

**FOR PUBLIC
RELEASE**

CHAPTER 1

Basic High Availability Concepts

*T*his book begins by taking an elementary look at high availability (HA) computing and how it is implemented through enterprise-level cluster solutions. We start in this chapter with some of the basic concepts of HA. Here's what we'll cover:

- What Is High Availability?
- High Availability as a Business Requirement
- What Are the Measures of High Availability?
- Understanding the Obstacles to High Availability
- Preparing Your Organization for High Availability
- The Starting Point for a Highly Available System
- Moving to High Availability

Basic High Availability Concepts

Later chapters will explore the implementation of HA in clusters, then describe HP's HA products in more detail. A separate chapter is devoted to concrete examples of business solutions that use HA.

What Is High Availability?

Before exploring the implications of HA in computer systems, we need to define some terms. What do we mean by phrases like “availability,” “high availability,” and “high availability computing?”

Available

The term **available** describes a system that provides a specific level of service as needed. This idea of availability is part of everyday thinking. In computing, availability is generally understood as the period of time when services are available (for instance, 16 hours a day, six days a week) or as the time required for the system to respond to users (for example, under one-second response time). Any loss of service, whether planned or unplanned, is known as an **outage**. **Downtime** is the duration of an outage measured in units of time (e.g., minutes or hours).

What Is High Availability?

Highly Available



Figure 1.1 *Highly Available Services: Electricity*

Highly available characterizes a system that is designed to avoid the loss of service by reducing or managing failures as well as minimizing planned downtime for the system. We expect a service to be *highly* available when life, health, and well-being, including the economic well-being of a company, depend on it.

For example, we expect electrical service to be highly available (Figure 1.1). All but the smallest, shortest outages are unacceptable, since we have geared our lives to depend on electricity for refrigeration, heating, and lighting, in addition to less important daily needs.

Basic High Availability Concepts

Even the most highly available services occasionally go out, as anyone who has experienced a blackout or brownout in a large city can attest (Figure 1.2). But in these cases, we expect to see an effort to restore service at once. When a failure occurs, we expect the electric company to be on the road fixing the problem as soon as possible.



Figure 1.2 Service Outage

What Is High Availability?

Highly Available Computing

In many businesses, the availability of computers has become just as important as the availability of electric power itself. **Highly available computing** uses computer systems which are designed and managed to operate with only a small amount of planned and unplanned downtime.

Note that *highly available* is not absolute. The needs of different businesses for HA are quite diverse. International businesses, companies running multiple shifts, and many Internet sites may require user access to databases around the clock. Financial institutions must be able to transfer funds at any time of night or day, seven days a week. On the other hand, some retail businesses may require the computer to be available only 18 hours a day, but during those 18 hours, they may require sub-second response time for transaction processing.

Service Levels

The **service level** of a system is the degree of service the system will provide to its users. Often, the service level is spelled out in a document known as the service level agreement (SLA). The service levels your business requires determine the kinds of applications you develop, and HA systems provide the hardware and software framework in which these applications can work effectively to provide the needed level of service. High availability implies a service level in which both *planned* and *unplanned* computer outages do not exceed a small stated value.

Basic High Availability Concepts

Continuous Availability

Continuous availability means non-stop service, that is, there are no planned or unplanned outages at all. This is a much more ambitious goal than HA, because there can be no lapse in service. In effect, continuous availability is an ideal state rather than a characteristic of any real-world system.

This term is sometimes used to indicate a very high level of availability in which only a very small known quantity of downtime is acceptable. Note that HA does *not* imply continuous availability.

Fault Tolerance

Fault tolerance is not a degree of availability so much as a method for achieving very high levels of availability. A fault-tolerant system is characterized by redundancy in most hardware components, including CPU, memory, I/O subsystems, and other elements. A fault-tolerant system is one that has the ability to continue service in spite of a hardware or software failure. However, even fault-tolerant systems are subject to outages from human error. Note that HA does *not* imply fault tolerance.

Disaster Tolerance

Disaster tolerance is the ability of a computer installation to withstand multiple outages, or the outage of all the systems at a single site. For HP server installations, disaster tolerance is achieved by locating systems on multiple sites and providing architected solutions that allow one site to take over in the event

What Is High Availability?

of a disaster. The multiple sites in a disaster-tolerant system may be distributed across a single campus, they may be located in different buildings in the same metropolitan area, or they may be dispersed as widely as across a continent or on opposite sides of an ocean. Solutions like these offer the greatest amount of protection for mission-critical data. Needless to say, they can be very expensive to implement, and they all require very careful planning and implementation.

5nines:5minutes

In 1998, HP management committed to a new vision for HA in open systems: 99.999% availability, with no more than five minutes of downtime per year. This ambitious goal has driven the development of many specialized hardware and software facilities by a number of vendors working in partnership. As of the year 2001, HP's own contributions include new generations of fault-resilient HP 9000 systems, improvements in the HP-UX operating system, new software solutions, and extensive monitoring tools that make it possible to measure downtime with a high degree of precision. Many of these improvements have been added back into the standard HP hardware and software products in a kind of "trickle-down" of technological improvement.

Not all users need every type of device or tool used to provide availability levels as high as 99.999%. Not all users wish to pay the price that such tools command in the marketplace. But everyone benefits from the effort to meet the goal of a very high degree of availability as the technology advances. Consider the analogy of race car engines: Even though you don't expect to see

Basic High Availability Concepts

a race car engine in a family sedan, the technology used in building and improving the race car engine eventually ends up improving the sedan anyway.

E-vailable Computing

The phenomenal expansion of Internet business activity has created the need to define yet another type of availability: **e-vailability**, the availability of a server to support fast access to a Web site. It is well known that at periods of peak demand, Web sites suffer performance degradation to the point that users cancel an attempted transaction in frustration at waiting too long or at being refused access temporarily. E-vailability is a combination of the traditional kinds of availability described earlier in this chapter and sufficient server performance and capacity to meet peak demands. Figure 1.3 shows the relationship of availability, performance, and capacity to achieve high levels of e-vailability.

By managing the components of e-vailability, you can allocate different levels of availability to different users depending on their standing as customers. For example, the premier customer class might be given the quickest access to a Web site, say under one second, whereas ordinary customers might get access in one to five seconds, and non-customers (simple Internet cruisers) might obtain access in five to ten seconds. Thus within the framework of e-vailability, HA can be a commodity that customers pay for, or it can be a reward for loyalty or high spending levels.

What Is High Availability?

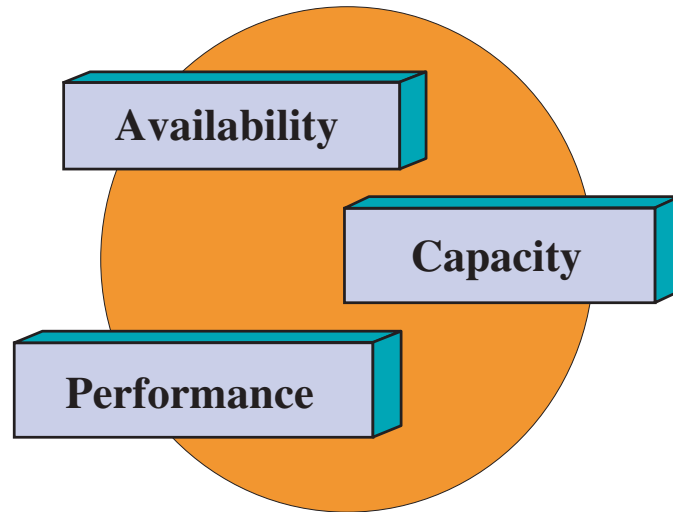


Figure 1.3 *Components of E-vailability*

Matching Availability to User Needs

A failure affects availability when it results in an unplanned loss of service that lasts long enough to create a problem for users of the system. User sensitivity will depend on the specific application. For example, a failure that is corrected within one second may not result in any perceptible loss of service in an environment that does on-line transaction processing (OLTP); but for a scientific application that runs in a real-time environment, one second may be an unacceptable interval.

Basic High Availability Concepts

Since any component can fail, the challenge is to design systems in which problems can be predicted and isolated before a failure occurs and in which failures are quickly detected and corrected when they happen.

Choosing a Solution

Your exact requirement for availability determines the kind of solution you need. For example, if the loss of a system for a few hours of planned downtime is acceptable to you, then you may not need to purchase storage products with hot pluggable disks. On the other hand, if you cannot afford a planned period of maintenance during which a disk replacement can be performed on a mirrored disk system, then you may wish to consider an HA disk array that supports hot plugging or hot swapping of components. (Descriptions of these HA products appear in later sections.)

Keep in mind that some HA solutions are becoming more affordable. The trickle-down effect has resulted in improvements in HP's Intel-based NetServer systems, and there has been considerable growth in the number of clustering solutions available for PCs that use the Windows and Linux operating systems.

High Availability as a Business Requirement

In the current business climate, HA computing is a requirement, not a luxury. From one perspective, HA is a form of insurance against the loss of business due to computer downtime. From another point of view, HA provides new opportunities by allowing your company to provide better and more competitive customer service.

High Availability as Insurance

High availability computing is often seen as insurance against the following kinds of damage:

- Loss of income
- Customer dissatisfaction
- Missed opportunities

For commercial computing, a highly available solution is needed when loss of the system results in loss of revenue. In such cases, the application is said to be *mission-critical*. For all mission-critical applications—that is, where income may be lost through downtime—HA is a requirement. In banking, for example, the ability to obtain certain account balances 24 hours a day may be mission-critical. In parts of the securities business, the

Basic High Availability Concepts

need for HA may only be for that portion of the day when a particular stock market is active; at other times, systems may be safely brought down.

High Availability as Opportunity

Highly available computing provides a business opportunity, since there is an increasing demand for “around-the-clock” computerized services in areas as diverse as banking and financial market operations, communications, order entry and catalog services, resource management, and others. It is not possible to give a simple definition of when an application is mission-critical or of when a highly available application creates new opportunities; this depends on the nature of the business. However, in any business that depends on computers, the following principles are always true:

- The degree of availability required is determined by business needs. There is no absolute amount of availability that is right for all businesses.
- There are many ways to achieve HA.
- The means of achieving HA affects all aspects of the system.
- The likelihood of failure can be reduced by creating an infrastructure that stresses clear procedures and preventive maintenance.
- Recovery from failures must be planned.

Some or all of the following are expectations for the software applications that run in mission-critical environments:

What Are the Measures of High Availability?

- There should be a low rate of application failure, that is, a maximum time between failures.
- Applications should be able to recover after failure.
- There should be minimal scheduled downtime.
- The system should be configurable without shutdown.
- System management tools must be available.

Cost of High Availability

As with other kinds of insurance, the cost depends on the degree of availability you choose. Thus, the value of HA to the enterprise is directly related to the costs of outages. The higher the cost of an outage, the easier it becomes to justify the expense of HA solutions. As the degree of availability approaches the ideal of 100% availability, the cost of the solution increases more rapidly. Thus, the cost of 99.95% availability is significantly greater than the cost of 99.5% availability, the cost of 99.5% availability is significantly greater than 99%, and so on.

What Are the Measures of High Availability?

Availability and reliability can be described in terms of numbers, though doing so can be very misleading. In fact, there is no standard method for modeling or calculating the degree of availability in a computer system. The important thing is to create

Basic High Availability Concepts

clear definitions of what the numbers mean and then use them consistently. Remember that availability is not a measurable attribute of a system like CPU clock speed. Availability can only be measured historically, based on the behavior of the actual system. Moreover, in measuring availability, it is important to ask not simply, “Is the application available?” but, “Is the entire system providing service at the proper level?”

Availability is related to reliability, but they are not the same thing. Availability is the percentage of total system time the computer system is accessible for normal usage. Reliability is the amount of time before a system is expected to fail. Availability includes reliability.

Calculating Availability

The formula in Figure 1.4 defines availability as the portion of time that a unit can be used. Elapsed time is continuous time (operating time + downtime).




Figure 1.4 Availability

What Are the Measures of High Availability?

Availability is usually expressed as a percentage of hours per week, month, or year during which the system and its services can be used for normal business.

Expected Period of Operation

Measures of availability must be seen against the background of the organization's expected period of operation of the system. The following tables show the actual hours of uptime and downtime associated with different percentages of availability for two common periods of operation.

Table 1.1 shows 24x7x365, which stands for a system that is expected to be in use 24 hours a day, 7 days a week, 365 days a year.

Table 1.1 *Uptime and Downtime for a 24x7x365 System*

Availability	Minimum Expected Hours of Uptime	Maximum Allowable Hours of Downtime	Remaining Hours
99%	8672	88	0
99.5%	8716	44	0
99.95%	8755	5	0
100%	8760	0	0

Basic High Availability Concepts

The table shows that there is no remaining time on the system at all. All the available time in the year (8760 hours) is accounted for. This means that all maintenance must be carried out either when the system is up or during the allowable downtime hours. In addition, the higher the percentage of availability, the less time allowed for failure.

Table 1.2 shows a 12x5x52 system, which is expected to be up for 12 hours a day, 5 days a week, 52 weeks a year. In such an example, the normal operating window might be between 8 A. M. and 8 P. M., Monday through Friday.

Table 1.2 *Uptime and Downtime for a 12x5x52 System*

Availability	Minimum Expected Hours of Uptime	Maximum Allowable Hours of Downtime During Normal Operating Window	Remaining Hours
99%	3088	32	5642
99.5%	3104	16	5642
99.95%	3118	2	5642
100%	3118	0	5642

What Are the Measures of High Availability?

This table shows that for the 12x5x52 system, there are 5642 hours of remaining time, which can be used for planned maintenance operations that require the system to be down. Even in these environments, *unplanned* downtime must be carefully managed.

Calculating Mean Time Between Failures

Availability is related to the failure rates of system components. A common measure of equipment reliability is the mean time between failures (MTBF). This measure is usually provided for individual system components, such as disks. Measures like these are useful, but they are only one dimension of the complete picture of HA. For example, they do not take into account the differences in recovery times after failure.

MTBF is given by the formula shown in Figure 1.5.



$$\text{MTBF} = \frac{\text{Total Operating Time}}{\text{Total Number of Failures}}$$

Figure 1.5 Mean Time Between Failures

Basic High Availability Concepts

The MTBF is calculated by summing the actual operating times of all units, including units that do not fail, and dividing that sum by the sum of all failures of the units. Operating time is the sum of the hours when the system is in use (that is, not powered off).

The MTBF is a statement of the time between failures of a unit or units. In common applications, the MTBF is used as a statement of the expected future performance based on the past performance of a unit or population of units. The failure rate is assumed to remain constant when the MTBF is used as a predictive reliability measure.

When gauging reliability for multiple instances of the same unit, the individual MTBF figures are divided by the number of units. This may result in much lower MTBF figures for the components in the system as a whole. For example, if the MTBF for a disk mechanism is 500,000 hours, and the MTBF of a disk module including fans and power supplies is 200,000 hours, then the MTBF of 200 disks together in the system is 1000 hours, which means about 9 expected failures a year. The point is that the greater the number of units operating together in a group, the greater the expected failure rate within the group.

Understanding the Obstacles to High Availability

It is important to understand the obstacles to HA computing. This section describes some terms that people often use to describe these obstacles.

A specific loss of a computer service as perceived by the user is called an **outage**. The duration of an outage is **downtime**. Downtime is either planned or unplanned. Necessary outages are sometimes planned for system upgrades, movement of an application from one system to another, physical moves of equipment, and other reasons.

Unplanned outages occur when there is a failure somewhere in the system. A failure is a cessation of normal operation of some component. Failures occur in hardware, software, system and network management, and in the environment. Errors in human judgment also cause failures. Not all failures cause outages, of course, and not all unplanned outages are caused by failures. Natural disasters and other catastrophic events can also disrupt service.

Duration of Outages

An important aspect of an outage is its duration. Depending on the application, the duration of an outage may be significant or insignificant. A 10-second outage might not be critical, but two

Basic High Availability Concepts

hours could be fatal in some applications; other applications cannot even tolerate a 10-second outage. Thus, your characterization of availability must encompass the acceptable duration of outages. As an example, many HP customers desire 99.95% availability on a 24x7 basis, which allows 5 hours of downtime per year. But they still need to determine what is an acceptable duration for a single outage. Within this framework, many customers state that they can tolerate single unplanned outages with a duration of no more than 15 minutes, and they can tolerate a maximum of 20 such outages per year. Other customers frequently wish to schedule *planned* downtime on a weekly, monthly, or quarterly basis. Note that allowing for planned downtime at a given level of availability reduces the number or duration of unplanned outages that are possible for the system.

Time Lines for Outages

The importance of HA can be seen in the following illustrations, which show the time lines for a computer system outage following a disk crash. Figure 1.6 shows a sequence of events that might take place when an OLTP client experiences a disk crash on a conventional system using unmirrored disks for data; when the disk crashes, the OLTP environment is unavailable until the disk can be replaced.

Understanding the Obstacles to High Availability

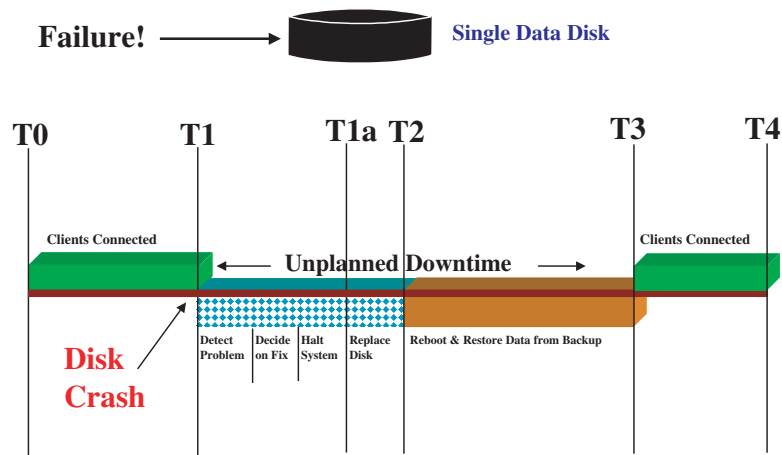


Figure 1.6 Time Line 1: Unplanned Downtime

The crash takes place at T1, and the user's transaction is aborted. The system remains down until T3, following a hardware replacement, system reboot, and database recovery, including the restoration of data from backups. This sequence can require anything from a few hours to over a day. In this scenario, the time to recovery may be unpredictable. Downtime is unplanned, and therefore out of the organization's control.

Basic High Availability Concepts

Figure 1.7 shows the same crash when the system uses an HA feature known as disk mirroring, which prevents the loss of service.

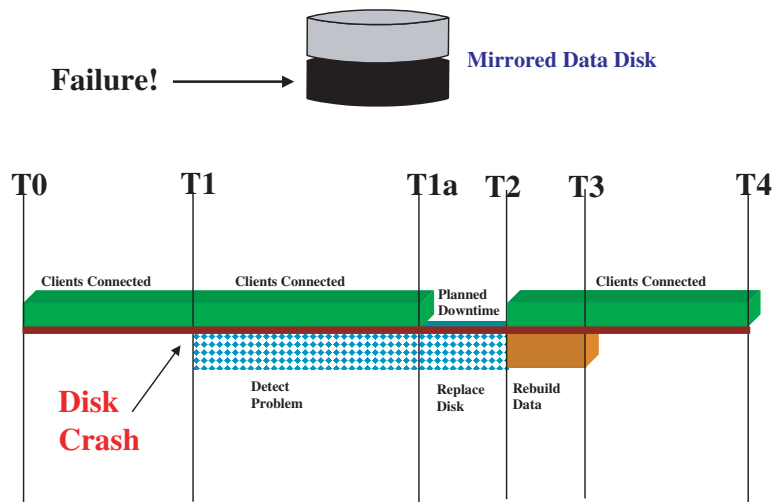


Figure 1.7 Time Line 2: Planned Downtime

When the crash occurs, the mirror disk continues to be available, so no data is lost, and service continues. Further, the replacement of the failed disk can be deferred to a period of planned maintenance. A significant difference between this scenario and the preceding one is that you can predict the amount of time needed for the repair, and you can plan the replacement for

Understanding the Obstacles to High Availability

the least inconvenient time. With disk mirroring, an unpredictable amount of unplanned downtime is replaced by a shorter known period of planned downtime.

A third scenario, shown in Figure 1.8, includes a disk array with hot-swappable disks. This configuration eliminates all downtime associated with the disk failure. (Similar results can be obtained with HA disk enclosures that allow hot plugging of mirrored individual disks.)

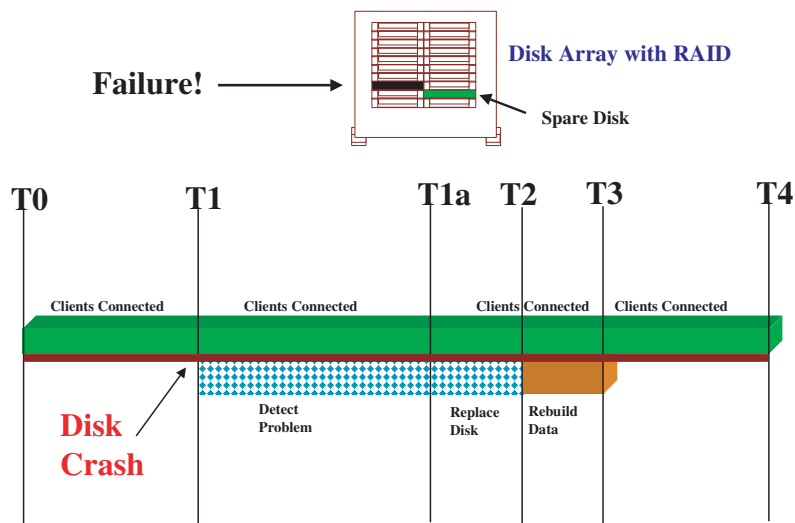


Figure 1.8 Time Line 3: Downtime Eliminated

When the crash occurs, a spare disk takes over for the failed mechanism. In this case, the disk array provides complete redundancy of disks, and the failed disk may be replaced by hot plug-

Basic High Availability Concepts

ging a new disk mechanism *while the system is running*. After the replacement disk is inserted, the array returns to the state it was in before the crash.

Causes of Planned Downtime

Planned outages include stopping an application to perform a scheduled backup or to install a software patch. Some others include:

- Periodic backups
- Software upgrades
- Hardware expansions or repairs
- Changes in system configuration
- Data changes

These outages do not normally cause problems if they can be scheduled appropriately. Some data processing environments can tolerate very little planned downtime, if any. Most can tolerate, and plan for, a regular down period every day or week.

An alternative to planned downtime is to carry out maintenance and other system operations while the system is on-line. Backup operations while the system is running are known as **on-line backups**. Hardware upgrades or repairs while the system is running are known as **hot-plug operations**.

Understanding the Obstacles to High Availability

Causes of Unplanned Downtime

The following are some common causes of unplanned outages:

- Hardware failure
- File System Full error
- Kernel In-Memory Table Full error
- Disk full
- Power spike
- Power failure
- LAN infrastructure problem
- Software defect
- Application failure
- Firmware defect
- Natural disaster (fire, flood, etc.)
- Operator or administrator error

As far as severity is concerned, an unplanned service outage has a far greater negative impact on the enterprise than a planned outage.

Severity of Unplanned Outages

The effects of unplanned outages include customers waiting in line during a computer crash, airplanes unable to take off or land because of an air traffic control failure, an assembly line coming to a halt, doctors unable to obtain patient data from a hos-

Basic High Availability Concepts

pital information system, and so on. In many cases, business is lost because transactions cannot be completed. An unplanned outage most often reduces customer satisfaction.

Figure 1.9 helps to define the total size of the problem. This figure shows the results of a 1998 Gartner Group survey that measured the causes of unplanned service outages.




Figure 1.9 *Causes of Unplanned Service Outages*

Traditionally, most people think only of hardware when they think of HA requirements. But as this survey shows, there are other major factors that must be considered when trying to design and implement a true HA environment within your enterprise. The categories of software and people (system and network management) must be factored into this requirement. There are also external elements such as electric and telecommunications services, climate, and weather that must be considered.

Designing for Reaction to Failure

High availability computer systems are designed to eliminate or minimize planned and unplanned outages. In any HA system, it is important to understand the different types of possible failures and how the system will respond to them. Not all outages are caused by failures, but failures will definitely cause outages unless you take steps to intercept them.

Identifying Points of Failure

Availability can be seen as a chain of services that must remain unbroken. Failures are breaks in the chain. The weak links are known as **points of failure**. For each link in the chain that is a possible point of failure, you can reduce the chance of outage by providing a backup or alternate link. This process is called **eliminating points of failure** in the system. The next chapter, “Clustering to Eliminate Single Points of Failure,” describes this process in some detail.

Preparing Your Organization for High Availability

Often, the greatest obstacle to HA computing is not a hardware or software failure, but a lack of process. In many respects, HA is a mindset as well as a technology, so the human dimension of the HA system will always be the source of additional failure

Basic High Availability Concepts

points. Therefore, it is essential to develop an organization that sees HA as the main priority and that has the skills to cope with the demands of the HA environment. A few suggestions are offered here. In addition, consulting services can assist in carrying out the necessary adjustments in organization that will make the move to HA successful for you.

Identifying Peak Demands and Potential Loss of Availability

In the Internet world of e-availability, you need to estimate the type of demand your system will be subjected to. Web site demand can grow rapidly, even suddenly, as many startups have discovered, so the system must be capable of rapid and sustained growth in response to the demand. Here are some questions to ask:

- What is the expected initial demand for the system?
- How can the demand be controlled?
- What is the maximum possible demand for the market?
- What is the expected rate of growth in the demand?

Planning should include guidelines for expanding service to meet increased needs as they occur.

Stating Availability Goals

To begin with, it is important to state availability goals explicitly. A service level agreement (SLA), negotiated with the users of the system, is a way of explicitly stating availability requirements in terms of services that are provided to clients

Preparing Your Organization for High Availability

within your organization. The SLA can state the normal periods of operation for the system, list any planned downtime, and state specific performance requirements.

Examples of items that appear in SLAs include:

- System will be 99.5% available on a 24x5x52 basis.
- Response time will be 1-2 seconds for Internet-connected clients except during incremental backups.
- Full backups will take place once each weekend as planned maintenance requiring 90 minutes.
- Incremental on-line backups will be taken once a day during the work week with an increase in response time from 2 to 3 seconds for no more than 30 minutes during the incremental backup.
- Recovery time following a failure will be no more than 15 minutes.

The SLA is a kind of contract between the information technology group and the user community. Having an explicit goal makes it easier to see what kind of hardware or software support is needed to provide satisfactory service. It also makes it possible to identify the cost/benefit tradeoff in the purchase of specialized HA solutions.

Note that the system architecture will be different for a large back-end database server than for a system of replicated Internet servers that handle access to a Web site. In the latter case, the use of multiple alternate server systems with sufficient capacity can provide the necessary degree of e-availability.

Building the Appropriate Physical Environment

Achieving HA requires some attention to the physical data processing environment. Since even small outages are not acceptable, it is important to control as much of this environment as possible so as to avoid problems with overheating, cable breakage, and physical jostling of the system. In addition, highly available systems should be physically secure, possibly under lock and key, and available by login only to authorized personnel.

Creating Automated Processes

Human intervention is always error-prone and unpredictable. Therefore, it is good policy in developing an HA organization to automate as many processes as possible. The following are good candidates for automation through scripts:

- Routine backups
- Routine maintenance tasks
- Software upgrades
- Recoveries following failures

The exact content of scripts for each of these processes will vary, but the use of automation will help prevent outages in the first place, and help restore service as quickly as possible when an outage occurs. In particular, recovery processes should be scripted and rehearsed so that the minimum time will be required when recovery is necessary.

Preparing Your Organization for High Availability

Another important use of automation is in the monitoring of processes that run on the HA system. Monitor scripts or programs can detect problems early and signal the need for corrective action. In some cases, a monitor script can be designed to initiate corrective action on its own, leaving a log message describing the action that was taken. When software is designed to facilitate monitoring in this way, the likelihood of a software failure decreases. Specialized software tools can also provide monitoring and early detection of problems.

Using a Development and Test Environment

When rolling out a new software module that is to run on a highly available system, it is critical to give the module a trial in a test environment before installing it. This avoids the significant risk of disrupting the HA system if the new component brings the system down. In other words, the HA system must be well insulated from software that is untested or is of unknown quality.

Maintaining a Stock of Spare Parts

Another useful tactic in maintaining an HA system is to keep on hand a stock of spare parts that can serve as replacements when hardware failures occur. This stock might include disk mechanisms, power supplies, LAN cards and other network components, and a supply of cables.

Purchasing Capacity on Demand

Another aspect of HA through the use of redundant components is **capacity on demand**. This is essentially a marketing model in which a system is delivered to you with extra capacity, but you pay only for the capacity you use. Then if you develop a need for more memory, additional system processor units (SPUs), or more LAN or I/O capacity, you can enable them simply by entering a lockword. Capacity on demand thus reduces the downtime associated with adding components.

Defining an Escalation Process

When a problem occurs, system administrators and operators must know how to decide on a course of action. This means knowing:

- When automated recovery is taking place
- When a system failure requires action by an operator or administrator
- When a support call is required
- When disaster recovery is necessary

Planning for Disasters

A final aspect of organizational planning for HA is to develop a clear strategy for dealing with a natural or man-made disaster. Under the pressure of a catastrophe, having a scripted, tested procedure ready to execute at a damaged site or at a remote recovery site can make a big difference in the organization's ability to recover.

Training System Administration Staff

System administrators must be trained to think in terms of HA, since the procedures used in the HA environment are frequently different from those for conventional systems. Administrators and operators also need special training to recognize and take correct action swiftly in the event of a component failure. This is especially important since failures are not common, and the “lights out” environment of many HA installations means that a system administrator may not experience a problem very frequently.

Using Dry Runs

One way of providing training is to conduct dry runs or rehearsals of recovery scenarios — simulating a problem and then walking through the solution on the development system.

Documenting Every Detail

Not least in importance in developing an HA environment is to document every detail of the hardware and software configuration and to create a procedures document that is periodically updated and reviewed by system administration staff. Whenever anything in the configuration is added or modified—including hardware components, software modules, and operator procedures—it should be recorded in this document.

Another important document that should be maintained and frequently reviewed is a log of all exceptional activity that takes place on the HA system. This log can include system log file

Basic High Availability Concepts

entries. It should also include a list of what corrective actions are taken on the system, by whom, with dates and times. Most importantly, the time required to return service should be carefully recorded for every failure that results in downtime. Planned downtime events may also be logged.

The Starting Point for a Highly Available System

A highly available system is built on top of highly reliable components. HP's enterprise-class servers have the following features, which are the first requirements for components that are to be made highly available:

- Basic hardware reliability
- Software quality
- Intelligent diagnostics
- Comprehensive system management tools
- Maintenance and support services

High availability is not guaranteed by these features, but together they improve the overall availability of the system.

The Starting Point for a Highly Available System

Basic Hardware Reliability

The best way to deliver HA is never to fail in the first place. HP has made a significant investment in designing and manufacturing extremely reliable components for its systems. This results, of course, in highly reliable servers. However, standard reliability will not alone meet the availability requirements of a mission-critical application. For example, all disk drives, being mechanical, go through a life cycle that eventually ends in device failure: no disk will perform forever. Therefore, specific HA storage solutions like disk arrays or disk mirroring are crucial to maintaining HA.

Software Quality

Software quality is another critical factor in the overall scope of HA and must be considered when planning a highly available processing environment. The presence of a software defect can be every bit as costly as a failed hardware component. Thus, the operating system, middleware modules, and all application programs must be subjected to a rigorous testing methodology.

Intelligent Diagnostics

Sophisticated on-line diagnostics should be used to monitor the operating characteristics of important components such as the disks, controllers, and memory, and to detect when a component is developing problems. The diagnostic can then proactively notify the operator or the component vendor so that corrective maintenance can be scheduled long before there is a risk of an

Basic High Availability Concepts

unplanned service outage. These intelligent diagnostics improve overall availability by transforming unplanned downtime into planned maintenance.

Comprehensive System Management Tools

System and network management are other major areas that need to be considered for minimizing outages. This is not a criticism of operators, but an acknowledgment of the complexity of today's computing environments. We are now building extremely complex networks that are difficult to manage without automated tools. A single small operator mistake can lead to serious outages.

An integrated set of system and network administration tools such as those provided on HP's OpenView platform can reduce the complexities of managing a multi-vendor distributed network. By automating, centralizing, and simplifying, these tools can significantly reduce the complexity of management tasks. Some tools also have the ability to detect and automatically respond to problems on systems, thus eliminating downtime. Chapter 3, "High Availability Cluster Components," describes some of these tools in more detail.

Maintenance and Support Services

Over time, it will become necessary to upgrade software and hardware components. Also, no matter how reliable a system is, components do fail. For these reasons, it is important to establish hardware and software support contracts with the suppliers of your system components. HP provides a large variety of support

Moving to High Availability

levels, including several that are specifically designed for HA users. Consulting is also available during all the phases of deploying an HA system.

Moving to High Availability

Starting with conventional, highly reliable systems, you can obtain HA in several ways:

- By providing redundancy of components
- By using software and hardware switching techniques
- By carefully planning all scheduled downtime
- By eliminating human interaction with the system
- By defining automatic responses to error conditions and events
- By using comprehensive acceptance tests
- By defining and practicing operator responses to unplanned outages that are not handled by automatic response

The use of redundant components eliminates single points of failure in a system, allowing a spare component to take over as needed. Software and hardware switching are what allow a spare component to replace a failed component. In addition, the HA system should attempt to avoid or reduce application outages for planned maintenance; if planned outages cannot be avoided, their duration should be minimized.

Basic High Availability Concepts

Eliminating human interaction allows you to create deterministic responses to error conditions: the same error condition always results in the same system response. The use of networked monitoring tools also lets you automate responses to errors.

Any proposed HA design should be thoroughly tested before being placed in production.

NOTE: If you are really concerned about HA, there is no room for compromise. The upper-most goal must be meeting the HA requirements, and other considerations, such as cost and complexity, take second place. It is important to understand these tradeoffs.

Summary

A highly available system must be designed carefully on paper. It is important to do the following *in the order specified*:

1. Define a goal for availability, including a detailed listing of your service level objectives for each application or service.
2. Identify the maximum duration of an acceptable outage.
3. Measure the availability of the current system, if one is in use. This includes understanding current statistics on availability, including planned and unplanned downtime. Be sure to use measurements consistently, and make sure everyone understands what the measure-

Summary

ments mean. Identify all the single points of failure in the current system.

4. Assess your applications. What improvements or changes in availability are desired? What are the costs involved?
5. In conjunction with the next two steps, choose an architecture for HA. More information on this subject follows in Chapter 2.
6. Purchase and configure HA component hardware and software and obtain support services.
7. Create or modify application software.
8. Choose system administration tools.
9. Design processes to be followed in the event of a failure.
10. Document these processes.
11. Train administrators and operators on these processes, and conduct dry runs to prove they work.
12. Document and record the state of the system.
13. Review existing processes on a regular basis.
14. Enforce a change management policy.

The following chapters can help you create an HA design. The emphasis will be on enterprise cluster solutions, which provide the redundancy and switching capabilities that are needed. In addition, HP Consulting can offer assistance in carrying out the analysis, architecting the environment, and choosing the appropriate HA cluster solution for your needs.