COMPAQ
White Paper

October 1998

Prepared by *Custom***Systems**
Enterprise High Availability
Segment

Compaq Computer Corporation

Contents

Introduction 3

Determining Availability Requirements 3

What is the Cost of Downtime? 4

Recovery Point and Recovery Time 5

What Causes Downtime? 6

Vulnerable Technologies 7

Availability Technologies 8

Architecting High-Availability Systems 9
-- Deployment Options 10
-- Service and Support Options 10

Putting it all Together 12

List of Sales Offices 14

Architecting and Deploying High-Availability Solutions

Abstract: The demand for high availability is growing. Long required for mission-critical applications in industries such as finance, process manufacturing, and telecommunications, high availability today is fast becoming a requirement in many other industries as well.

This white paper provides an overview of the combination of factors that defines high-availability computing requirements. It describes methodologies that can assist you in architecting and deploying the right level of availability across your information-technology environment. And it describes how Compaq can assist organizations of any size in achieving their high-availability computing goals.

Help us improve our technical communication. Let us know what you think about the technical information in this document. Your feedback is valuable and will help us structure future communications. Please send your comments to: customsystems@digital.com

Notice

Architecting and Deploying High-Availability Solutions
White Paper prepared by CustomSystems Enterprise High Availability Segment

First Edition (October 1998)

Introduction

When the systems that run your organization[1] are down the costs can be devastating. Lost opportunities. Lost revenues. Failure-to-perform fees. Non-compliance penalties. Plus stranded fixed costs you have to keep on paying whether your people are productive or not.

Potentially more damaging is a nearly incalculable loss: the loss of good will. Customers, partners, and suppliers affected by system shutdowns perceive your organization as poorly run and ill-suited to meeting their needs.

Thus, the starting point for any discussion of availability has to be the cost of information system downtime to your organization. The higher the cost of downtime, the more robust the availability environment needs to be. And the more successful your environment is in delivering the level of availability required by the organization, the faster the return on your investment.

The purpose of this white paper is to provide an overview of the factors that – taken together – define your availability needs and how Compaq Computer Corporation can meet them.

Definitions

Before proceeding with a discussion of high availability, it is helpful to define a few terms.

Availability: The ratio of the total time a functional unit is capable of being used during a given interval to the length of that interval. It is the proportion of time a system is productive which implies performance.

Mission Critical: A term applied to information systems upon which the success of an organization depends and the loss of which results in unacceptable functional or financial harm.

Mean Time Between Failure(MTBF): a statistically derived length of time a user may reasonably expect a component, device, or system to work between two incapacitating failures.

Reliability: A measure of how dependable a system is once you actually use it. Reliability can also be considered the sum of availability and data integrity.

Determining Availability Requirements

Determining an organization's availability requirements and architecting a system to meet them is a multi-step process.

1. Determine the cost of downtime. (*Page 4*)

2. Understand recovery in terms of point and time. That is, when is recovery necessary in the system's operations and how long a time exists between the point of failure and recovery. (*Page 5*)

3. Focus on the events that can have a negative impact on the ability to keep an application -- and an organization -- up and running. Understanding these events is essential for developing the right high-availability environment for your organization. (*Page 6*)

4. Understand the vulnerability of various systems with respect to the above. (*Page 7*)

5. Once you understand the events that lead to downtime, determine which technology areas you will need to focus on to achieve the level of availability required. (*Page 8*)

6. When all these factors are understood, architect and deploy the availability environment. (*Page 9*)

[1] As used in this white paper, organization means any entity that requires computing technology to achieve its goals and conduct its operations. Examples include businesses, departments of businesses, academic institutions, research facilities, or military units.

Many of our customers find it more cost-effective to engage Compaq for architectural and deployment activities.

1. What is the Cost of Downtime?

You need your Information System to survive in a world fraught with risk. A world where off-chance and down-right failure can bring your operations grinding to a halt.

What happens to your organization when the system goes down? The range of answers runs from "inconvenient" to "catastrophic."

If your answer is closer to "catastrophic" than to "inconvenient" then you should read on. If your answer is closer to "inconvenient" then perhaps you should read on to see if things are really as rosy as they seem. In fact, most organizations underestimate -- or have not calculated -- the impact of downtime on their business. The Gartner Group (1998) studied downtime costs for a variety of industries. The table below summarizes the findings.

Industry	**Application**	**Average Cost per Hour of Downtime**
Financial	Brokerage Operations	$6,500,000
Financial	Credit Card Sales	$2,600,000
Media	Pay-per-View	$1,150,000
Retail	Home Shopping (TV)	$ 113,000
Retail	Catalog Sales	$ 90,000
Transportation	Airline Reservations	$ 89,500

In order to measure the impact of downtime, let's ask a basic question that helps quantify the level of availability you might need.

Who and what gets hurt when a system goes down?

Processes: Vital business processes may be interrupted, lost, corrupted, or even changed. Such processes might include order management, inventory management, financial reporting, transactions, manufacturing, human resources, life-sustaining medical systems, extended 911 identification, ATM operations, and more.

Programs: Both long- and short-term revenue can be affected. Key employee or customer activities might be missed or lost.

Business itself: In this age of electronic commerce, if prospects or customers can't access your site, chances are they'll never come back. And the chances are good that they'll end up with your competitor. Customers lost forever!

People: Lives can be lost; employee benefits missed with adverse impact on morale; governmental program problems might harm citizenry; and even battles can be adversely affected if vital information is lost, corrupted, or late.

Projects: Hundreds of thousands of person-hours of work can be lost, deadlines can be missed, change-orders might be skipped with devastating results.

Operations: Those who manage the daily activities of an organization may find themselves without needed data, with lost information, with standard activities lost, or with key reports missing or corrupted.

Loss can be measured in more than money. But if money is the measure then the figures can be astounding. In a recent study, the Standish Group (1998) reports that costs of downtime typically range from $1,000 to $27,000 per minute. What's more, they report that in some cases, the cost of downtime for a single incident has exceeded $10,000,000. And if you consider the estimates of the Gartner Group as noted, the costs can be in the *Billions!* Think about what downtime means to your organization.

2. Recovery Point and Recovery Time

High availability means different things to different people. At the high end it is called "continuous availability" or "nonstop computing" and has come to mean something on the order of 99.999% uptime, some five minutes a year of downtime. Pretty impressive. But what is your definition of high availability? Perhaps you don't need "five-nines" but you'd like to come as close as you can. Your requirement may not be for continuous computing 24 hours per day, 365 days per year, but you may require that when your system is in operation it *cannot* go down. An airborne surveillance and target acquisition system might be in operation for only eight hours over the forward edge of battle but it better be available every second that it's there. Or a retail operation that does 90% of its business during a holiday season had better not go down during those few weeks or months. Each type of availability may demand very different requirements.

The first thing to keep in mind, though, is that defining availability depends on your needs in terms of *Recovery Point* and *Recovery Time*.

While the inherent reliability of Information Systems has been increasing, things still do happen that cause applications to stop. Disaster Recovery specialists tend to examine the impact possibilities in terms of Recovery Point -- the amount of "acceptable" loss -- and Recovery Time -- the amount of time needed to get back in operation. Recovery Point is most important in *data-centric* operations where the loss of data is unacceptable. Recovery Time is most important in *transaction-centric* operations where realtime continuity is key.

Do you need fast recovery, or recovery to the exact state prior to the failure… or both? What is the impact on your operations measured by a Recovery Point standard? If you don't resume processing right where you left off will it be inconvenient? Damaging? Catastrophic? What is the most effective and efficient method to use to recover the information? What is the impact on business measured in Recovery Time? If you don't resume processing within a second will it be inconvenient? Damaging? Catastrophic?

Thus the recovery strategy you use depends on this assessment of Recovery Point and Recovery Time. The diagram below displays four availability options measured in those terms.

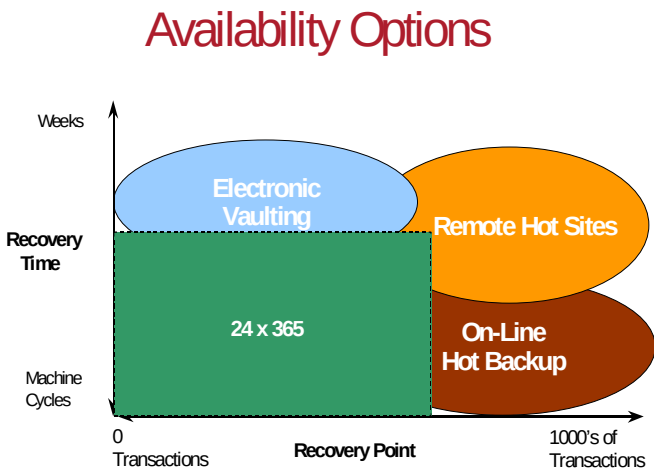

Remote Hot Sites (functional locations geographically distant from the primary operations center) are an option if the Recovery Point and Recovery Time for an application are not very critical. An example might be a billing application where the monthly statements could be delayed in mailing with minimum impact on a business

Electronic Vaulting (method of electronically storing, managing, and protecting data in a computer "vault" which is located off-site in a physically secure location) is an option if Recovery Point is more important than Recovery Time; if, for instance, an indeterminate amount of data cannot be lost or historical data needs to be available online for reference. An example might be an inventory application where the most current transactions are recoverable by other means and the application can be restarted where it left off using the historical data as a basis for inventory status. This is a good example of a *data-centric* operation.

On-line Hot Backup (data backup that is conducted while the system is in full operation) is necessary if the Recovery Time is more critical than the Recovery Point A good example is an on-line traffic or production control system where history is not as important as the current state of the situation. In air-traffic control, where the planes were five minutes ago is not as critical as where they are now, because in five minutes they may have moved 50 miles each, but in what directions? This is a classic example of a *transaction-centric* operation.

24 x 365 (continuous availability) is the only viable option where both the Recovery Point and Recovery Time are critical for an application.

Using the criteria of Recovery Point and Recovery Time, which state of availability is right for your organization?

3. What Causes Downtime?

After looking at your information systems, the user community, and the cost of downtime, you can determine the level of availability you need. Now it is time to focus on the events that can have a negative impact on your ability to keep an application – and an organization – up and running.

Component faults due to hardware, software, or interoperability issues. While the industry has come a long way in reducing Mean Time Between Failure (MTBF) rates for individual hardware, packaging , and mechanical components, the interdependent nature of today's multivendor and networked solutions makes them vulnerable to hardware, software, and network interoperability problems.

Administrative intervention. Just because it's planned downtime doesn't mean it's not downtime. Management tasks like system maintenance, database backups, index builds, table reorganizations, cache changes, application/operating system updates, system re-configuration, and a physical move may require that a system be brought down. Or the intervention itself may cause a failure.

Building-level incidents. In addition to system problems, disasters affecting a site or building, such as fire, power loss, or flooding, can interrupt service by damaging systems, robbing them of power, or preventing access to them.

Metropolitan area disaster. Disasters, such as floods, fire, and blackouts, can also affect whole cities, impacting systems located throughout the metropolitan area.

Regional events. Computing can also be interrupted by disasters that affect systems across an even a larger region. Hurricanes, earthquakes, or geopolitical disruptions can cause outages over hundreds of square miles.

Do you know the probability of each of these events affecting your operation? Do you know what will happen to your applications, particularly those in the “24 x 365” zone in each of these cases? Do you know it can cost less than the alternative to minimize the negative impact that could occur? Understanding these factors is crucial to determining the level of availability required by your organization.

4. Vulnerable Technologies

Once you've looked at risk and cost, recovery, and failure events, you then need to turn your attention to the following technology areas.

Hardware: Is the concern uninterruptible power supply, memory integrity, I/O, or processor failure?

Operating system: Some operating systems tend to be more stable than others, nevertheless, problems do occur that can affect an operating system and cause it to corrupt a process. Once the operating system is in failure everything above it such as databases, applications, processes, and more can be lost.

Storage: Is your business data-centric, transaction-centric, or both? If it is data-centric you do not want to lose what is in your storage system; you don't want to lose access to that vital data; and, perhaps worst of all, what happens if that data becomes corrupted?

Database: Whether your business is data-centric or transaction-centric the loss or corruption of a database can spell disaster. This brings up the question about the need to replicate data.

Network: Whether it is backing up data over wide-area connections, the high-speed connections needed for transaction processing, or the continuous availability of information already gathered, a robust network infrastructure is a foundation on which to build truly high-availability environments.

Management: How much does it cost to manage this whole system? What happens if you are no longer able to deal with intrusions? How can you easily, or at least with minimum disruption, upgrade, expand, or tune your system to ensure that your enterprise operates at peak profitability?

Application: Is the application always there for you? Can you upgrade it without taking your system down? Can the application take advantage of the underlying high-availability elements? Interdependencies are also a concern. Storage failure, for example, can corrupt a database leading to operating system hang-up.

Looking at the Causes and Vulnerable Technologies Together

Once you have gained an understanding of which causes of downtime are of most concern and which technologies are most at risk, you can begin to analyze your total environment and architect the high-availability system that is right for you. This can be done in-house or with the assistance of a third party. Compaq has considerable experience in this area and stands ready to assist you.

Availability Matrix

The matrix below is a simple yet effective way to quantify your availability requirements. The top row displays the locations where failure events occur and the leftmost column lists the technology areas that must be addressed to minimize the loss of availability.

	Component	*Administrative*	*Building*	*Metropolitan*	*Regional*
Hardware					
Operating System					
Storage					
Database					
Network					
Management					
Applications					

Recovery Point and Recovery Time are implicit to the Availability Matrix because there are products and capabilities to help with the analysis, and to establish and implement the practices necessary to assure that

the high-availability goals, once met, are maintained. These capabilities are very often described as *Availability Reviews*, *Disaster Recovery Services,* or *Business Impact Analyses* and are primarily consulting services to help you understand and/or implement a high-availability environment. In addition, at each intersection in the matrix there are specific products designed to address the level of availability you require.

What is most critical to know is that these technologies are available and that they play key roles in implementing high-availability solutions. Your choice of an appropriate partner that has "been there, done that" can make the difference in implementing the solution.

5. Availability Technologies

As we have seen thus far, there are a number of important elements to consider in architecting a high-availability solution: the cost of downtime; trade-offs between recovery point and time; potential negative events. Only after these elements have been qualified and quantified can the next step be taken: choosing a technical strategy for achieving the level of availability required by a particular environment. The next few paragraphs provide a quick overview of some availability technology options.

Power supply
Implementing solutions for any of the situations described requires some element of redundancy. At the component level the most basic element is the power supply to the system. Uninterruptible Power Supplies (UPS) are a common tool to deal with the possibility of power outages. Multiple power sources, dual battery feeds, and power connections should also be considered in planning for component failures.

Storage
Once the power supply is assured, the next element to consider in a particular environment might be storage. The ability to maintain redundancy of the data and applications is key to any recovery situation. After all, physical storage devices are electromechanical and that in itself makes them more failure-prone than other elements in the environment.

Database
Database replication is a combination of hardware and software implementations specifically focused at protecting the data in your database. It may, or may not, require redundant physical storage. It provides application-transparent database backup and may also provide for use of multiple storage media.

Backup and Restore is typically a capability that integrates a number of database and storage technologies and media with the intent of providing both database backup and fast restoration capabilities.

Among other things, a TP monitor manages system resources, replication, load balancing, and failover capabilities. TP monitors also enhance security and assure that transactions are completed before confirming their completion by using a two-phase commit process.

Processors, operating systems, and interconnects
Fault-tolerant (FT) is a term used to describe the ability of a server or solution to tolerate failures. FT includes both Hardware and Software components.

Hardware fault-tolerant: The ideal fault-tolerant solution is where the recovery point is instantaneous through the use of alternate data paths to deliver continuous availability. To deliver continuous availability, the components of the system need to operate in a fast-fail mode, that is, to *identify* the problem quickly, *isolate* it from the integrity of operation of the total solution, and *recover* by using alternate paths. (Note that this ability to route around failures allows the user to select a lower level of support response time as mean-time-to-repair (MTTR) is not a critical component in ensuring the availability of the solution.)
Software fault-tolerant: The ability to recover from software failures. The main cause of server failures in today's architectures is software. The ability to tolerate software failures is a key requirement for the delivery of continuous availability.

Clusters should be able to overcome the administrative failures described above. Specifically included is the fact that the application should be available while items like maintenance, hardware upgrades, software upgrades, and database updates are going on. Then there is the question of how many nodes a cluster can support as well as where the nodes physically reside.

Fibre Channel, SCSI 2, UltraSCSI, and ServerNet technologies improve storage throughput and allow storage devices to be separated by greater distances while providing for the efficient, protected movement of data between clusters and storage subsystems.

Channel Extenders do for CPUs what Fibre Channel and Ultra SCSI do for storage -- they provide for greater separation enhancing tolerance for disasters.

Wide-Area and Global Clusters are separated by greater distances. The 500-miles-and-beyond capabilities that exist today easily qualify them for dealing with Regional events.

Networks
By their very nature, networks are capable of spanning enough distance. When you look at the redundancy and virtual capabilities that are available, they qualify as Regional solutions in the Availability Matrix. A network in which a two-phase commit process assures a transaction is not complete until the remote node confirms it certainly makes this high availability.

The Compaq Technology Brief entitled *Eliminating Single Points of Failure and Enabling Rapid Recovery in Server Subsystems* (ECG0094/1197) provides more detail on increasing availability in basic components including processor, expansion, memory, primary and secondary storage, power, cooling, and management.

6. Architecting High-Availability Systems

Analyzing and defining the requirements for a high-availability environment that meets your unique needs is not a trivial process. Add to that the challenges of solution deployment and the task can seem overwhelming because -- in addition to issues surrounding availability -- other factors must be considered, such as:

Performance and scalability: Can the solution deliver the performance required? Can it scale appropriately to meet the size of the application?

Affordability: Is the cost of purchasing, implementing, or managing the solution prohibitive or consistent with the operational loss you wish to prevent?

Manageability: Does the organization have the knowledge and resources necessary to manage the solution over the long run?

Service and support: What level of service and support is available? Is the infrastructure in place to support the solution reliably and affordably – wherever it is located?

Turn to Compaq for the Answers
Compaq unquestionably has the broadest range of availability products and solutions on the market today. We can provide better answers *anywhere* on the Availability Matrix with open, reliable solutions and technologies from hand-held PCs to the most scalable and available fault-tolerant and clusters systems.

The professionals in the Compaq Enterprise Computing division helped *invent* the concepts of disaster tolerance, fault tolerance, and continuous availability almost two decades ago and have been delivering and supporting systems in the most demanding mission-critical environments ever since. The combined development, delivery, and support experience of Compaq, DIGITAL, and Tandem personnel is unsurpassed in the marketplace.

Consider these compelling facts: Among customers with the highest levels of availability requirements – those with an hourly cost of outage of $100,000/hour or more—the choice is Compaq by a factor of nearly three times that of our nearest competitor.

Our line of availability platforms is recognized[2] as the most robust set of offerings in the industry. And through Compaq integration programs, we are extending the high levels of availability found in these high-end servers to the Microsoft Windows NT environment. Compaq is the leader and major driving force in the acceleration of Windows NT Server as an enterprise solutions platform.

Deployment Options

As we have seen, the process of architecting a high-availability environment involves the understanding and analysis of multiple factors, many of which may be unique to your organization. While the majority of a solution might be based on Compaq's wide range of off-the-shelf products, unique factors require unique deployment solutions attained through some level of tailoring or customization. The *Custom***Systems** organization at Compaq specializes in tailoring off-the-shelf products to meet complex needs.

Pre-packaged Systems

Most vendors ship systems with little or no up-front testing. But Compaq is different. Our Ready to Go clusters, for instance, are factory-configured and tested solutions. They eliminate the need for the consolidating, compiling, testing, and tuning steps typically associated with building and debugging a complex system from scratch. As a result, cluster solutions are online faster and ready to run your choice of applications.

Tailored Solutions

High-availability systems interoperate within increasingly complex software, hardware, storage, and networking environments. Your requirements may go beyond a standard solution. In that case, the proven engineering and testing principles employed by Compaq deliver a reassuring measure of certainty that your systems will work reliably in the real world and meet your definition of mission critical.

For example, we can help you determine how your application will perform on large-scale configurations. Working together, we can do a one-time assembly and test in our labs to validate the entire system solution. You'll have the answers you need prior to committing valuable time and precious resources. What's more, if you're developing or planning to deploy a specialized application, we can give you a detailed look at all or part of your development efforts. This enables you to recognize any potential problems and compare design alternatives.

A key advantage to you is Compaq's ability to integrate the complete system in our factories. This ensures that you receive a fully functioning, characterized, and tested system that is ready to be installed and turned on as soon as it arrives at your door. Factory integration reduces your risk and accelerates your implementation speed. What's more, the disruption to your operation that might be caused by extended installation, testing, and pre-functioning repair is greatly reduced.

The result for you is reduced risk and a high degree of confidence that the combination of applications and enterprise computing platforms you plan to deploy will work together to meet your mission-critical requirements.

Service and Support Options

Compaq also offers our customers one of the most powerful IT service and support organizations in the world -- including 27,000 IT experts in 114 countries and a network of over 30,000 highly skilled resellers. We offer a suite of proactive and reactive services that will keep your Windows NT, OpenVMS, UNIX, Intel, Alpha, NonStop Himalaya, Integrity/XC, and Integrity S-series environments highly available and operating at peak performance. Our service capabilities include

---

[2] *The Real McCoys: Enterprise Clustering*. The Standish Group International, Inc. 1998
Product Analysis, MCSC: Can this Wolf Lead the Pack? DataQuest, April 27, 1998

- Total support for servers, network operating system, applications, switching components, and PCs
- Multivendor support for a diverse range of products including networking equipment, applications, and peripherals
- Unrivaled expertise in Microsoft support: A Microsoft Solution Provider and Authorized Support Center, we have the world's largest concentration of Microsoft certified engineers in the world.
- Unequaled warranty terms and options

With more than three decades of experience, Compaq offers a full suite of processes, methods, and tools designed to create high-availability environments that meet your mission-critical requirements. With our broad range of technologies, we can help you define your availability requirements, target the best platform or blend of platforms for deployment, and deliver an integrated, enterprise-wide system. The people who know availability best can help you assess how to improve the availability of your existing environment or to design an infrastructure from the ground up. Let's take a look at how we do this.

Compaq Availability Review
Identify risks to availability in your environment
If your organization depends on the availability of your technology environment, downtime can result in loss of revenue, penalties and fees, and damage to your organization's image. Any reduction in availability affects productivity. Inevitably, time and change will threaten the stability of any technology environment and in turn, your operations. A Compaq Availability Review can help you guard against downtime exposure and improve your availability.

The Compaq Availability Review offers unique value through a combination of expertise, innovative tools, and experience that produces a proactive solution for high-availability needs. Compaq availability consultants train intensively on the factors affecting uptime, solutions to improve them, and the tools and methods employed during an Availability Review.

In partnership with your team, our consultants compare the current state of your information technology installation with your business and availability goals and help you plot the course to achieve your goals. This process entails

- Investigating all factors of your IT installation related to availability
- Calculating the cost of downtime for critical environments
- Simulating multiple improvement scenarios to reach the maximum return on investment
- Recommending actions that can minimize risk and improve availability

Compaq can analyze any vendor's technology, in any configuration, in over 100 countries around the world. And we can tailor your analysis report to suggest improvements that will satisfy specific and differing availability requirements across your organization. In the process, we

- Gain a full understanding of your unique business, operational, and availability risks
- Accurately examine downtime consequences across complex, dynamic, and global computing environments
- Identify and prioritize availability improvements to minimize risk, optimize uptime, and gain control of your IT assets
- Simulate the addition of new systems to determine the impact on your environment

An Availability Review is the best first step toward optimal uptime and its financial rewards.

Availability Partnership
Availability Partnership is an annual service that provides a fully customized plan to maintain and improve your availability levels over time and throughout changes.

Availability Partnership offers these services:

- A partnership service plan based on a complete assessment of your business and service needs
- Availability monitoring and reporting, including trend analysis
- Proactive configuration management to identify and prevent problems
- Change impact analysis and consulting
- Electronic Site Management guide, a web-based tool for ongoing management of your environment

Mission-Critical Services
In addition to Availability Review and Availability Partnership, Compaq offers an array of proactive and reactive hardware and software support services designed to achieve your target uptime level.

Focused on preventing the causes of system-, network-, or environment-based outages, Compaq Business Critical Services are available for Windows NT running on multivendor Intel or Alpha systems, DIGITAL UNIX, NonStop Himalaya, Integrity/XC, Integrity S-series, and OpenVMS. And for all of these platforms, we offer an uptime guarantee – our commitment to sharing the risk with you.

A formal contract between you and Compaq, this guarantee commits Compaq to keeping your business critical OpenVMS and DIGITAL UNIX systems operational 99.99% of the time. Compaq backs the uptime guarantee where it counts the most — on the bottom line: If we don't deliver on the promised availability level, you don't pay the full service price. Best of all, entering into an Uptime Guarantee costs you nothing extra. Once you've met the required service and serviceability levels, the Uptime Guarantee is yours at no additional charge.

Business Continuity Solutions
This complete portfolio of services is designed specifically to help you deal with the environmental and human-error events that can wreak havoc on the IT environment your business depends on for its day-to-day operations—sometimes for its survival. Our Business Continuity Solutions start with comprehensive planning. Information gathered in the planning phase helps you identify and prioritize the areas of your business that would require the fastest recovery. You then select the best method of implementing your well-defined recovery plan from a range of available protection and recovery capabilities, including

- Contingency Planning to develop a strategy for fast recovery with minimum disruption to your operations
- Electronic vaulting for storing critical data at a secure site
- Disaster-tolerant services that link two data centers into a single resilient solution
- Hot-site Service with access to a fully equipped data center, allowing you to conduct business as usual during the recovery period
- Mobile recovery provides recovery configurations at a site you select
- Recover-All equipment replacement services offer a unique combination of services and insurance to cover you for damage or loss, plus provide FAST equipment repair/replacement
- A guaranteed, 24-hour equipment replacement service that can provide a complete duplicate configuration of your environment overnight, available in select locations

Putting it all Together

An approach to architecting your high-availability solution

Moving from analysis to architecting to implementation is an iterative process of eight steps. It is vital to consider each step in turn and recognize that once your reach an end point, long-term management then becomes the key to success. Eventually, as your organization and its requirements evolve, you will again need to begin the cycle and move to whatever the next implementation might be for you.

Regardless of where your organization is in this cycle, Compaq can help you address the requirements at that point and then move on. Through many years of experience, Compaq has developed the capability to assist organizations in creating and implementing this process. The following graphic and description introduce the eight steps in achieving the level of high availability most appropriate for your needs at any point in time.

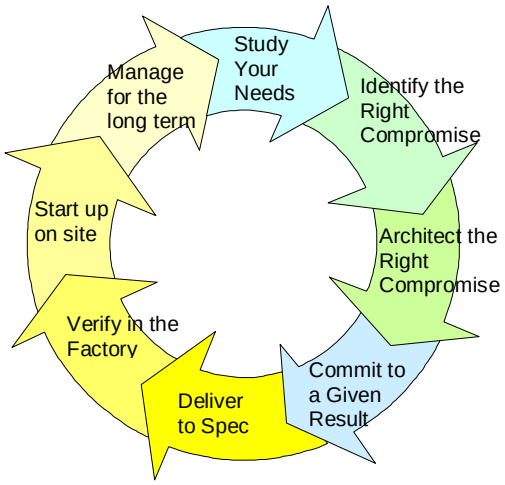

Study your needs
As we have discussed in this paper, architecting a high-availability environment requires a methodical approach beginning with an understanding of your needs. The definition of your needs includes the cost of downtime to your organization and your requirements in terms of Recovery Point and Recovery Time. Consider also the ramifications of managing your high-availability environment in a multi-year time frame.

Identify the right compromise
What is the opportunity cost of downtime versus the price of the solution over the life of the system? Should you spend less now for a short term fix or make the major investment that will carry you well into the future?

Architect the right compromise
Once a balance has been struck among the issues described above, the next task is to turn that compromise into a design with the appropriate weight given to the individual availability technologies.

Commit to a given result
With design in hand, an organization must commit itself to carrying through on the plan whose end result is nothing less than the level of availability deemed necessary. This requires some degree of systems engineering – solutions modeling, characterization, proof of concept.

Deliver to spec
At this point, the vendor or vendors supplying the components of the availability solution must be equally committed to the end result. A vendor with strong internal technical project management and quality control is essential in this endeavor.

Verify in the factory
Don't wait until the solution is delivered to your site and installed to see if it works. A competent solutions provider will have the capability to factory pre-stage, integrate, and test the complete solution.

Start up on site
If the solution has been factory integrated and tested in a tightly controlled technical environment the on-site start-up should be a low-risk and non-disruptive exercise that accelerates the pace of implementation.

Manage for the long term
After implementation, the challenge is to maintain the level of availability originally architected. This is accomplished through a program of consistent internal controls and long-term relationships with support and service providers.

Conclusion

When it comes to "putting it all together" Compaq is unique in the full range of performance-driven, affordable, high-availability computing products and services offered. We're in business to architect and deploy complete, integrated solutions worldwide and we're ready to start today. Give us a call or visit us on the Web. Consult the following list to find the number of the nearest sales office or Web site in your region.

In the Americas

	Country	*City/Region*	
Argentina:	+54	1	331 7500
Canada:			800 344 4825
Chile:	+56	2	671 8161
Mexico:	+52	5	448 1713
United States:			800 344 4825
All other countries:			954 360 6470

On the web at www.compaq.com/customsystems

In Europe, the Middle East & Africa

	Country	*City/Region*	
Belgium:	+32	2	729 71 49
Denmark:	+45	45	17 23 16
France:	+33	4	50 09 41 14
Germany:	+49	89	95 91 16 85
Italy:	+39	2	66 18 84 43
Netherlands:	+31	30	283 24 54
Spain:	+34	91	583 40 55
Switzerland:	+41	22	709 51 14
United Kingdom:	+44	0118	920 46 20
All other countries:	+33	4	50 09 41 14

On the web at www.europe.digital.com/customsystems

In the Asia/Pacific region

	Country	*City/Region*	
Australia:	+61	2	9561 6333
China:	+86	10	6849 2883 ext. 7165
Hong Kong:	+852		2805 3111
India:	+91	80	337 4785
Japan:	+81	3	5349 7369
Korea:	+82	2	3771 2944
Malaysia:	+60	3	751 3235
Philippines:	+63	2	840 6349
Singapore:	+65		290 7600
Taiwan:	+886	2	776 8902
Thailand:	+66	2	654 0788 ext. 3323

On the web at www.asia-pacific.digital.com/customsystems
(Australia) www.digital.com.au