CHAPTER 1

Introduction to High Availability Networking

Recently, the growth of the Internet and the use of computing systems to run businesses have blossomed in a way that few would have expected. In 1990, contacting individuals at some companies (mostly computer companies) via electronic mail was possible. Just a decade later, thousands of highly reputable companies were offering consumers the ability to actually purchase products from their web sites on the Internet. This fantastic growth of the Internet and networking has and is changing our lives in many ways.

As the Internet and networking become more a part of our lives, we are becoming dependent on them. We rely on them, and thus, we need them to be highly reliable. When you say that you want a *highly reliable network*, you are saying that you want your network to work all the time—you want it to be *highly available*.

As you proceed through this book, you will find that I have attempted to cover the subject of high availability using only arithmetic and algebra. This is by design and is the key reason I believe many people will be able to understand the material. All the other books I have read about reliability and availability use advanced mathematics such as calculus. While most of us can perform calculus when absolutely required, calculus is not easy to remember, nor is it something most enjoy. I hope that you, the reader, appreciate that what we give up in accuracy and process, we make up for in ease of understanding and accomplishment.

For those of you that want to move on to advanced reliability and availability topics, references in the front of this book will guide you.

Why Do We Need High Availability?

We have come to depend on the use of computers, access to the Internet, and the help of our favorite Internet sites. Many people regularly shop for things on the Internet. They expect to be able to go Internet shopping 24 hours per day, seven days per week.

If you have ever purchased anything online, you have probably felt a bit worried as you typed in your name, address, and credit card number—you feel as though you are giving out information that will result in yet more junk mail and the addition of your name on even more lists that cold callers use. The first-time registration process at most Web sites involves entering a wealth of personal information. On subsequent visits to the same Internet merchant, only a small portion of this information is normally requested. If you go back to the same merchant, you are unlikely to be added to more of those lists. Your privacy is retained. This first-time registration is a small but important barrier to closing the deal for the Internet merchant.

Imagine that you are a merchant on the Internet and someone visits your site and your competitor's site while shopping. Subsequently that person decides to purchase an item that you and your competitor are offering for a similar price. If the person registers on your site and buys the item, he or she very likely will be back to your site the next time a similar item is needed because he or she has already gone through the registration process. If the person registers with your competitor's site, then it is likely that he or she will buy that item from your competitor the next time it is needed. Now imagine that the consumer decided to buy from your competitor because your site was down. You not only lost the sale during the downtime, but you also lost the customer. In sales, this is about as bad as it gets.

Customers don't care why they were unable to access your site to place an order. Whether it was a Web server, network, service provider, or some other problem means nothing to them. All they care about is the fact that your site was unavailable and your competitor's site was available. Keep this story in the back of your head, and you will have no trouble understanding the cost of downtime. When the time comes to make the decision between price and availability, it will be much easier when you understand both the cost of downtime, as well as the cost of making downtime go away. You might say that competition is only a mouse click away!

Today's networks don't just carry sales transactions and business information. In fact in some places, you can pick up a telephone, dial a telephone number, and never even realize that the telephone call went over a data network instead of through the traditional telephone company infrastructure. Voice traffic is becoming yet another part of data networks, and data networks are becoming part of the telephone system.

In some cities, consumers can purchase their telephone service from a cable company instead of a traditional telephone company. These people depend on networking equipment to deliver emergency services. Imagine if you picked up your telephone to dial 911 and didn't get a dial tone. When someone's life depends on a network, network availability is crucial—it can be a matter of life or death. Network availability is now a member of an exclusive club, consisting of hospitals, ambulances, and doctors—things that can save your life.

What Is High Availability?

Have you ever been surfing the Internet and had to wait for a Web page for a minute or two? Have you ever noticed what seemed like a few minutes before anything happening? Now I don't mean waiting for a page to load because of your slow connection and pretty pictures on the page; that is a result of a slow connection. I am talking about those times when everything is going just fine and then, all of a sudden, things just seem like they stopped for a minute.

If you have experienced this, then you understand what it is like to notice some sort of network downtime. Of course, we all know this sort of thing happens all the time. A minute here, a minute there, and nobody actually worries about it too much. But imagine how you would feel if the page stopped loading and didn't budge for an hour. Imagine this happened while you were doing research for a project that needed to be done in 30 minutes. If you are like most folks, a minute here or a minute there won't bother you very much, but an hour might be a different story. An hour might irritate you enough to call someone and ask him or her what is happening with your Internet service.

Now think about your telephone. Do you remember the last time you picked up your telephone and there was no dial tone? Do you remember the last time you were on the telephone and you were disconnected? Some folks remember these things. However, most people have a hard time remembering the last time they had telephone trouble. It's hard to remember because it doesn't happen very often. If you do remember the last time you had no dial tone, how did you feel during that experience? In most cases, the feeling of having no dial tone is considerably more traumatic than having no access to the web. No dial tone equals no help in an emergency. No dial tone is a big deal to everyone. No web access is usually a minor inconvenience.

If you are about to build a network that provides dial tone, then you are going to have to build a highly available network. People will not stand for loss of dial tone because dial tone is something that they require all the time. Networks that provide dial tone must be highly available, or they will not be used. The customer will find a way to replace the network with something more available if they sense the network is not reliable.

Attracting and Keeping Internet Customers

Imagine that you are the customer and that you have a choice of buying your telephone service from your traditional telephone company, your cable television company, or your Internet service provider (ISP). Most people would say that they want to buy their telephone service from the telephone company. After all, the telephone company has been the most reliable service of all their services. The telephone company is even more reliable than the power company in most cities.

Now let us complicate the issue. Say that you can buy your phone service from your ISP for 60 percent less money than you have been paying for your phone service from the phone company. Most people will consider changing services for that much of a discount. The offer is even more attractive if you have two telephone lines. Maybe you can move the second line to the ISP, but leave the first line with the older, more established service. Perhaps you want to just try out the new discounted service, but only on your extra line.

Everyone I know would be willing to put his or her second line on the alternate source for a while. But everyone I asked said that he or she would switch back to the telephone company if the new service proved unreliable. Even at a 60 percent cost discount, high availability is a requirement. Reliability is not something that most people are willing to give up, no matter how much less money they have to spend for their phone service.

High Availability and Government Regulation

In the United States, telephone service is regulated by the government. Telephone companies are required to report downtime of their systems to the Federal Communications Commission (FCC). If a phone company has been negligent or had too much downtime, the government imposes a fine.

Life or Death and High Availability

A friend of mine used to work for a telephone company as a manager of a team that fixed the telephone network whenever it broke. He was under pressure to fix the network very fast whenever it wasn't working.

He once told me, "You haven't ever felt pressure and pain until you've been standing in front of an FCC review panel explaining why the phone system was down when a citizen needed to use it for a 911 phone call. There is no feeling as bad as knowing that someone died because your network was not working."

If you are planning to build a network that carries 911 traffic, you can expect that the government is going to be regulating the network at some point. For now, data networks are not regulated. But someday soon as more data networks carry voice and more of them make it possible for people to call 911 over them these networks will be regulated. Network repair folks and network managers are going to have to report downtime to the government.

Network downtime is going to be a big deal—yet another reason why high availability has to be designed into a network.

Presenting and Describing High Availability Measurements

There are two main ways to state the availability of a network: the *percentage method* and the *defects per million method*. Once you understand the basics of these methods, you will understand what you are reading when you get an availability analysis report. Both of these methods use figures like MTBF (Mean Time Between Failure) and MTTR (Mean Time To

Repair). The following sections describe these two methods as well as the terms MTBF and MTTR.

The Percentage Method

You have probably heard the term *five 9s* in relationship to the availability of a network. When someone says this, he or she is really saying that the device or network is 99.999 percent available. In fact, *99.999 percent availability* is a sure sign that the person is using the percentage method.

NOTE When you are actually doing the math of calculating availability, you are likely to see a number that looks like 0.99999, rather than 99.999 percent. Remember that you have to multiply by 100 to come up with a percentage, which is how most people are used to seeing availability described.

The essential use of the availability percentage is to figure out how much downtime you are going to have over a year-long period. You determine downtime by multiplying the number of minutes in a year by the percentage of availability. This gives you the minutes per year that you will be operational. The balance is the downtime you can expect.

Because there are 365 days per year, 24 hours per day, and 60 minutes per hour, we can calculate that there are 525,600 minutes per year. However, this does not account for leap years, which have an extra day. The way that we will account for leap years, since they happen every fourth year, is to add one fourth of a day to every year. This results in 525,960 minutes per year, which is the number that is used in all the calculations in this book. 525,960 minutes per year is important enough, in availability calculations, that you will likely have it memorized before long—as you become well versed in performing availability calculations.

In addition to the number of minutes per year, *annual reliability* should be understood. Annual reliability is the number of times each year that a device fails. When you know the MTBF for a device, you can divide that MTBF by the number of hours in a year (8766) to predict the average number of failures per year. We will be using this knowledge when we predict how many minutes a network is out of service while it switches from a broken device to a working device in a redundant situation.

Because we know the number of minutes in a year and because we now understand that availability is a percentage, we can calculate downtime for a year based on the availability number. Table 1-1 describes how the number of 9s relates to uptime and downtime.

Table 1-1 *Number of 9s; Uptime and Downtime*

Number of Nines	Availability Percentage	Minutes of Uptime per Year (Percentage * 525,960)	Minutes of Downtime per Year (525,960 – Uptime)	Annual Downtime
1	90.000%	473,364	52,596	36.5 days
2	99.000%	520,700.4	5259.6	3.5 days
3	99.900%	525,434.0	525.96	8.5 hours
4	99.990%	525,907.4	52.596	1 hour
5	99.999%	525,954.7	5.2596	5 minutes
6	99.9999%	525,959.5	0.52596	32 seconds

As you can see, for each 9 in the availability percentage, a significant increase in performance is achieved. It is often rumored that after the second 9, each additional 9 costs twice as much. That is to say, if you want to go from three 9s to four 9s, the amount of money you spend building your network is going to double! But remember, double the money buys you 10 times more availability.

The Defects per Million Method

The second way to state availability is by using the defects per million (DPM) method. Using this method, we describe the number of failures that have occurred during a million hours of running time for a device or a network. It is common to see this method used for existing large networks.

With the DPM method, we can report issues of reliability that the percentage method would have difficulty tracking. Because DPM is often used for existing networks, we can use it to measure partial and full network outages. We can also measure the million hours in terms of hours of operation of the network, the hours of operation of the devices (added together) that comprise the network, or perhaps even the hours of use that the users get from the network.

In order to clarify the DPM method, let us work through a couple of short examples. Assume that your network consists of 1000 hubs, switches, and routers. Assume that you count any degradation in performance as an outage. Also assume that we base our figures on 8766 hours per year (accounting for leap years) and that our failure reports are done monthly. To determine hours of operation per month, multiply 1000 devices times 8766 hours per year and divide by 12 to compute 730,500 hours of operation each month. As you can see in Figure 1-1, you compute one million hours divided by the number of operating hours to get the number of defects per million hours for a single defect. Then you multiply this result by the number of defects to get the total defects per million hours. So if we had two failures during a month, we would report that as 2.74 DPM for the month.

Figure 1-1 *Determining Availability Given Two Failures in One Month*

Hours per year = 8766 (Accounts for leap years)

Number of Devices = 1000

Accumulated Hours per Year = 8,766,000 hours

$$\text{Accumulated Hours per Month} = \frac{8{,}766{,}000}{12}$$

$$= 730{,}500 \text{ hours}$$

Converting 2 Defects per 730,500 hours:

$$\frac{1{,}000{,}000}{730{,}500} = 1.3689$$

$$2 * 1.3689 = \boxed{2.74 \text{ Defects per Million}}$$

Another way of reporting failures using the DPM method would be to base the million hours on network use. Let us assume that our network is large and constantly growing. Let us also assume that we bill our customers for each hour they use the network. Let us say that we have several thousand customers and that over the period of a month, they accumulate 1,200,000 hours of operation. In this network, if we had two failures, then the resulting DPM would be 1.67 DPM. Figure 1-2 shows you the calculations.

Figure 1-2 *Another DPM Calculation Example*

Total Network Hours = 1,200,000

Total Network Failures = 2

$$\frac{1 \text{ million}}{\text{Actual time}} = \frac{1{,}000{,}000}{1{,}200{,}000}$$

$$= .83333$$

$$\text{DPM} = .83333 * 2$$

$$= 1.67$$

MTBF, MTTR, and Availability

So far, you have learned about failures per million hours of operation and about downtime as a percentage. In order to have a feeling of the health of a network, relating the number of failures and the length of each failure to each other is helpful.

MTBF is a number that you have probably seen on product documentation or in some other specification for a product. It describes the number of hours between failures for a particular device. A similar term is Mean Time to Failure (MTTF), and it describes the amount of time from putting a device into service until the device fails.

Many people confuse these two terms. Using MTBF instead of MTTF and MTTF instead of MTBF in most situations makes very little difference. Technically, the mathematical equation for calculating availability that we will use in this book should use the term MTTF, according to the historic standards on the subject. You will nearly always be able to get an MTBF number about products you wish to purchase. Finding MTTF numbers about products is difficult. Technically, the companies stating MTBF about their products are very likely to be giving you MTTF and not even know it.

In keeping with this minor industry oversight and to simplify our lives, this book uses MTBF in place of MTTF and discards MTTF completely because it will make very little difference in our calculations. Before getting into an example that shows exactly how little difference it makes to switch MTBF and MTTF numbers, you have learn about MTTR and the availability equation.

MTTR is the amount of time (on average) that elapses between a network failing and the network being restored to proper working order. In most cases, MTTR includes a little bit of time to notice that the network has failed. Then it includes some time to diagnose the problem. Finally, MTTR includes some time to perform the appropriate action to fix the network and a small amount of time for the repairs to bring the network into proper working order. In an ideal world, the timing to detect, diagnose, and repair a network problem will be measured in minutes. However, sometimes things happen in the night and no one notices for hours. Sometimes the first diagnosis is wrong and several hours are wasted fixing something that isn't broken. The key point here is to remember that there are three phases to fixing a network problem:

- Detection
- Diagnosis
- Repair

You can calculate percentage availability directly if you have both the MTBF and the MTTR numbers for a particular device, as shown in the availability equation in Equation 1-1.

Equation 1-1 The Availability Equation

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

If we use Equation 1-1 to work a couple small examples, you can easily see how the difference between MTBF and MTTF makes no difference for our purposes.

Let us assume that a device we want to measure for availability has an MTTF of 200,000 hours. Let us also assume it has an MTTR of six hours. Let us further assume that the MTBF would be 200,006 hours, which is mathematically correct. By plugging all these numbers into the availability equation in Equation 1-1 and using a calculator, we get the two results, 0.9999700009 and 0.9999700017999. Now if we round these numbers, we get 0.99997 and 0.99997—the same. Figure 1-3 shows the work behind this.

Figure 1-3 *Comparing MTTF and MTBF for Availability*

$$\text{MTTF Availability} = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}}$$

$$\text{Our MTBF Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

$$\text{MTTF Availability} = \frac{200{,}000}{200{,}006}$$

$$= 0.9999700009$$

$$\text{Our MTBF Availability} = \frac{200{,}006}{200{,}012}$$

$$= 0.9999700017999$$

Both Equations Round to 0.99997

Relating the Percentage and DPM Methods

With MTBF and MTTR, it is possible to convert between the percentage and DPM methods. Since availability is a percentage and DPM is not, we need MTTR in order to do the conversion. MTBF will help because we use it to arrive at the percentage availability number. The best way to illustrate the conversion is through an example.

Assume that we have a network that is special because all the routers and switches in it have the exact same MTBF: 200,000 hours. Furthermore, assume that the MTTR for these devices is six hours because a support contract guarantees the replacement of any failed device in six or fewer hours. Using the availability equation, we can easily arrive at the same answer we got for these figures in the preceding section—an availability of 99.997 percent. Taking 99.997 percent and multiplying by the number of minutes in a year, we get the answer—15.78 minutes per year of downtime per device.

In order to predict annual downtime using the DPM method, we determine the total amount of downtime over one million hours and then convert that to annual downtime. First we divide 1,000,000 by 200,000 (MTBF) to get 5, which is the number of defects per million. Then we multiply that by the downtime (6 hours) to get the total downtime per million hours (30 hours). To convert to downtime per year, we simply calculate the ratio of "hours in a year" over one million hours, and multiply that by our 30 hour result. This produces .26 hours or 15.78 minutes, which matches our annual downtime result using the percentage method.

Although the conversion between percentage and DPM methods can be done, it usually is not. The DPM method is most often used to measure existing network performance. The percentage method is most often used to predict network performance based on a design. The reasoning for this is based on simplicity of use and the ease of getting the data to put into the equations.

For the percentage method, we simply take the MTBF provided by the manufacturer of the product or products in question. Then we estimate the MTTR based on the contract we sign with the provider of the network. From this we can estimate the predicted availability of the network or network device.

Figure 1-4 *Relating the Percentage and DPM Methods*

$$\text{MTBF} = 200{,}000 \text{ hours}$$

$$\text{MTTR} = 6 \text{ hours}$$

$$\text{Availability} = \frac{200{,}000}{200{,}006}$$

$$= 0.99997$$

$$\text{Annual Downtime} = (1 - 0.99997) * 525{,}960 \text{ minutes per year}$$

$$\boxed{\text{Annual Downtime} = 15.78 \text{ Minutes}}$$

$$\text{DPM} = \frac{1{,}000{,}000}{200{,}000}$$

$$= 5$$

$$\text{Downtime per Million Hours} = 5 * 6$$

$$= 30$$

$$\text{Failures per Year} = \frac{8766}{1{,}000{,}000}$$

$$= 0.008766$$

$$\text{Downtime per Year} = 0.008766 * 30$$

$$= .26 \text{ hours}$$

$$\boxed{\text{Annual Downtime} = 15.78 \text{ Minutes}}$$

For the DPM method, we accumulate millions of hours of run time (by machine or by use) and then note each failure that occurs.

Either way creates a scenario by which we can determine and compare performance over time or performance from one manufacturer to the other.

This book uses the percentage method because we are mainly concerned with answering the question: How available will this network be if we build it according to this design? Because the answer to this question is a prediction based on available information such as MTBF and suggested service contracts, the percentage method is the simplest way to provide the answer.

Additional Considerations in Measuring Availability

Although considering every possible scenario is beyond the scope of this book, some additional thoughts about measuring availability are in order. The following sections introduce the concepts of partial outages and process failures (that is, human error). Partial outages are mostly omitted from this book as they are an advanced availability subject and incredibly difficult to predict. The method for including human error-induced downtime is introduced here and described in more detail in Chapter 4, "Factors That Affect Availability."

Analyzing Time Spent in Network Outages

If you decide that you want to measure your network's availability after it is built, you might also decide that you want to know more than just the uptime and downtime. You might also want to know how much time is spent in partial outages. You can arrive at this data in several ways, but the following way is my favorite.

First you must know the type and quantity of devices on your network. You must also decide which you want to measure. Each time your network has an outage of any sort, you must log it and categorize it. You need to collect the following information for each outage:

- The device (including model, serial number, hardware version, software version, etc.)
- The nature of the outage (complete failure or partial outage)
- The time the outage began
- The time the outage was recognized
- The time the outage was diagnosed
- The time the solution was implemented
- The time the network was fully restored to normal operation

From this information, you can derive plenty of statistical information over time that will enable the reduction of downtime.

The first thing is to calculate DPM by counting the number of failures (of a particular device type) against the accumulated operating hours (of that particular device type). Next, take the average length of the outages (in hours) to determine the MTTR for that particular device type. This gives you the basic data that enables you to compare performance over time for your network. In Chapter 4, in the human error and standard processes section, we will perform an example that should clarify these ideas for you.

Next, compare the different devices and make sure that only those devices with the best performance are used in your network. If you find that some routers have a higher failure rate than other routers, then you should figure out why or possibly switch to the more reliable model. If a particular router is not as reliable as you would have expected, you should contact the manufacturer. Cisco Systems, for example, has improved the reliability of its routers through both hardware and software upgrades based on this type of feedback from customers.

Another way to improve availability is to improve your processes for handling network failures. Within a particular type of equipment and outage, analyze each of the time segments between the failure and the restoration. You are looking for the average and the range for each of the segments. If a particular segment shows a large range, consider it suspect. Compare the fast restoration cases (within this time segment) to the slow restoration cases to see why they were different. The reason for this difference is obvious.

Here is simple example. Say you notice that when routers fail, the average time to recognize the problem is five minutes, but sometimes it takes nearly an hour. This leads to further research. Every time a router fails between 8 a.m. and noon and between 1 p.m. and 5 p.m., the problem recognition time average is five minutes and the range is from three to seven minutes. Every time a router fails from noon to 1 p.m., the recognition time average is 35 minutes and the range is from four to 61 minutes. This should tell you that at least one person should remain in the network operations center during lunch hour.

This same method also works when a particular problem is difficult and some of your staff need more training to detect or solve the problem at the same speed as your senior staff. Finally, this method might point out a problem where the manufacturer could improve the troubleshooting capabilities of the product.