

Remote Networking: Driving Down Costs, Driving Up Reliability

David Weiss, CEO
Dataprobe

Natural disasters, human error, malicious behavior, even the complexity of the systems themselves - all are obstacles that challenge us in trying to build resilient networks that meet our core objectives in an increasingly competitive marketplace. In order to build the most effective computer and communication networks, contingency planning is not only required, but essential. Networks with remote sites, however, present particularly unique challenges when undertaking contingency planning and need to be carefully addressed.

Key Terms in Contingency Planning

Mission Critical - The core activities within an organization that are required for survival; these activities generate revenue and represent the 'raison d'etre' of the organization. In order to know what is Mission Critical, it is important to first understand what the mission is. We do a lot of work with the Federal Aviation Administration, for example, and they are very clear about what their mission is: "To provide the safest, most efficient aerospace system in the world". A clear mission is key to helping drive what is Mission Critical.

Disruption - An unplanned event that interrupts the normal flow of a business functions for an appreciable length of time.

Disaster - When we think of disaster, we tend to think of the 'biggies': hurricane, tornado, fire, flood, earthquake, and now terrorism. But an event does not necessarily need to be large-scale or catastrophic to qualify as a disaster. If the function being affected is 'Mission Critical' and the duration of the disruption is unacceptable, then that meets the classic definition of disaster, no matter what the cause.

Availability - Availability is the measurement of success in achieving our uptime goals. It is the percentage of time that a system is ready to do its assigned function. This does not necessarily mean that it is always doing its function; just that it is ready if called upon.

Measuring Availability

The classic way to measure availability is to compare the time the system is working to the time that it is not. The working time is stated as the Mean Time Between Failure or MTBF. MTBF is expressed in a number of hours, and represents the Mean or Average time a system will perform continuously. A MTBF of 100,000 hours represents over 11 years of continuous operation.

The downtime is expressed as Mean Time to Return or MTTR; this does not mean the time it takes to repair the system but rather the time it takes to restore the system to readiness. There is a big difference between the two. MTTR is the length of time that it takes to get the system back up and running. It could be as simple as a reboot of a computer, about 30 seconds, or the complete rebuilding of a data center facility.

If a network card in a workstation breaks, it might take an hour to diagnose the problem, and swap out a new one. It might take three hours to run down to the store and buy another one, so the MTTR for the workstation in that case is an hour or so. The MTTR is one to three hours for the workstation. The NIC card itself is probably not worth repairing, so the MTTR on the card itself is forever.

In order to restore a system, it takes five key steps:

Detection - Discovering that a problem has occurred

Notification - Letting the right people know what has occurred

Response - Determining a course of action

Action - Repairing, replacing, rebooting the system, etc.

Recovery - Restarting the system and verifying its operational readiness

Each step takes time, each adding to the MTTR. The formula for Availability is the Uptime, or MTBF, divided by itself, Plus the Downtime, or MTTR.

If we have an MTBF of 100,000 hours and a MTTR of 30 minutes, the Availability is 100,000 divided by 100,000.5 or 99.999 percent. 99.999 percent is often referred to as 'five nines'- the holy grail of availability. AT&T's stated goal for its central office IES switch was five nines, or one day outage every 40 years. No switch was in service that long, but they worked pretty reliably. Other PBX manufacturers had similar availabilities, but the systems as a whole performed much worse. That is because the phones, cables, power supplies and other related gear were not included in those measurements.

When you look at a system and not just each component, the availability of each element is added together. So a system with 10 components, each with five nines of availability, all add together to achieve only 99.99 percent or four nines of availability. As systems get even more complex, the availability just keeps going down and down.

The standard measure of availability does not take into account that not all outages are alike. With this in mind, you will need to determine what is better (or worse) for your organization - one long outage, or several shorter ones? If you had to budget five minutes a year for downtime, would five one-minute outages be better than one outage of five minutes? The answer to that depends on the specifics of your situation, your core mission. You will also need to factor in if time of day makes a difference on the impact of an outage to your enterprise. Some activities are 24/7/365, some are not. Again, not all downtime is alike.

So, what can you do to prepare for the inevitable unplanned downtime that threatens our systems? Plan!

Contingency Planning Model

Business Continuity Planning (BCP) is a relatively new term, but the practice has been around at least as long as Noah's Ark. In a nutshell, it is the methodology used to create a plan for how an organization will resume partially or completely interrupted critical function(s) after a disaster or disruption. It is a structured way to assess critical processes and threats and to build a program of detection, notification, restoration and recovery that can be implemented immediately when a disaster or major disruption occurs. It is also a scalable practice that can be used at any level of an organization.

Business Continuity Planning is more comprehensive than Disaster Recovery, which primarily focuses on major natural disasters, such as hurricanes, tornadoes and earthquakes. BCP, on the other hand, encompasses outages due to legal and labor actions, loss of key personnel, shortages of critical manufacturing components; in short, anything that can adversely affect the core mission of the organization.

A methodology of contingency planning is clearly spelled out in the Contingency Planning Guide for Information Technology Systems, a highly informative document from the National Institute of Standards and Technology.

The guide lays out a sound seven-step approach to building a continuity plan specific to IT systems:

1. Develop the contingency planning policy statement. A formal department or agency policy provides the authority and guidance necessary to develop an effective contingency plan.

2. Conduct the business impact analysis. The BIA helps to identify and prioritize critical IT systems and components. A template for developing the BIA is also

provided to assist the user.

3. Identify preventive controls. Measures taken to reduce the effects of system disruptions can increase system availability and reduce contingency life cycle costs.

4. Develop recovery strategies. Thorough recovery strategies ensure that the system may be recovered quickly and effectively following a disruption.

5. Develop an IT contingency plan. The contingency plan should contain detailed guidance and procedures for restoring a damaged system.

6. Plan testing, training, and exercises. Testing the plan identifies planning gaps, whereas training prepares recovery personnel for plan activation; both activities improve plan effectiveness and overall agency preparedness.


7. Plan maintenance. The plan should be a living document that is updated regularly to remain current with system enhancements.

The Contingency Planning Guide for Information Technology Systems also defines three phases that govern the actions to be taken following a system disruption:

1. **Notification/Activation**. The Notification/Activation Phase describes the process of notifying recovery personnel and performing a damage assessment.

2. **Recovery**. The Recovery Phase discusses a suggested course of action for recovery teams and personnel to restore IT operations at an alternate site or using contingency capabilities.

3. **Reconstitution**. The Reconstitution Phase outlines actions that can be taken to return the system to normal operating conditions.



Downtime Causes

Hardware malfunction - Hardware malfunction is the number one cause of system downtime. This type of malfunction includes both IT systems themselves (hard drives, power supplies, communications equipments, etc.) and also hardware associated with infrastructure, such as HVAC and power systems.

Software malfunction - This type of malfunction includes not only operating system and application code crashes, but bugs that cause loss of function and also firmware bugs in imbedded systems, the microprocessors that run everything from printers, hubs, modems, etc.

Natural disasters - While they may wreak the most havoc, surprisingly, natural disasters only account for approximately three-percent of unplanned downtime.

Malicious acts - Malicious acts, which can include vandalism, employee misdeeds and malware (viruses, Trojan horses, spyware, etc.) are more prevalent than ever. Malware attacks are now targeting imbedded network devices like printers that were previously overlooked. Mobile phones now link to the Internet and can connect into corporate networks, prompting new threats. Even elevator control systems and vending machines are being targeted at an alarming rate.

Human error - Human error accounts for almost a third of downtime. This is not surprising since there are so many ways humans can make mistakes—from a janitor plugging in a vacuum to the wrong outlet, to an under trained employee reconfiguring a router.

The Cost of Downtime

The most important thing to remember is that contingency planning is also a human effort, and subject to the same levels of error as any other activity; constant training and testing will continually improve the process. In order to make a justification for the expense of contingency planning, and the remedies that stem from that planning, a clear understanding of the costs associated with downtime is required.

The formula for calculating downtime cost is: Frequency times Duration times the Cost per hour.

Frequency - The number of times that downtime is expected to occur, can be estimated from the Mean Time Between Failure.

Duration - The length of the outage or the Mean Time to Return.

Cost - The total cost of the downtime, in dollars per hour. Cost can come in many forms, each of which has to be taken into consideration:

Recovery Costs - the cost of replacing the damaged items over time

Revenue Loss - the revenue that is not generated during an outage

Productivity Loss - loss of productivity (machinery, employees, etc.) pre- and post-outage

Loss of Future Revenue - loss of long-term revenue, loss of market share and loss of opportunities

Loss of Confidence - loss of confidence from customers, partners, vendors, the investment community, stockholders and other key stakeholders

Loss of employment - loss of staff due to downtime (probably the most costly!)

Bottom Line: Downtime is Expensive

A 2005 study of 80 large organizations by Infonetics Research found that overall downtime costs averaged 3.6 percent of annual revenue. A Forrester Research survey found that almost two-thirds of respondents could not even provide an estimate of their downtime costs. Of those that did, 43 percent of companies estimated their downtime costs at 10,000 to \$100,000 per hour, and 7 percent of companies pegged it at over 1 Million dollars per hour.

Remote Site Perspective

Remote sites are more difficult to manage than traditional sites-especially in the area of contingency planning-for a number of reasons. First, by their very nature, they are remote; most likely in unpopulated areas, with great distances between them and the staff responsible for them. Remote sites may have difficult access, either because they are in difficult-to-reach locations, like mountaintops, or because they have extreme security requirements, like airports. These constraints make it difficult to access these sites in an expedient fashion for planned maintenance and certainly for unplanned downtime.

Second, remote sites may have a limited function, and therefore, a smaller investment in resources, which tends to be overlooked in the planning process. At remote sites, there may not be standby resources available to take over. Contingency resources might also be centrally located, or expensive to have in readiness, as there needs to be one spare at every location. In central sites, a small number of contingency resources can serve many production systems.

Finally, the functions of the remote site may be specific to the geography. The pump has to be where the well is. The cell tower needs to be located relative to the rest of the network. You just can't move the functions to a hot standby site. The fix has to be at the site itself.

Acceptable Downtime?

Downtime is inevitable. Even at five nines of availability, it is going to happen. But there are two ways to minimize downtime, and often one is severely overlooked. Often most- if not all- of the emphasis is placed on minimizing the likelihood of downtime; by building high availability systems. This, of course, is a sensible but expensive course. What is sometimes overlooked, is the recovery side of the equation, the MTTR.

Reaching just 99.5 percent availability costs about two and a half times more than standard systems. Reaching five nines of availability costs a lot more than that.

If you can reduce the duration of downtime, it can have as much a positive effect, at perhaps a substantially lower cost. Cutting MTTR to one tenth has the same effect on availability as multiplying the MTBF by a factor of 10. Going from one hour to six minutes of downtime might be a lot easier than going from 100,000 hours MTBF to 1 Million.

Minimizing Duration

There are a number of targeted initiatives that can help dramatically reduce the duration of downtime occurrences.

For starters, automating recovery operations can achieve tremendous improvements in shortening failure duration. Failover switching is designed to detect trouble in active systems and automatically move operations to hot standby systems. Protection switching reroutes communications equipment to diverse routed services from alternate providers. These automatic processes can be integrated into existing management schemes like SNMP, or be standalone dedicated managers, specifically for providing rapid recovery in the event of downtime.

Standby power sources like UPS systems, solar power and self-starting generators can provide non-stop operations in the event of utility mains power failure. Adequate surge protection is also important. Replacing a surge protector after a lightning strike is generally much quicker and easier to do than replacing the critical equipment it is meant to protect.

Automatic or remote reboot capabilities allow for a quick restart of failed equipment. Sometimes this is all that is needed to restore critical systems. It can certainly be the first line of defense in determining the severity of the problem. Reboot systems can be integrated into

both AC and DC power distribution schemes.

Accessing remote sites with console port access or remote KVM systems also allows basic troubleshooting, reconfiguration and troubleshooting without having to send a technician to remote locations.

The Human Element in Contingency Planning

The most important element of contingency planning is the human element. When things go wrong, having the right people in place who can make the right decisions and follow the right procedures, can make all the difference in the world. No amount of hardware can replace or adequately compensate for ill-trained, under managed staff.

Having a clear set of instructions in a readily available format and location is vital to getting staff focused on problem resolution with minimum distraction. In major disasters, some staff may be unavailable, and in general, people will be more concerned, and rightly so, with their families and loved ones.

Conclusion

Businesses today rely heavily on their remote networks to perform multiple critical functions. When downtime occurs it can have a profoundly negative impact on the organization, striking at the heart of employee productivity, customer satisfaction and corporate profitability. In order to optimize your network for maximum availability, it is essential to employ an optimal mix of planning and preparedness, cross training, redundancy in staff and advanced equipment. These measures will ultimately help to drive down costs and drive up reliability and ensure the longevity and long-term success of your organization.

About Dataprobe

Dataprobe (www.dataprobe.com) is a producer of disaster prevention, disaster recovery and remote site management technologies. Since 1969, Dataprobe has been designing, manufacturing and marketing solutions that provide maximum uptime, minimize downtime and provide efficient management of remote systems and facilities.

DON'T MISS A SINGLE ISSUE!



For a new subscription,
or to renew your
current subscription go to:

www.RemoteMagazine.com/r-sub.htm

REMOTE

Site & Equipment Management