

Paper #4130

## **High Availability Service Strategies for the Next Millennium**

Jeff Louie  
Hewlett-Packard Company  
100 Mayfield Avenue  
Mountain View, CA 94043  
(415) 968-5600

and

Neal Clapper  
Hewlett-Packard Company  
3404 East Harmony Road  
Fort Collins, CO 80525-9599  
(970) 299-3800

## Introduction

### New Parameters for Availability

As boundaries blur for international trade and more companies go global, IT processes will need to be “up” more of the time. Workers will continue to conduct their business in mobile, remote environments. Corporate decision-makers, dispersed around the globe, will need enterprise-wide access to information that previously was obtained in batch mode. Indeed, faced with pressure to decrease their time-to-market, companies will readily embrace new technologies, pushing today’s data centers off mainframes and onto distributed, open systems.

To accommodate all this mobility and expansion, more companies will need to operate around the clock. A Client/Server Computing Worldwide Technology Analysis by *Dataquest* (September 30, 1996) estimated that more than one-third of all worldwide businesses are run on a seven days-a-week, 24-hours a day basis. That number is expected to grow about 60% by the end of the decade.

#### **Figure 1. Percent of Worldwide Businesses Operating on a 24x7 Basis**

Why? It all boils down to service—customer service. Most products have become commodities, with little differentiation in cost or quality. With competitors leapfrogging each other, the only way to get ahead is to be more “available” to customers. As FedEx discovered when other delivery companies mastered its ability to “ship any package anywhere, anytime,” the company could differentiate its overnight services by providing customers with such conveniences as “point and click” tracking.

What once was a competitive advantage is now a baseline requirement or commodity. For most global companies, service has become *the* product. The key to providing such service is a well-planned strategy.

### Paradigm Shifts from Mainframe to Open Systems

Increasingly, high availability is being delivered by a very different kind of computing architecture. Proprietary mainframes have been replaced or off-loaded by open, networked systems supplied by a range of vendors. With the explosive growth of the personal computer, the computing industry has emerged from its relative infancy and consolidated around industry-wide standards.

Today’s global enterprise could not survive without these standards. Now technology is built from interchangeable parts and common reusable components from a number of suppliers. Like the telephone, today’s open, distributed computing allows connections among people anywhere in the world.

Such free-flowing, transparent connections would not have been accepted in the traditional hierarchy of business. But as business giants move from authoritarian rule to team-based structures, they need support systems that can operate as distributed workgroups in real-time, 24 hours a day. In short, their systems need to be flexible and highly available so that companies can satisfy their

customers and shareholders. Consider, for instance, the computing shift that Chemical Bank of New York recently undertook. Guiding every transaction at Chemical Bank is the philosophy that the customer is boss. With the goal of forming long-term customer relationships, employees had to be equipped with flexible, distributed systems for their decision support.

According to the international bestseller, *Paradigm Shift*, by Don Tapscott and Art Caston, Chemical Bank became one of the first organizations to establish an enterprise architecture for IT. Earlier attempts at choosing technology standards had not gained wide acceptance from a well-entrenched and very dispersed IT community, which included other acquired regional banking operations. Faced with tremendous challenges, Bruce Hasenyager, senior vice president at Chemical Bank, crystallized IT's dilemma this way: "Is it possible to create an architecture for IT when we haven't completed our strategic business and re-organization planning?"

It was agreed that the best situation for Chemical was to be ready to support "any conceivable change in business direction" from its executives. Hasenyager's team decided to develop architectural principles and build generic technology models as a framework for establishing technology standards. This required buy-in from the entire IT community, which meant executive sponsorship. Four task forces composed of senior managers and professionals were formed, each meeting approximately one day a week over three months.

What was initially grudging participation became enthusiastic support. One of the early outputs of this process was a list of principles, each with the rationale and implications for the entire organization. Both management and IT professionals concurred that Chemical Bank could not afford to stay locked into the centralized technology of its past. Not only did it shift from proprietary mainframes to open systems, but it replaced its hosts with client/server networks.

"Vested interests tend to fight change," reflected author Don Tapscott in a presentation to staff at Hewlett-Packard Company. "New paradigms are often received in coolness, even mockery or hostility. The leaders of the old are often last to embrace the new. I believe every company needs to find within itself the leadership for transformation."

## **New Requirements for IT**

Delivering smooth service rests largely on the shoulders of IT managers. Not only must their systems support the growing demands of their business, but their technology needs to run at peak efficiency at all times.

Forced to balance business needs with changing technology, IT managers must meet customer demands with strong support services. From a business viewpoint, consolidations, rapid growth in transaction volumes, additional needs for large data repositories and continually changing business alliances all put increased pressure on IT to meet customer demands for availability. Correspondingly, the growth of open systems, electronic commerce and the Internet bring

additional technology demands. The pressure has never been greater for IT. First, IT managers must justify every upgrade, every investment in new technology by real business needs. Then, regardless of cost or available resources, they must ensure their systems do not fail—not only for customer processes that generate revenue, but for the increasing base of internal users that make the organization more efficient.

## **What is High Availability?**

Availability has come a long way. What once were costly, complicated options to make computing resources accessible are now critical components in day-to-day business activities. From global corporations to small businesses, enterprises in virtually every industry depend on the availability of information systems. Consider this sampling:

Airlines, hotels and other ticket- or reservation-driven industries rely on their automated reservation systems to sustain their businesses;

Banking and securities firms depend on information systems to conduct trades, monitor the stock market, respond to changing market conditions or to process payments;

Newspapers, magazines and other media rely on networked systems to report news as well as for online research; and

Home shopping via catalogues, the Web or television are all based on information systems.

High availability translates directly into greater levels of productivity, increased competitiveness, better control and an increase in annual revenue. It saves both time and resources. In fact, without high availability, a business cannot maintain its mission-critical applications.

## **Hierarchy of Availability**

So what is high availability? According to International Data Corporation, “A system is considered to be highly available if, when failure occurs, data is not lost and the system can recover in a reasonable amount of time.”

**Figure 2. Hierarchy of Availability**

The following hierarchy defines high availability in terms of the amount of downtime per occurrence. As a customer moves up this hierarchy, the risk associated with either planned or unplanned downtime increases and the requirement for availability increases.

**Reliable systems** provide the baseline for measuring availability. The service interruption at this level of the hierarchy varies from 30 minutes to four hours, in addition to database recovery. Users must repair failed components, reboot, log on to the system and then restart and recover applications such as the database. Planned maintenance may occur during operation hours.

**Highly available systems** increase system availability by eliminating single points of failure (SPOF) with add-on products. Downtime is reduced to 10 to 30 minutes, excluding database recovery. Users must reboot, log onto the system and then restart and recover applications. Planned maintenance generally is scheduled outside of operation hours and may be rescheduled as required.

**Fault resilient systems** are able to reduce downtime or service interruption to one to five minutes per occurrence. This level of resiliency is achieved through a multi-system architecture. Multiple, independent computer systems are arranged into a networked configuration in which they interoperate through high availability software. Users will need to log on to the system and restart and recover applications.

The use of **redundant hardware** eliminates SPOFs and minimizes routine maintenance. Planned maintenance, such as operating system upgrades, can be accomplished by moving the applications from one system to another. The second system continues to service requests while the original system is updated, thus maximizing availability. Other tasks, such as database backup, are performed online. Unplanned failures are detected automatically; services that originate on the failed system are moved to another system within the environment.

**Continuously available systems**, which top the hierarchy, provide 7×24 availability, regardless of planned or unplanned downtime. The service interruption associated with continuously available systems is three to 15 seconds. The recovery is fully automated and transparent. In fact, transparent application recovery and transparent client reconnection to applications are key elements of continuously available systems. Planned maintenance is performed online.

## Creating High Availability

High availability doesn't just happen; it must be built on reliability, supportability and manageability. To maintain and improve systems without disruption requires a strong combination of reliable technology, strong support and a flexible infrastructure.

With each new generation of systems, however, the cost of availability is likely to decrease as new technology incorporates high availability features. By the end of the decade, IS organizations can anticipate that high availability solutions will be standard features of application servers. However, the value of IT will be its ability to provide end-to-end service.

**Figure 3. The Decreasing Premium for High Availability**

## The Strength of Service Level Agreements

Today, there is a more acute alignment of IT and business goals company-wide. Despite an increasing number of end users that rely on IT for their business processes, the percentage of overall IT spending actually has decreased—as much as 30% in the last ten years. While IT jobs are much more complex, end users are not concerned with “behind-the-scenes” technology—protocols, components, libraries. Instead, they simply want to know that the applications will enable them to be productive.

In a distributed environment, end users are subject to hundreds of applications variables ranging from middleware and databases to development tools. Unfortunately, there are no standards on managing or integrating into new and legacy environments. This is further exacerbated by the realization that IT organizations may not have planned for IT management when client/server applications (particularly mission-critical applications) were developed.

These factors combine to create a new role for central IT organizations. While internal lines of business (LOB) are demanding a commitment for uptime from their IT organizations, they also may regard the internal IT organization as a commodity. If the LOB can't obtain the service it needs from IT internally, it is likely to show no reticence in outsourcing IT services.

IT increasingly is forced to be more proactive and responsive to internal users. Today's IT departments must consider its core competencies and how it can leverage these strengths across the company to bring added value to the organization. Likewise, it also must evaluate its weakest areas. In its effort to meet the needs of its users, IT, in turn, must manage the enterprise from an end-user's perspective, looking to outside vendors for guarantees of uptime.

In many instances, IT must function with an actual service level agreement (SLA) with its various internal users within the company. With a strong internal SLA, IT departments can return to the business of supporting their users, knowing that their external vendor SLAs, in turn, will support them.

FedEx depends on its tracking systems to be available to customers at all times. Without reliability, FedEx could not perform. Its results -- measured by the percentage of packages delivered on-time -- generate FedEx's revenue and differentiates the company from its competitors. In a unique SLA, FedEx allows any user to access its web page to determine the location of a package. If a response is not available within 30 minutes, FedEx will refund the delivery charge.

## Equation for Success

Availability and reliability equal performance—performance that can be measured. These are not luxuries; they are essential to a business' survival. System outages can damage future relationships with customers. If clients can't reach a business—especially if the interface is a web server—they may give up and go to a competitor. As if lost business were not tragic enough, system outages may result in misinformed or incomplete decisions. If a company's server is down, critical information is not reaching key decision-makers. Judgments made on partial data could have a detrimental impact on operations.

“The ultimate scenario in today's computer industry is continuous operation,” reports the September 30, 1996 issue of *Client/Server Computing Worldwide*, published by Dataquest. “In this environment, the computer system never goes down and is always available. When a failure occurs, other systems take over the operations, and the user is unaware of any disruption.”

## **Heal Thyself, But Not With Hardware Alone**

If the goal of high availability is guaranteed up-time to the customer—internal or external—the ultimate in tools are integrated environments that can predict outages and begin corrective action before the system starts to go down. In a predictive, self-healing IT world, a human *somewhere* will be alerted and the component can be fixed before failover procedures begin.

The particular needs of a business and its operating environment will determine a company's approach to high availability. The company should begin by identifying key processes and elements within its environment that require high availability. If the calculated costs of downtime are high or if downtime affects critical business issues, the company will likely determine that there is a strong need for comprehensive high availability services. It also is necessary to understand the technology that is needed to support a high availability environment.

The strength of an IT environment is limited by the strength of its weakest link. If any one component is weak—regardless of whether it is hardware, the operating system, the network, middleware or applications—the solution for the entire enterprise can collapse.

Building effective, highly available systems for decision support requires a strong foundation. First, the technology must be reliable and resilient enough to withstand maximum performance with both internal and external change. Second, the supportability must be designed into the technology. Third, the systems must be capable of easy, automatic management so that human error, which accounts for 20% of all unplanned downtime, will have minimal effect.

This foundation, which makes up the new paradigm, is supported by three pillars: reliability, built-in supportability and availability management.

### **Figure 4. Three Pillars of Availability**

## **Reliability**

Reliability applies to the technology. It is designed into products, both hardware and software and is available upon purchase of those products.

Each business environment requires a different set of parameters to fine-tune corrective actions according to different needs. Moreover, a high availability solution should prove itself reliable over time at numerous local and remote installations.

Certainly, there should be no single point of failure in a high availability solution. The solution can minimize the risk of downtime if it includes a contingency plan in the event a failure does occur.

## **Built-In Supportability**

Built-in supportability, the second pillar driving availability, encompasses both predictive capabilities as well as remote diagnostic tools coupled with training and education. If built-in supportability is incorporated into the technology and the staff are well-educated, problems can be diagnosed remotely and resolved faster.

Predictive capabilities include designing in features (such as online diagnostic tools or diagnostic software) to avert problems before they occur. Such efforts can include software that alerts staff to impending failures or disks that can be swapped or added while the system is still on-line. Reactive efforts are necessary to restore availability in the event of a failure. The clear objective is to get the system functioning again *before* solving the technical problem.



## Availability Management

The third pillar represents one of the greatest challenges creating availability—availability management. It is achieved through internal processes and external support contracts. It is the effort to increase the solution availability and the overhead required to operate the high availability environment itself. Availability management includes the ability to manage and operate the computing environment as well as the proactive processes for performance management, configuration management and change management.

The objective of a high availability solution is to ensure that application availability to the end user is not compromised by any changes. Operation should be non-intrusive; system users should not have to know it is in place. Availability management must integrate with major management frameworks and allow rolling upgrades—minimizing downtime.

To lessen the frequency and impact of human error, high availability systems must be managed effectively. The operating system must be able to monitor and recover key resources and services as well as integrate with a complete set of communication protocols. However, proven processes that can isolate a problem, actively manage its parts and prevent a situation from escalating also are important. For example, network node managers can be purchased to handle the network, automatically monitoring and restarting network connections and specific adapters.

## HP's High Availability Service Strategy

### Developing a High Availability Service Alliance

In contemplating high availability services for mission-critical environments, a company is actually seeking a support alliance. In such an alliance, the company wants more personal attention, a proactive partner who will look out for their interests and who regards the company's downtime as seriously as if it were its own. The right partner can assist with technology selection, deployment and exploitation.

It is imperative that when selecting a partner for such an alliance, a company must consider the range, reach and reputation of all candidates. A partner must have a comprehensive knowledge and expertise, as well as extensive offerings, to meet the company's current and future availability needs.

A company's relationship with its service provider extends beyond just technical consulting. The partner also must understand the cultural, regulatory and market requirements of the organization's specific business. The right support alliance will contribute to a company's overall business goals and actually may enable them to meet these goals more effectively and efficiently.

As businesses implement more complex applications and migrate more business-critical applications to client/server environments, they will rely more heavily on IT for business basics. Similarly, concerns about system maintenance and support costs also will continue to grow. A partner for availability services not only must keep the system running, it also must fix it fast and fix it once in the event of an outage.

Service providers must deliver what they promise. Fast system recovery must be *the top* priority. A solid support alliance can deliver value to every aspect of a company's organization. It can identify potential areas where availability may be at risk and then minimize these exposures. But the right people with the right processes and tools must be available to support the company's technology.

## **HP's Solid Record of Availability Services**

Hewlett-Packard Company can help companies to achieve high availability. HP is positioned to help customers design, implement and operate their IT infrastructure to ensure appropriate levels of availability for mission-critical applications.

HP has had success in the high availability server market dating back to the first available commercial RISC operating systems running MPE. Because HP-UX is a stable, mature operating system, recovering applications and systems has eased with experience. Since 1985, when the HP 3000 evolved, to today's high-end, HP 9000 enterprise-wide servers, a range of global businesses have turned to HP for reliability and support. From small mail order houses to the world's largest insurance companies, reliance on HP has provided availability so that businesses can do what they do best—serve their customers.

## **HP's Strategy for Reliability**

Reliability consists of creating highly available hardware and software while minimizing the impact of a failure, the need for planned downtime and the time to diagnose and repair any outages. HP's reliability solutions include preventive features and failover capabilities built into the hardware or software that minimize planned downtime and diagnostics.

Preventive features, such as predictive software and self-healing technology, reduce the system failure rate and eliminate the risk of human error. These features and easy-to-use operations also can include patch tools (to proactively distribute, manage and track patches), as well as better testing of maximum configurations and targeted third-party solutions.

Failover capability minimizes the impact of a failure with hardware and software resiliency. Planned downtime is reduced with online patching and online addition or replacement of disks, LAN, WAN and power supplies. Hardware diagnostics and software core dump analysis minimize the time to diagnose and repair problems.

Reliability is a strong foundation for the entire spectrum of HP products. HP 9000 servers are designed from the ground up to provide industry-leading levels of reliability. HP 9000 D-Class Enterprise Servers come standard with memory page deallocation and hot-plug internal disks to enable online replacement or addition of disks and integrated uninterruptible power supply (UPS). Like the larger K- and T-Class Enterprise Servers, they also provide CPU deallocation and automatic restart.

A broad spectrum of other products also enhance HP 9000 reliability and high availability. MC/ServiceGuard, HP's industry-leading, high availability software protects network configurations and keeps mission-critical applications running.

HP also enhanced its MC/LockManager software with Oracle's Parallel Server relational database product. This enhancement, which enables four HP 9000 servers to be clustered together, provides a robust, integrated solution for data integrity and application availability. Also, its performance scalability provides protection for mission-critical applications on top of HP-UX. For the Internet and intranet, HP's Praesidium Security strategy provides centrally managed end-to-end security. VirtualVault, one component of this strategy, acts as a gateway between the Internet and a company's intranet for Web implementations that require access to client/server or database applications on the company's intranet.

HP's latest suite of leading cluster management tools, Network Connection Policy Manager (NCPM) makes access to a network environment as easy as accessing a single system while providing connection time load balancing.

## **HP's Strategy for Built-In Supportability**

HP is committed to providing comprehensive, built-in supportability that includes predictive capabilities, remote diagnostic tools and educational efforts. The primary focus for achieving this through technology is centered on four well-defined programs:

- High Availability Observatory
- Patch Management Technology and Processes
- Support Delivery Infrastructure
- Embedded System Supportability Capabilities

### **High Availability Observatory**

The main goal of the High Availability Observatory (HAO) is to observe the health of a mission-critical customer's environment and take proactive and reactive steps to keep their systems up-and-running, fix their system faster when a problem occurs, and manage their system events. The specific technologies planned for the HAO are High Availability Configuration Tracker, AutoNisp, Failure Event Notification, HP-UX Remote Service Node and Secure Router Connectivity.

The High Availability Configuration Tracker will gather system configuration data on a daily basis enabling the HP Response Center engineers to take proactive configuration-related actions and deal with reactive problems more effectively.

AutoNisp is a tool that collects and transmits a logical view of a customer's network map enabling the HP Response Center engineers to monitor changes to the customer's network and provide better network support in the event of a problem.

Failure Event Notification is a capability that is being added to HP's Hardware Predictive solution that will inform the customer and HP Response Center engineers of failures due to disk arrays or a MC/ServiceGuard switchover.

HP-UX Remote Service Node is a computer system that sits at the customer's site for collecting pertinent support information and reporting it back to HP. While Secure Router Connectivity is a conduit which provides secure, two-way connectivity between HP Response Centers and end-user customers.

### **Patch Management Technology and Processes**

Many efforts from previous years are being combined to finally begin tackling this ever-present customer concern. Extension Software is a bundled solution which takes integrated, well-tested patch bundles and includes them on the latest software distribution media. Custom Patch Manager is a tool that is provided in HP's Electronic Support Center, an on-line electronic support solution, that allows a customer to upload their specific configuration and then let the customer filter and recommend patches.

### **Support Delivery Infrastructure**

HP Personnel are extremely important when considering the satisfaction of the customer. Several tools are being provided to HP's Account Support Engineers (ASE) and Hardware Support Engineer (HSE) to ensure their success and efficiency in supporting large, high availability customers. Each Engineer is equipped with a mobile laptop for use when traveling and on-site calls. Information that is delivered are improved account management features, customer commitment manager, improved inter-HP communication tools, problem resolution enhancements, and remote access to HP systems.

### **Embedded System Supportability**

Base level capabilities to improve system availability for hardware and software are tools for troubleshooting CPU failures, cache error recovery capabilities built into new hardware platforms, improved diagnostic platform and increased backporting of patches to prior O/S versions.

HP's long term direction is to provide customers the best support and diagnostic tools to keep their system running and/or fix their system fast when it does go down. These solutions include new analyzers, NT support, reliability improvement, several sophisticated analysis capabilities that will be integrated into an enterprise-wide system management solution, as well as the ability to provide continuous computing.

To meet the challenges of integrating UNIX and Windows NT environments, HP developed a joint strategy with Microsoft called "Simplified Enterprise Computing." With this worldwide integration initiative, HP gives customers a means to leverage existing systems while reducing the costs of operating multiple environments.

This new strategy consists of three primary elements:

- Products and services that simplify the management of customers' computing environments and help to reduce their total cost of ownership;
- Collaboration on offering Windows NT-based solutions to HP customers; and
- Solutions for heterogeneous computing environments that use multiple operating systems including UNIX, Windows NT and others.

Educational efforts also are an integral component to HP's strategy for built-in supportability. Because human errors account for 20% of unplanned downtime, it is critical that IT staff are adequately trained. HP offers a range of educational and training programs that will better prepare personnel to minimize errors and meet the challenges of high availability. These programs include such topics as The Fundamentals of the UNIX System, Inside the HP-UX Operating System, Hands On With High Availability, MC/ServiceGuard, and OpenView Network Node Manager.

### **HP's Strategy for Availability Management**

HP offers a broad continuum of services that enable customers to scale and manage their specific availability needs based upon their computing environment. These services begin with Standard Support, step up to Critical System Support and extend to Business Continuity Support.

**Figure 5. Availability Management Model**

### **Standard Support**

HP's Standard Support includes software and network support, licenses for updates, software media and documentation. It ranges from eight hours a day, five days a week to 24 hours a day, 365 days a year. Beyond Standard Support, HP also offers two special high availability services: Critical System Support and Business Continuity Support.

### **Critical System Support**

HP's Critical System Support (CSS) is a suite of services for businesses running enterprise-class computing environments that require very high systems availability. It is a modular, flexible service designed for companies that want a proactive response and immediate reactive support beyond HP's standard services.

CSS provides hardware, software and network support for critical systems by reducing the frequency of systems failures, and helping to recover systems when problems occur. As the first vendor to offer a six-hour hardware call-to-repair commitment, HP assigns a CSS support manager who leads a team of experts in maintaining high-availability computing environments. These engineers provide guidance and assistance on technical and operational issues. The team's efforts result in fewer system problems, shorter planned downtime and a more effective use of technology.

In the event of an outage, the CSS system recovery team responds immediately with access to HP's critical-parts network and immediate dispatch of personnel to repair system hardware within six hours. Phone-in software assistance is available 24 hours a day, 365 days a year.

CSS also offers quarterly patch management reviews and an operations check-up. The operations check-up looks at the entire environment to evaluate any single point of failure and makes recommendations to create a more robust high availability environment.

HP has created the Mission-Critical Server Suite Program to incorporate best-in-class products and services to help customers achieve new levels of availability. The foundation environment provides uptime assurance at 99.95% availability. It includes the hardware, operating system, middleware inter-node connectivity, pre-selected storage, and Critical System Support. Additional assured availability configurations will be offered.

### **Business Continuity Support**

HP Business Continuity Support Services (BCS) is the industry's most comprehensive high availability support program with a four-hour call to restoration. It features maximum on-site presence and an outage prevention program tailored to each customer's specific needs. This service is designed to identify and prevent problems *before* they can have an impact on operations. If a problem should occur, the service can respond quickly to restore the system.

BCS features an assigned account team who is experienced in high-availability hardware, software and network support. These specialists include an account support manager, trained specialists from HP's phone-in response center, R&D engineers, support delivery specialists and HP senior management.

The team begins by conducting a full assessment of a company's operating environment to understand specific business needs. They identify precisely how to achieve the required level of systems availability as well as areas of potential risk to maximize critical systems performance and availability.

BCS also provides daily reviews to determine which patches the customer needs to change. The account team is committed to the customer's service objectives. It helps plan for changes and manages obstacles in the IT environment. The team has access to a regularly updated electronic profile of the customer's critical IT operations worldwide, 24 hours a day, seven days a week.

The emphasis of BCS is to keep the customer's high availability systems operating. It means rapid response and recovery before paperwork; proactive management over a reactive crisis mode; and reduced business costs and outages.

HP's business process model for high availability services delivers value to every aspect of a company's operation management. It identifies availability exposures and minimizes these risks. This comprehensive model encompasses all of the requisite steps for managing a high availability environment: configuration management, change management and user support. It includes service planning and management, infrastructure control and deployment, operational processes and problem management.

## **Other Complementary Services**

HP's high caliber Worldwide Customer Support Operation (WCSO) and Professional Services Organization (PSO) provide a range of proactive and reactive high availability services including consulting, maintenance, outsourcing and disaster recovery.

### **HP's Selective Outsourcing Approach**

HP has devoted much of its global support organization to providing outsourcing to key customers. These services utilize HP's proven IT management skills, tools and processes to perform daily systems and network management, administration and operation for selected portions of customers' IT environments.

HP is committed to the concept of "selective outsourcing." That is, HP works with its customers to help them craft an outsourcing solution that complements their existing IT resources. HP is able to manage just those portions of the infrastructure that the customer wants a vendor to manage.

This frees IT organizations to pursue other objectives that might be more strategic to the business or more relevant to the skills of the IT organization. It also helps IT organizations to accelerate the rate of technological transformation, particularly in complex, mission-critical environments.

HP's selective outsourcing services can be deployed to manage specified geographic units, business units or technology platforms (servers, networks, desktops or designated application environments). Every outsourcing agreement is documented with comprehensive performance objectives and metrics, including the availability requirements for critical applications and platforms. This helps to provide customers with a customized solution to meet all of their availability and performance requirements.

### **Professional Services**

HP's Professional Services Organization (PSO) provides high availability consulting based on HP's IT Service Management Framework. This framework incorporates the relationship between the users of the IT processes themselves, and the foundation elements (configuration and change management).

The PSO is available to provide expertise on infrastructure and processes necessary to deliver service and manage information, depending on the customer's specific high availability environment and needs. These consulting services range from the assessment of current process weaknesses through the design, planning and implementation of a complete IT infrastructure.

The PSO also offers a wide variety of educational courses that explore a broad range of issues that affect a high availability environment to specific high availability product management and configuration (for example, MC/ServiceGuard).

### **Disaster Recovery**

In the disaster recovery arena, HP, in conjunction with its partners, offers a full portfolio of business recovery services that include HP Backup Recovery Services and a telecommunications network. These services feature HP resources; an annual rehearsal for disaster recovery and review; and disaster site restoration.

The benefits of HP Business Recovery Services are obvious: it enables customers to maintain critical computer operations and ensure special needs are met, even in the throes of a disaster. HP also helps customers to refine their recovery plans and restore normal operations faster. HP has two backup recovery facilities in the U.S. and one in Canada. Coordinated HP account teams of highly trained and experienced staff are available to assist during rehearsals and actual recoveries.

## **Optimizing the High Availability Environment**

Hewlett-Packard Company not only offers the best comprehensive support services to meet availability needs across the enterprise, it also understands the concept of a support alliance. HP builds an alliance with each customer based on a commitment to focus on the customer's needs. In this alliance, HP works as a team with the customer to understand and improve that business' particular computing environment.

### **HP's Range, Reach and Reputation**

HP's support service delivers value to every aspect of a customer's operation management—from service planning and management to infrastructure control and deployment to operational processes to problem management and user support.

This service begins with HP's dedicated management. With over 23,000 service and support employees, 40 customer education centers and over 600 support offices in more than 120 countries, HP has the range, reach and reputation for comprehensive enterprise mission-critical support in open distributed computing.

Availability services can be a hard expense to justify for many companies. HP's long-standing reputation as a leader of quality service and products, however, offers companies considering a support agreement a solid track record for a return on their investment.

### **Multi-Vendor Support**

While a single point of contact offers many advantages in manageability as well as cost efficiency, it is imperative that the single vendor is equipped to deal with any and all complexities of the IT environment. HP has created strategic alliances with a variety of partners to provide customers with one-stop reporting and resolution.

As the first, and only one of three, worldwide **Microsoft** authorized support centers, HP is differentiated from the thousands of other Microsoft authorized service companies. HP engineers receive the highest level of training and technical backup available on Microsoft products.

By providing support for **Informix**, the second largest database vendor worldwide, HP is able to create a single point of contact for customers.

HP is the first worldwide support provider for **Netscape's** entire line of Internet and intranet software.

As an authorized service center for **Novell**, HP can provide a broad range of support services to help customers realize the full benefits of a Novell NetWare computing environment.

HP, in conjunction with the leading application provider SAP, has created global Competency Centers to address the joint complexities of high availability environments.

## **Case Studies**

### **Airlines Never Sleep**

At Ansett Airlines in Australia, key systems are required 24 hours every day of the year. "An airline never sleeps," said Robert Beswick, Ansett's Mid-Range System Manager. "Few industries are as adverse to risk as the airline industry. Some of the departments relying on our systems work triple shifts and use applications 24 hours a day. Other users are spread in a band of time zones. For some of our systems, the spread of users soon will become global."

Among the most critical systems for Ansett are scheduling and commercial support. "If scheduling went down, our operations would be impacted in less than an hour," said Beswick. "If such commercial applications as Yield Management and Core Financials were to go down, it would affect the work of many hundreds of users."

To ensure availability and minimize risk, Ansett uses HP's open systems and employs a number of high availability technologies from HP and its partners: MC/ServiceGuard clusters, mirrored disks, RAID disks, Oracle replication and Smart UPS systems.

"We place a very high value on proactive support and change management," said Beswick. "We avoid the collateral operational risks that can emerge from unsuccessful strategic alliances. That said, reactive support in an emergency is absolutely critical. We have to be satisfied that any support organization we deal with is able to meet that requirement before we award a contract."

HP has proven that it can provide both the reactive as well as the proactive support required.

### **Availability Keeps Distributor Ahead of Competition**

Merisel, Inc., a leading distributor of computer hardware and software products based in El Segundo, Calif., distributes a full line of 25,000 products to more than 45,000 resellers through North America. With reported 1996 sales of \$5.5 billion, Merisel has stringent service level requirements, with sub-second response time being a must. Such requirements place tremendous demands on the company's hardware technology to provide greater and greater throughput capabilities.

Merisel has a variety of different systems and technology: client/server Sybase systems, client/server Oracle systems, mainframe systems as well as server technology running in-house LAN-based applications. The company is currently converting from its legacy mainframe systems to HP 9000 T520s and K220s running an SAP R/3 environment. This conversion will allow Merisel to facilitate its North American strategy and improve the accuracy, accessibility and availability of information.

Merisel has been operating the HP/SAP environment in Canada since 1995. It has four HP 9000 Series K220 application servers that support over 400 users. In this environment, the response time averages between 0.6 and 0.7 seconds. Merisel's performance monitoring proves that 90% of all transactions are



less than one second; and 99.5% are less than 10 seconds. These transactions are based on an average of 450,000 calls per day. The U.S. operation anticipates four times this number of dialogs.

Mary West, Merisel Vice President of Information Systems, commented, “HP and SAP certainly have given us a lot of advantages. Between the IS organization and the business community there is a service level expectation. The expectation is that you must be able to process orders at all times and *never* have any downtime.”

She confided, “We know when our competitors’ systems are down because our order volume increases dramatically—and vice versa! It is absolutely critical that we have back-up order entry and high availability because we literally cannot afford downtime.”

According to West, Merisel uses numerous tools from HP to proactively manage its systems (Network Node Manager, GlancePlus and PerfView). “We’re very oriented to preventive action, but if we have a problem, we can recover quickly with our support processes and products like MC/ServiceGuard and EMC Symmetrix. HP is able to give us the service

that we need.”

## **High Availability Gives a Kick to Hosiery Company**

The consequences of system downtime for a manufacturing company like Sara Lee Hosiery, a market leader primarily in women’s leg coverings, may not be a crisis. It can, however, have a costly and significant impact to the overall business operations.

This Winston-Salem, NC-based company, whose brands include Hanes and L’Eggs panty hose, is currently in the throes of migrating 80 of its 120 legacy mainframe systems to SAP manufacturing and logistics programs on HP servers. When completed, SAP will be running the bulk of Sara Lee Hosiery’s business from manufacturing facilities and human resources to payroll, sales, marketing, data collection and financials.

John Zaski, Director of Information Technology for Sara Lee Hosiery, explained, “We have a lot of facilities that run multiple shifts, so sometimes we need true 24-hour availability. Based on that, our intent is to have a highly available system because any downtime is very difficult for the business to absorb. By moving to the SAP environment, we are putting a tremendous burden on HP to ensure that our systems are available on a real time basis *all* the time.”

No stranger to HP’s reliability, Sara Lee Hosiery has had HP equipment in its distribution centers for years. “We’ve been very, very pleased with the reliability of the HP equipment without any of its high availability solutions,” he noted. “We’re confident that the equipment is sound. But with the added risk of implementing a mission-critical application like SAP, we are looking to implement solutions like MC/ServiceGuard as additional insurance.”

As Sara Lee Hosiery’s IT moves more mission-critical applications to the HP environment, the company is addressing service level agreements so its users will have clear expectations. In an effort to reduce the amount of risk, it also is looking into upgrading its support contract.

“We are a lean and mean IT department. We’re constantly pushed to do things in the most productive manner as possible. Our intent is to forestall major system availability issues by having the right proactive processes in place. HP’s increased level of support will help mitigate the risk and ensure that we really do have an environment that is as risk-free as possible in terms of system availability.”

## **Conclusion: A Total Service Solution**

This is a time in IT technology when there is increasing vulnerability and more potential areas of exposure and failure. Businesses must determine their optimum service strategy and availability levels across their IT environment to ensure maximum uptime for mission-critical applications. By understanding that availability is as much a business decision as it is a technological one, businesses will realize that they cannot afford to operate without high availability services.

Reliable technology alone cannot substitute for strong support services and integrated management. All three must come together as a total solution.

Because each environment is unique, there is no single right combination to achieve high availability -- except with Hewlett-Packard Company. As one of the world's most respected computer companies, HP has been lauded consistently by *Fortune Magazine* for its high level of quality services, financial stability and value. HP is recognized worldwide not only for its superior products, but also for its leadership and technical strength. Consider HP's track record for service and support as ranked by these publications:

The 1996 *DataPro* Survey placed HP Number 1 worldwide in overall support satisfaction for UNIX operating systems;

In 1995, *Software Magazine* found HP to have the best service and support customer satisfaction among system/platform providers; and

That same year, *Datamation* put HP's overall service and support among the best of all U.S. IT vendors.

HP partners with its customers as a trusted advisor to help make technology choices. Then HP can help customers capitalize on the value of their technology investments in availability. Finally, HP can help maximize customers' current investment in people, infrastructure and partnerships.

As a leading provider of services, HP works with its customers to design and deliver complete end-to-end solutions that include hardware, software, middleware, networking, third-party products, technical consulting and other services that meet mission-critical high availability needs.

HP knows how to optimize the high availability environment. HP is available to help companies plan, then deploy the latest technology trends into their IT environment. Drawing on its own experience and the strength of its partners', HP has the reach, the reputation and the resources to give its customers every competitive advantage.

Reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under the copyright laws.

Printed  
Paper #4130

in

USA

GK0697

Technical information in this document is subject to change without notice.

© Copyright Hewlett-Packard Co. 1997  
All Rights Reserved. Reproduction, adaptation, or translation without prior written  
permission is prohibited except as allowed under the copyright laws.

Printed in USA GK0697  
Paper #4130