High Availability White Paper

EQUILIBRIUM
IT SOLUTIONS, INC.

This document provides an overview of high availability best practices for mission critical applications.

Author: George Quinlan, Senior Consultant

Background - High Availability

Network availability (commonly referred to as “high availability”) is the design and measurement of a network in terms of the accessibility of the network services. The network must be available to access the services and applications running on it. High availability simply refers to the goal to keep the network available all of the time. Network uptime describes the availability of the network.

Benefits of designing networks for high availability:

- Prevents financial loss
- Prevents productivity loss
- Reduces reactive support costs
- Improves customer satisfaction and loyalty

Businesses measure their network downtime in terms of average cost per hour. For example, if a portion of a credit-card transaction network goes down such that businesses are unable to swipe credit cards for sales, the credit-card company might end up losing millions of dollars per hour.

Devices such as ATM machines, web services, and automated check-in machines at airports require constant availability. If these machines are down, then they cannot conduct business and revenue is affected.

Common terms for discussing availability are “24x7x365” and “five 9s”. The phrase “24x7x365” refers to keeping the network up 24 hours a day, 7 days a week, 365 days a year. This demand reflects several trends:

- Businesses are international. While people in the US sleep, their coworkers in Australia and Japan conduct business, and they need access to network resources.
- Web presence lets companies keep their shops “open” 24 hours a day.

Five 9s refers to the measurement of availability in terms of a percentage: 99.999 percent. This measurement implies that the network is available 99.999 percent of the time (and not available for .001 percent of the time). This type of measurement made sense in the mainframe world (where its use began) in which it measured a set of hosts. However, today's networks are distributed and consist of hundreds and thousands of devices. In this case, DPM (defects per million) hours of operation is a more realistic measurement, or the number of defects per million hours of operation.

In terms of availability, the following table shows how the measurement translates into downtime per year.

Availability	***DPM***	***Downtime per Year***
99.000%	10,000	3 days, 15 hours, 36 minutes
99.900%	1000	8 hours, 46 minutes
99.990%	100	53 minutes
99.999%	10	5 minutes

Five 9s availability means that the network is unavailable for a total of only 5 minutes per year! How do you design a network such that devices and applications are always working?

First, what contributes to unavailability of the network?

- Human error
- Failed devices
- Bugs
- Power outages
- Service provider outages
- Natural disasters
- Backhoes
- Acts of war or terror

- Upgrades, scheduled maintenance, or hardware replacements

Note that most of these examples are unplanned and thus generally outside of the control of the network administrator. Human error tends to be the leading cause of network outages. It is the design of the network that allows (or prohibits) the network to be available during these planned and unplanned outages.

Best Practices for Avoiding Downtime

Reaching any sort of constant uptime does not happen if the following factors exist:

- Single points of failure
- Outages resulting from hardware and software upgrades
- Long recovery times for reboots or switchovers
- Lack of tested spare hardware on site
- Long repair times due to a lack of troubleshooting guides, a lack of training, available external support and maintenance contracts
- Excessive environmental conditions
- Redundancy failure (failure not detected, redundancy not implemented)
- High probability of double failures
- Long convergence time for rerouting traffic around a failed trunk or router in the core

Because outages do (and will) occur, the goal of the network administrators is to reduce the outage to as short a time as possible.

The following design practices increase network availability:

Concept	*Example*
Hardware redundancy	You achieve redundancy with redundant hardware, processors, and line cards; devices acting in parallel; and the ability to hot-swap cards without interrupting the device's operation.
Software availability features	Availability features include Hot Standby Router Protocol (HSRP), nonstop forwarding, spanning trees, line-card switchovers, fast route processor switchovers, and nondisruptive upgrades.
Network and server redundancy	Redundant data centers mirror each other so if one data center (with it's servers, applications, databases, and networking gear) becomes unavailable, the network automatically reroutes to a redundant data center with minimal loss.
Link and carrier availability	Carrier availability comes from multihoming servers, multiple connections between switches and routers, and subscriptions to several different service providers.
Clean implementation, good documentation, cable management	You can take steps to minimize the chances of human error. Cleanly implementing a network (by labeling cables, tying cables down, using simple network designs and up-to-date network diagrams, etc.) helps prevent human error.

Backup power and temperature management	Using uninterruptible power supplies (UPSs) on primary network and server equipment ensures that when the power goes out, you have an alternative power source to keep the devices operational. Keeping devices in temperature-controlled rooms ensures that extreme temperature and moisture do not contribute to an outage.
Network monitoring	Monitoring the network, servers, devices, and applications allows systems administrators to detect problems or outages quickly, which minimizes downtime. The goal is to detect problems before they affect the network's ability to pass traffic. Admins typically use network-management software to monitor the network as well as detect trends.
Reduction of network complexity	Selecting a simple, logical, and repetitive network design over a complex one simplifies troubleshooting and network growth. It also reduces the chance of human errors. This step includes using standard released software, well-tested features (as opposed to bleeding-edge technology), and good design sense.
Change control management	Change-control management is the process of introducing changes to the network in a controlled and monitored way. This step includes testing changes before moving them onto the production network, researching software upgrades for known bugs, making a back-out plan in case a change causes a failure, and making one change at a time.
Training	Nothing is more important than a properly trained staff. This step significantly reduces human error by eliminating mistakes made out of ignorance.

At-A-Glance-High Availability

A highly available network means that a network and its applications are both operational and accessible at all times. As more businesses use networking to conduct day-to-day business, networking becomes critical to business. To put this in perspective, look at the cost of a network outage. The following numbers reflect the cost of one hour of downtime for various business functions:

- ATM Fees: $14,000
- Package shipping: $28,000
- Teleticket sales: $69,000
- Airline sales: $89,500
- Catalog sales: $90,000
- Credit-card authorization: $2.6 million
- Brokerage operations: $6.24 million

Designing a network for high availability does the following:

- Prevents financial loss
- Prevents productivity loss
- Reduces reactive support costs
- Improves customer satisfaction and loyalty

What Affects Network Availability?

The following three types of errors are the most common causes of network failures:

1. **Operational errors** account for 40% of network failures; they are usually the result of one of 3 factors:
    a. Poor change-management process
    b. Lack of employee training
    c. Lack of good documentation.
2. **Network failures** account for 30% of network failures, and they include single points of failure.
3. **Software failures** account for 30% of network failures. They can be caused by software crashes, unsuccessful switchovers, or latent-code failures.

How Do You Measure Network Availability?

The two most common methods for measuring availability are “number of 9s” and defects per million (DPM). Number of 9s refers to the measurement of availability in terms of a percentage (and not available for the other percent of the time). It is important to emphasize that as we move up the scale, the amount of resources and expenditure increases quite significantly, and having redundant servers is helpful, but money and resources must be spent in every “best practices” area to get to the highest levels of availability, and we will not get to 100% or even 99.999% by simply clustering servers, or adding a second data center.

Availability	***DPM***	***Downtime per Year***
99.000%	10,000	3 days, 15 hours, 36 minutes
99.900%	1000	8 hours, 46 minutes
99.990%	100	53 minutes
99.999%	10	5 minutes

Best Practices

Hardware redundancy means redundant hardware, processors, line cards, and links. You should design the network such that critical hardware (e.g. core switches) has no single points of failure. Hardware availability also allows you to hot-swap cards or other devices without interrupting the device's operation (online insertion and removal).

Reduction of Network Complexity

Although some redundancy is good (and necessary), overdoing it can cause more problems than it solves. Selecting a simple, logical, and repetitive network design over a complex one simplifies the availability to troubleshoot and grow the network. There is a trade-off between expense and risk. A good design maintains the proper balance between the two extremes.

Software Availability

Software availability refers to both reliability-based protocols, such as Spanning Tree Protocol (STP) and HSRP (Hot Standby Router Protocol), and reliable code and nondisruptive upgrades. STP, HSRP, and other protocols provide instructions to the network and to components of the network on how to behave in the event of a failure. Failure could be a power outage, a hardware failure, a disconnected cable, or any number of things. These protocols provide rules to reroute packets and reconfigure paths. Convergence is the process of applying these rules to the resolution of any such network errors. A converged network is one that, from a user standpoint, has recovered from a failure and can process instructions and requests.

You should thoroughly test and use software in a real (quarantined) or simulated real environment before putting it on the network. Avoid "bleeding-edge" or inadequately tested code. You should also follow procedures for introducing new or updated code. Shutting down the network, loading new code, and hoping it all works is usually a bad idea. You should first introduce new code on segmented, noncritical parts of the network. Plan for the worst case when loading new code.

Link and Carrier Availability

Another component in building highly available networks is understanding your service provider's plans and policies for network availability. For business-critical applications, it might be worthwhile to purchase a secondary service from an additional service provider. You can sometimes use this second link for load sharing as well.

Clean Implementation and Cable Management

This best practice might seem like a waste of time when first implementing a network, but disorganized cabling and poor implementation can increase the probability of network disasters and can hinder their timely resolution.

You can save time, grief, and money by taking some simple steps, such as labeling cables, tying cables down, using simple network designs, and keeping up-to-date network diagrams. Good documentation and clean implementation leads to quicker problem resolution.

Systems / Network Monitoring

Monitoring the network servers and devices allows network administrators to detect problems or outages quickly, which contributes to minimizing network downtime. The goal is to detect problems before they affect the network's ability to pass traffic.

Network and Server Redundancy

Redundant data centers mirror each other such that if one data center (with it's servers, databases, and networking gear) becomes unavailable, the network automatically reroutes to a redundant data center with minimal data loss.

Training

Nothing is more important than a properly trained staff. Up-to-date, comprehensive training can significantly reduce failures. Human error will always exist, but you can limit that exposure through documentation, good design practices, and training. Human error can be forgiven; ignorance cannot.

Change-Control Management

Always expect the worst when first installing upgrades. It saves you time and trouble in the long run. Introduce all changes to the network in a controlled way. This method includes testing

changes to the network before moving them onto the production network, researching software upgrades for known bugs, and following a backup plan in case the change causes a failure or doesn't implement correctly.

Conclusion

Designing a highly available system is complex as there are many design elements to consider. A common approach is to leverage third-party experts for one-time projects to ensure the project is completed right the first time.

About Equilibrium

Founded in 2004, half of Equilibrium is dedicated to providing IT project consulting and the other half dedicated to providing IT services. Equilibrium specializes in IT strategy, security, cloud, datacenter upgrades and ongoing operations.

Visit us at EQinc.com

Request a Free Consultation: http://eqinc.com/request-a-free-consultation

For more information: info@eqinc.com, 773-205-0200