and that's just the kind of failure S.M.A.R.T. is designed to predict.

S.M.A.R.T.-capable drives use a variety of techniques to monitor data availability. These techniques vary from one manufacturer to another. For example, a S.M.A.R.T. drive might monitor the fly height of the head above the magnetic media. If the head starts to fly too high or too low, there's a good chance the drive could fail. Other drives might monitor different conditions, such as ECC circuitry on the hard-drive card or soft-error rates. Depending on the circumstances, some drives might monitor all or none of these conditions. If one of IBM's S.M.A.R.T.-capable drives predicts it is going to fail while it's still under warranty, IBM will repair or replace it at no additional cost to you.

Netfinity Manager PFA Service

When you use the Netfinity Manager PFA service to monitor all the PFA-enabled disk drives installed on your system, you will instantly be notified when a PFA-message is generated. Also, you can configure this service to automatically generate an alert when it receives a PFA message.

The Predictive Failure Analysis Window. Each PFA-enabled physical drive is represented by an object in the Netfinity Manager Predictive Failure Analysis window. The Netfinity Manager PFA service uses two objects to help you quickly determine the status of each disk drive as shown below.

Object	Description
Solid Disk Drive	Normal: The drive has not reported any PFA messages.
Shattered Disk Drive	Warning: The drive has reported one or more PFA messages and might be failing.

Information that will help you identify the drive includes adapter, physical unit number (PUN), logical unit number (LUN), physical drive value, logical drive value and size.

The PFA Options for Drive Window. Use the PFA Options for Drive window to view additional information about the selected PFA-enabled drive and to configure Predictive Failure Analysis service options specific to the selected drive.

Detailed Disk Drive Information. The PFA Options for Drive window duplicates the drive-specific information from the Predictive Failure Analysis window, and also provides the vendor ID, product ID, product revision level and status.

Predictive Failure Analysis Options. In addition to providing detailed drive information, the PFA Options for Drive window allows you to:

- Configure PFA alert-generation options for this drive
- Simulate a PFA warning message for this drive
- Reset the drive from "Warning" status to "Normal" status.

DASD PFA on a RAID System. IBM Netfinity PFA messages generated by PFA-enabled disk drives that are part of a RAID array cannot be detected by the PFA service. However, PFA-enabled disk drives can be monitored using the System Monitor service's attribute monitors. The attribute monitors settings can be used to set up an alert for each disk drive in the RAID array.

The S.M.A.R.T. hard-disk drives communicate to the RAID adapter when a PFA event occurs. The RAID adapter passes this information to Netfinity Manager's attribute monitors, which send out the alert for the PFA event.

Conclusion

IBM has a rich heritage in high-availability systems and features based on our large-enterprise systems. Over the years, IBM has played a leadership role in pioneering and standardizing PFA for DASD through the use of S.M.A.R.T. technology. Now, beginning with the Netfinity 5500, we're extending PFA capabilities in the Netfinity server marketplace by providing PFA technology for fans, memory, processors, VRMs and power.

IBM's PFA features are automatically enabled on Netfinity servers. Each PFA algorithm provides monitoring of conditions that indicate the need for preventive maintenance. When a condition for maintenance is detected, the customer is notified and the event logged. Part replacement is covered under the IBM 3-year, on-site warranty and can be performed without unscheduled downtime, thereby helping you reduce the total cost of ownership.

IBM's experience with enterprise computing, our leadership role in the development and implementation of new technologies, 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 IBM Netfinity directions, products and services, refer to the following white papers and information briefs, available from our Web site at **www.ibm.com/netfinity**.

IBM Netfinity Technology Trends and Directions

IBM Netfinity PCI Hot-Plug Solutions

IBM Netfinity Cluster Directions

IBM Netfinity 8-Way SMP Directions

IBM Netfinity Fibre Channel Directions

IBM Netfinity Ultra2 SCSI Directions

IBM Netfinity Servers and Intel Architecture

IBM Netfinity Server Quality

System Management for Servers

At Your Service...Differentiation beyond technology

IBM Netfinity Storage Management Using Tape Subsystems

IBM Netfinity ServerGuide for Netfinity and PC Servers

IBM Netfinity System Management Processor

Hard Drive Interfaces (Information brief)

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

Printed in the United States of America
7-98

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.

IBM and Netfinity are trademarks of International Business Machines Corporation in the United States and/or other countries.

This publication may include typographical errors and technical inaccuracies. The content is provided "AS IS," without express or implied warranties of any kind including merchantability or fitness for a particular purpose.