

No Time for Downtime

We gain strength, and courage, and confidence by each experience in which we really stop to look fear in the face . . . we must do that which we think we cannot.

— Eleanor Roosevelt

The need for high availability did not originate with the Internet or e-commerce. It has existed for thousands of years. When Greek warships or merchant ships sailed to discover new lands or business, the captains carried spare sails and oars on board. If the primary sail failed, the crew would immediately hoist a replacement and continue on their way, while they repaired damaged sails. With the advent of electronic sensors, the spare parts employed in industrial systems did not need human intervention for activation. In the early twentieth century, electric power-generating plants automatically detected problems, if any, in the primary generator and switched to a hot standby unit.

With the recent explosive growth of the Internet and our dependence on information systems, *high availability* has taken on a new meaning and importance. Businesses and consumers are turning to the Internet for purchasing goods and services. People conduct business anytime from their computer. They expect to buy clothes at 2 a.m. on the Web and expect the site to function properly, without problem or delay, from the first click to the last. If the Web site is slow or unavailable, they will click away to a competitor's site. Business globalization caused by the Internet adds another layer of complexity. A popular online store, with business located in Bismarck, North Dakota, may have customers in Asia who keep the seller's servers busy during quiet hours in the United States. Time zones, national borders, and peak and off-peak hours essentially disappear on the Web.

As computers get faster and cheaper, they are being used for more and more critical tasks that require 24-7 uptime. Hospitals, airlines, online banking services, and other service industries modify customer-related data in real time. The amount of online data is rapidly expanding. It is estimated that online data will grow more than 75 percent every year for the next several years. The rapidly increasing demand for placing more and more data online and the constantly decreasing price of storage media have resulted in an increase of huge amounts of critical information being placed online.

Employees and partners depend on data being available at all times. Work hours have extended beyond the traditional 9-to-5, five days a week. Intranet servers such as e-mail, internal applications, and so forth, must be always up and functional for work to continue. Every company has at least one business-critical server that supports the organization's day-to-day operation and health. The unavailability of critical applications translates to lost revenue, reduced customer service and customer loyalty, and well-paid, but idle, workers. A survey of 450 Fortune 100 companies (conducted by the Strategic Research Division of Find/SVP) concluded that U.S. businesses incur about \$4 billion of losses per year because of system or network downtime.

In fact, analysts estimate that every minute of Enterprise Resource Planning (ERP) downtime could cost a retailer between \$10,000 and \$15,000. Systems and data are not expected to be down, not even for maintenance. Downtime literally freezes customers, employees, and partners, who cannot even complete the most basic daily chores.

The requirements for reliability and availability put extreme demands on servers, network, software, and supporting infrastructure. Corporate and e-commerce sites must be capable of processing large numbers of concurrent transactions and are configured to operate 24-7. All components, including both the server hardware and software, must be configured to be redundant.

And what happens when no one can get to the applications? What happens when data is unreachable and the important servers do not want to boot up? Can you shut down your business and ask your employees to go home? Can you tell your customers to go somewhere else? How is it that no one planned for this scenario? Is it possible to recover from this? How long will it take and how much will it cost? What about reputation among customers? Will they ever come back? Why doesn't this happen to your competitors?

As you can see, it happens all the time and all around us. Following are some events that have occurred over the last few years. They expose our total dependence on computer systems and utter helplessness if critical systems are down.

- In April of 1998, AT&T had a 26-hour frame relay-network outage that hurt several business customers. In December of 1999, AT&T had an 8-hour outage that disrupted services to thousands of AT&T WorldNet dial-up users.

- In early 1999, customers of the popular online stock trading site ETrade could not place stock trade orders because the trading sites were down. At the same time, there were a few outages at The Charles Schwab Corporation because of operator errors or upgrades. Schwab later announced a plan to invest \$70 million in information technology (IT) infrastructure.
- In June of 1999, eBay had a 22-hour outage that cost the company more than \$3 million in credits to customers and about \$6 billion (more than 20 percent) in market capitalization. In January of 2001, parts of the site were again down for another 10 hours.
- In August of 1999, MCI suffered about 10 days of partial outages and later provided 20 days of free service to 3,000 enterprise customers.
- Three outages at the Web retailer `amazon.com` during the busy holiday-shopping season of December 2000 cost Amazon more than \$500,000 in sales loss.
- Denial-of-Service and several virus-induced attacks on Internet servers continue to cause Web site outages. On July 19, 2002, a hacker defaced a page on the U.S. Army Research Laboratory's Web site with a message criticizing the Army's organization for bias to certain nations.
- Terrorist attacks in Washington, D.C., New York, London, and cities around the world in recent years have destroyed several data centers and offices.

Businesses everywhere are faced with the challenge of minimizing downtime. At the same time, plans to enhance service availability have financial and resource-related constraints. Taking steps to increase data, system, and network availability is a delicate task. If the environment is not carefully designed and implemented, it would cost dearly (in terms of required time, money, and human resources) to build and manage it.

To increase service availability, you must identify and eliminate potential causes of downtime, which could be caused by hardware failures, network glitches, software problems, application bugs, and so forth. Sometimes, poor server, application, or network performance is perceived as downtime. Service expectations are high. When someone wants to place a phone call, he or she picks up the phone and expects a dial tone within a few seconds. The call must connect within one second of dialing and there should be no dropped connections. When surfing the Web, users expect the first visual frame of a Web page within a few seconds of accessing the site. All systems, especially those related to consumers and critical operations, should always be ready and must operate with no lost transactions.

But potential causes of downtime abound. The entire IT infrastructure is made up of several links, such as user workstations, network devices, servers,

applications, data, and so forth. If any link is down, the user is affected. It then does not matter if the other links in the chain are available or not. *Downtime*, in this book, is defined as an end user's inability to get his or her work done. This book examines ways to enhance service availability to the end user and describes techniques for improving network, data, server, and application uptime.

Availability is the portion of time that an application or service is available to internal or external customers for productive work. The more resilient a system or environment is, the higher the availability is. An important decision is the required availability level. When you ask a user or project manager how much uptime he or she needs for the application, the reflex answer is "One-hundred percent. It must always be available at all times." But when you explain the high costs required to achieve 100 percent uptime, the conversation becomes more of a two-way negotiation. The key point is to balance downtime cost with availability configuration costs.

Another point is the time duration when 100 percent uptime is necessary. Network Operations Center (NOC) and 24-7 network monitoring applications and e-commerce Web sites require 100 percent uptime. On the other extreme are software development environments, used only when developers are accessing the system. If, on occasion, you take development systems down (especially at night or on weekends), and if you warn your users well in advance, downtime is not an issue.

Table 1-1 illustrates how little time per year is afforded for planned or unplanned downtime as availability requirements move closer to 100 percent. Suppose a server, "hubble," has no special high-availability features except for RAID-1 volumes and regular backups and has 98 percent uptime. The 2 percent downtime is too high and, therefore, it is clustered with "casper," which also has 98 percent uptime. Server "casper" is used only 2 percent of the time when hubble is down. The combined availability is 98 percent plus 98 percent of 2 (0.98×2), which is 1.96 percent. These add to a theoretical service uptime of 99.96 percent for the two-node cluster.

Table 1-1 Total Allowable Downtime for Planned or Unplanned Events

PERCENTAGE UPTIME	PERCENTAGE DOWNTIME	DOWNTIME PER YEAR	DOWNTIME PER MONTH
98%	2%	7.3 days	14 hours 36 minutes
99%	1%	3.65 days	7 hours 18 minutes
99.9%	0.1%	8 hours 45 minutes	43 minutes 45 seconds
99.99% ("four nines")	0.01%	52.5 minutes	4 minutes 22 seconds
99.999% ("five nines")	0.001%	5.25 minutes	26 seconds

In reality, several other factors affect both servers, such as downtime during failover duration, power or network outages, and application bugs. These failures will decrease the theoretical combined uptime.

As you move down the table, the incremental costs associated with achieving the level of availability increase exponentially. It is far more expensive to migrate from a “four-nines” to a “five-nines” (99.99 percent to 99.999 percent uptime) configuration than to move from 99 percent to 99.9 percent uptime.

Causes of Downtime

About 80 percent of the unplanned downtime is caused by process or people issues, and 20 percent is caused by product issues. Solid processes must be in place throughout the IT infrastructure to avoid process-, people-, or product-related outages. Figure 1-1 and Table 1-2 show the various causes of downtime. As you can see, planned or scheduled downtime is one of the biggest contributors (30 percent). It is also the easiest to reduce. It includes events that are pre-planned by IT (system, database, and network) administrators and usually done at night. It could be just a proactive reboot. Other planned tasks that lead to host or application outage are scheduled activities such as application or operating system upgrades, adding patches, hardware changes, and so forth.

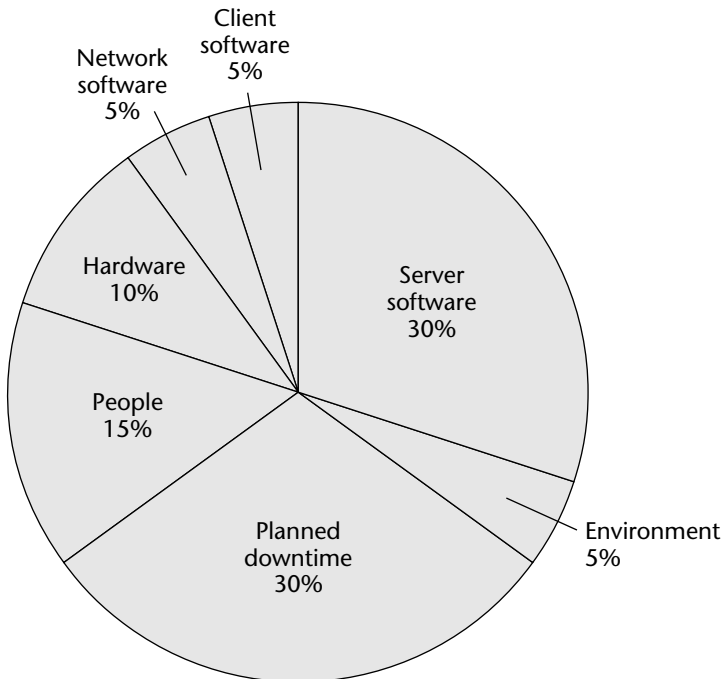


Figure 1-1 Causes of downtime

Table 1-2 Causes of Planned and Unplanned Downtime

CAUSES OF PLANNED DOWNTIME	CAUSES OF UNPLANNED DOWNTIME
Backup	Extended planned downtime
Replace or upgrade hardware	Human error
Software or application upgrade	Application failure
Software maintenance or reconfiguration	Operating system failure
Operating system upgrade	Hardware failure such as disk, CPU, memory
Patch installation	Incompatibility/conflict between application parameters

Most of these planned events can be performed without service interruption. Disks, fans, and power supplies in some servers and disk subsystems can be changed during normal run-time, without need for power-offs. Data volumes and files systems can be increased, decreased, or checked for problems while they are online. Applications can be upgraded while they are up. Some applications must be shut down before an upgrade or a configuration change.

Outages for planned activities can be avoided by having standby devices or servers in place. Server clustering and redundant devices and links help reduce service outages during planned maintenance. If the application is running in a cluster, it can be switched to another server in the cluster. After the application is upgraded, the application can be moved back. The only downtime is the time duration required to switch or failover services from one server to another. The same procedure can be used for host-related changes that require the host to be taken off-line. Apart from the failover duration, there is no other service outage.

Another major cause of downtime is people-related. It is caused by poor training, a rush to get things done, fatigue, lots of nonautomated tasks, or pressure to do several things at the same time. It could also be caused by lack of expertise, poor understanding of how systems or applications work, and poorly defined processes. You can reduce the likelihood of operator-induced outages by following properly documented procedures and best practices. Organization must have several, easy-to-understand how-tos for technical support groups and project managers. The documentation must be placed where it can be easily accessed, such as internal Web sites. It is important to spend time and money on employee training because in economically good times, talented employees are hard to recruit and harder to retain. For smooth continuity of expertise, it is necessary to recruit enough staff to cover emergencies and employee attrition and to avoid overdependence on one person.

TALES FROM THE TECH TURF: ON-CALL WOES

One organization I worked at rebooted their UNIX servers every Sunday morning at 4 a.m. to clear memory, swap, and process tables. Of course, sometimes the boxes would not boot up all the way and the NOC had to call someone at a weird hour. Later, the reboot time was moved to 6 a.m. This was done to avoid application-related problems on systems with high uptime. This was initially implemented due to a Solaris problem on Suns that had not been rebooted in the last 350 days and were running an old release of Oracle Database.

Avoiding unplanned downtime takes more discipline than reducing planned downtime. One major contributor to unplanned downtime is software glitches. The Gartner Group estimates that U.S. companies suffer losses of up to \$1 billion every year because of software failure. In another survey conducted by Ernst and Young, it was found that almost all the 310 surveyed companies had some kind of business disruption. About 30 percent of the disruptions caused losses of \$100,000 or more each to the company.

When production systems fail, backups and business-continuance plans are immediately deployed and are every bit worth their weight, but the damage has already been done. Bug fixes are usually reactive to the outages they wreak. As operating systems and applications get more and more complex, they will have more bugs. On the other hand, software development and debugging techniques are getting more sophisticated. It will be interesting to see if the percentage of downtime attributed to software bugs increases or decreases in the future. It is best to stay informed of the latest developments and keep current on security, operating system, application, and other critical patches. Sign up for e-mail-based advisory bulletins from vendors whose products are critical to your business.

Environmental factors that can cause downtime are rare, but they happen. Power fails. Fires blaze. Floods gush. The ground below shakes. In 1998, the East Coast of the United States endured the worst hurricane season on record. At the same time, the Midwest was plagued with floods. Natural disasters occur mercurially all the time and adversely impact business operations. And, to add to all that, there are disasters caused by human beings, such as terrorist attacks.

The best protection is to have one or more remote, mirrored disaster recovery (DR) sites. In the past, a fully redundant system at a remote DR site was an expensive and daunting proposition. Nowadays, conditions have changed to make it very affordable:

- Hardware costs and system sizes have fallen dramatically.
- The Internet has come to provide a common network backbone.
- Operating procedures, technology, and products have made an off-site installation easy to manage remotely.

TALES FROM THE TECH TURF

In 1995, a building in downtown Oklahoma City was destroyed by a terrorist. Many offices and data centers lost servers and valuable data. One law firm had no off-site backup of its data and lost its records. The firm was in the business of managing public relations for its clients. It was unable to restore any data pertaining to its customers. Sadly enough, it went out of business within three months. Another commodity-trading firm had a remote mirror site just outside the city and was able to bring up its applications on standby servers at the remote site. It quickly transferred its operations and business to the secondary site.

To protect against power blackouts, use uninterruptible power supplies (UPS). If Internet connection is critical, use two Internet access providers or at least separate, fully redundant links from the same provider.

Cost of Downtime

Organizations need to cost out the financial impact caused by downtime. The result helps determine the extent of resources that must be spent to protect against outages. The total cost of a service outage is difficult to assess. Customer dissatisfaction, lost transactions, data integrity problems, and lost business revenue cannot be accurately quantified. An extended period of downtime can result in ruin and, depending on the nature of the business, the hourly cost of business outage can be several tens of thousands of dollars to a few million dollars. Table 1-3 provides some examples of downtime costs.

Table 1-3 Average Cost of Downtime per Hour for a Variety of Industries

INDUSTRY	BUSINESS OPERATION	DOWNTIME COST RANGE (PER HOUR)	AVERAGE COST OF DOWNTIME (PER HOUR)
Financial	Brokerage operations	\$5.6M to 7.3M	\$6.45M
Financial	Credit card/sales authorizations	\$2.2M to 3.1M	\$2.6M
Media	Pay-per-view TV	\$67K to 233K	\$150K
Retail	Home shopping (TV)	\$87K to 140K	\$113K
Retail	Home catalog sales	\$60K to 120K	\$90K
Transportation	Airline reservations	\$67K to 112K	\$89.5K
Media	Telephone ticket sales	\$56K to 82K	\$69K

Table 1-3 (continued)

INDUSTRY	BUSINESS OPERATION	DOWNTIME COST RANGE (PER HOUR)	AVERAGE COST OF DOWNTIME (PER HOUR)
Transportation	Package shipping	\$24K to 32K	\$28K
Financial	ATM fees	\$12K to 17K	\$14.5K

Source: Computer Economics, Inc., Irvine, California, www.computereconomics.com

It is important to arrive at reasonable estimates of financial losses that could be incurred during an outage. The total cost is calculated by the sum of losses in areas of labor, revenue, and downtime. A good starting point is to collect employee-related statistics from human resources and accounting departments regarding salary, duration of past outages, number of employees in each group, and annual gross revenue from online or telesales.

Employees continue to receive full pay even during planned or unplanned outages. The productivity of the employee is rated to be higher than the salary. Labor cost during an outage is calculated using the following equation:

$$\text{Labor Cost} = \text{Number of Employees} \times \text{Hours of Outage} \times \text{Average Hourly Pay Rate}$$

Several employees within a department and with similar salary can be grouped together. A department with 50 employees that cannot function for a week (assuming 40-hour weeks and average employee expense of \$100 per hour) incurs a loss of \$200,000 for the week. Then there is the cost of overtime that must be paid to catch up on the work. That doubles the loss to \$400,000.

If you estimate the loss of revenue to a company whose sales rely on server and application availability, the lost revenue is calculated using the following equation:

$$\text{Lost Revenue} = \left(\frac{\text{Gross Annual Revenue}}{\text{Annual Business Hours}} \right) \times \text{Percentage Impact} \times \text{Hours of Downtime}$$

The first two elements of the equation provide the revenue generated per hour. The percentage impact allows you to scale the hourly revenue based on whether the lost customers can be partially or entirely recovered, or whether they would cause more damage by creating negative publicity among those who never witnessed the downtime firsthand. During a downtime, customers call in and the operator politely tells them, “Our systems are down. Please call back after a few hours.” Some customers prefer to complete the purchase sooner. They will go to a competitor, buy what they want, and do not need to ever call back. Some will call back. If 75 percent do call back, then the loss and

percentage impact is only 25 percent. One way of collecting data on the percentage of customers who called back is to look at the amount of revenue generated above and beyond the normal orders immediately following the outage. If it is 50 percent more, that means 50 percent of the people called back.

Several other effects of downtime are impossible to quantify. Some customers who were merely inconvenienced may tell the story to friends and recommend to them never to shop at that place. Percentage impact is difficult to estimate and is different during different outages.

The results from the equations convey only part of a company's losses. Satisfied customers become loyal customers and dissatisfied customers do not. Let's say a customer would have made a \$100 purchase and would have repeated that just once a year. Using a discounted cash flow rate of just 15 percent, the present value of those purchases over a 20-year period would have been \$719. The company, therefore, suffered a total loss of more than seven times the first lost sale value.

Downtime causes several other losses that are difficult to predict. Companies that host applications or servers for customers have strict service level agreements (SLAs) that require them to credit customer accounts upon incurring any service disruption. Downtime also causes missed deadlines and penalties. It adversely affects stock price, goodwill among customers and partners, and employee morale.

For servers connected to the Internet, the greatest cause for concern is security. Hackers are quick to take advantage of vulnerabilities. Malicious code attacks have significant economic impact. Table 1-4 shows the adverse economic impact caused by cyber attacks.

Table 1-5 shows the economic impact of specific incidents. To date, the "I Love You" Bug outbreak in 2000 has caused the highest financial damage.

TALES FROM THE TECH TURF: THE WEEKEND MUSIC

On July 25, 2002, the Recording Industry Association of America (RIAA, www.riaa.com) openly backed the U.S. Peer-to-Peer Piracy Prevention Act proposed by Congressman Howard Berman. The bill enables music copyright holders to "block, divert or otherwise impair" networks believed to be involved in music piracy. The bill was criticized for various reasons. On the Internet, it is nearly impossible to tell good computers from bad. Many innocent servers on the Internet are "stolen" for file sharing and participating in "distributed denial of service" (DDoS) attacks. On July 26, 2002, the RIAA Web site was buried under a flood of DDoS attacks from hundreds of computers. The attack continued until 2 a.m. on July 29. The bill allows RIAA and music copyright owners to engage in precisely this kind of denial-of-service attacks against networks that distribute illicit copies of music.

Table 1-4 Annual Economic Impact Due to Cyber Attacks

YEAR	WORLDWIDE ECONOMIC IMPACT (\$ U.S.)
1995	0.5 billion
1996	1.8 billion
1997	3.3 billion
1998	6.1 billion
1999	12.1 billion
2000	17.1 billion
2001	13.2 billion

Source: Computer Economics, Inc., Irvine, California, www.computereconomics.com/

The use of the Internet is rapidly expanding around the world and across cultures, and it spans a wide range of economic and educational strata. With such widespread use, cyber attacks will only become more frequent. Server downtime and adverse economic impact will become more pronounced. Law enforcement agencies around the world face the challenge of investigating, prosecuting, and combating cyber crimes. The liberal use of the Internet has sparked discussions and laws at the highest government levels in almost all countries. The Internet conflicts with local laws governing commerce and openly accessible content. Many countries (such as France, Saudi Arabia, and China) have attempted to control the use of the Web.

Table 1-6 shows the number of Internet users by continent for 2002 with projected growth through 2006. Internet users are those who are connected via LANs, dial-up modems, DSL, cable modems, ITV, wireless connections, and so forth.

Table 1-5 Economic Impact Analysis per Incident

YEAR	CYBER ATTACK CODE NAME	WORLDWIDE ECONOMIC IMPACT (\$ U.S.)
1999	Explorer	1.02 billion
1999	Melissa	1.10 billion
2000	I Love You (Love Bug)	8.75 billion
2001	SirCam	1.15 billion
2001	Code Red(s)	2.62 billion
2001	Nimda	635 million

Source: Computer Economics, Inc., Irvine, California, www.computereconomics.com/

Table 1-6 Internet Users by Continent

CONTINENT OR REGION	2002 (MILLIONS)	2003 (MILLIONS)	2004 (MILLIONS)	2005 (MILLIONS)	2006 (MILLIONS)
North America	212	223	234	244	256
South and Central America	25	32	43	59	80
Europe	163	196	225	240	257
Middle East/Africa	9.2	10.7	11.6	12.6	13.6
Asia and Pacific regions	151	204	238	273	313
TOTAL WORLDWIDE	560.2	665.7	751.6	828.6	919.6

Source: Computer Economics, Inc., Irvine, California, www.computereconomics.com/

Key Points

Following are some key points discussed in this chapter:

- With the fast-growing dependence on computer systems, availability is a requirement for the entire IT infrastructure. Businesses view availability as the single metric of overall system performance.
- Unplanned service outages are expensive and are caused by hardware and software failures, human errors, Internet viruses and attacks, natural disasters, and human-caused crises.
- Financial losses incurred by downtime depend on the business operation.

