



Next-Generation Networks

In today's always-connected environment, corporate networks must be able to provide critical applications and business services without interruption. Learn how resilient networks are able to deliver on the promise of high availability.

Table of Contents

- 1 Executive Summary
- 2 What is a Resilient Network?
- 3 Planning for a Resilient Network
- 4 Considerations for High Availability
- 6 Designing a Resilient Network
- 6 Managing and Growing the Network
- 7 CDW: Your Partner for Resilient Networks

Executive Summary

Businesses are more dependent than ever on their network infrastructure. The Internet has enabled many organizations to go global and operate 24x7x365. Employees, customers and partners have expectations that the information and interaction they require from the network will be available to them whenever they need it, wherever they are.

For IT and networking professionals, this means building networks with the goal of eliminating downtime and providing uninterrupted access to the business resources supported by the network — applications, websites, databases, you name it. When thinking about network architecture, it is critical that IT decision-makers understand which applications and business processes are most important in terms of application uptime. This requires careful analysis, planning and network design.

Fortunately, today's next-generation technology enables network resiliency that offers unprecedented levels of high availability for your most important IT and business resources. By designing, building and deploying networks with high levels of redundancy, IT and network professionals can ensure a highly available network and protect their organizations against the lost business and diminished productivity that results from network downtime. This paper discusses the importance of high availability, resilient networks and how to go about planning and designing them based on the needs and requirements of your business.

What is a Resilient Network?

The goal of any network design is to eliminate downtime and ensure the availability of critical applications to employees, partners and customers. Resiliency is the ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation.¹ In network design and planning, resiliency is built on both the combination of hardware and software used and the methodology employed to provide uninterrupted access to applications and services during network disruptions or hardware failures.

One of the key goals of any resilient network is to eliminate single points of failure by creating multiple pathways, through the LAN and WAN, as well as the Internet. Resiliency can also be achieved by building redundancy into the network hardware: Downtime is often reduced and minimized through the utilization of backup processors, switches and routers. The last thing you want is the entire organization to be at risk because a router or switch suffers an outage. In addition, many of today's networking hardware products are available with hot-swappable components, such as power supplies and fan trays.

The key in building and designing a resilient network is understanding what your organization means by "an acceptable level of service." In most cases, IT and network design professionals discover that an acceptable level of service is actually a moving target — what may be acceptable for one part of the organization may not be acceptable for another. For instance, employees in the accounting department at headquarters may go home at night and would never notice an interruption in service between 3 a.m. and 4 a.m. local time. However, if the same company is conducting global online sales, then any downtime of an hour — at any time of day — could result in millions of dollars in lost sales opportunities.

High availability in a network is typically specified as a percentage of time a network is available, expressed as a number of nines²:

- Three nines (99.9 percent) — 10 minutes of downtime per week
- Four nines (99.99 percent) — 1 minute of downtime per week
- Five nines (99.999 percent) — 6 seconds of downtime per week

Not every network and application is going to require five-nines availability — most do not, in fact. Not every IT or network professional will be charged with accomplishing that task. However, focusing on network resiliency and application uptime is more important than ever in today's business environment. The growth of the Internet and the accompanying ubiquity of both access and devices mean that many businesses are operating globally and are dependent on their networks for revenues, customer service, product procurement and other mission-critical functions that could be crippled if service were unavailable — even for seconds.

Even with a network of three nines, you can expect that your organization will have 500 minutes of downtime a year — more than eight hours. Calculators that attempt to measure the true cost of downtime focus on a wide range of factors, including lost revenue, labor costs, service outages and number of people affected. Downtime impacts goodwill, and can also result in fines and penalties and force companies to spend money on public relations to recover from significant outages. Some downtime can be planned for, but unplanned outages can be devastating. As seen in Figure 1 (Typical Hourly Cost of Downtime by Industry), a single hour of downtime can cost an organization millions of dollars, depending on how dependent it is on the network.

Figure 1. Typical Hourly Cost of Downtime by Industry (in U.S. dollars)

Brokerage Service	6.48 million
Energy	2.8 million
Telecom	2.0 million
Manufacturing	1.6 million
Retail	1.1 million
Healthcare	636,000
Media	90,000

Sources: Network Computing, the Meta Group and Contingency Planning Research.

¹Reliance and Survivability in Communications Networks: Strategies, Principles and Surveys of Disciplines, March 2010

²How CAPNet Achieved High Availability, Cisco IT Best Practices, 2008

Best Practices in High Availability

While most applications and networks don't require five-nines availability, there are clear benefits to achieving five nines and clear business challenges that require it — think Amazon, eTrade or other organizations that conduct essentially all of their business on the Internet. Cisco is another company that has strived for five-nines capability on its CAPNet, which is a global backbone network that serves tens of thousands of Cisco employees and subcontractors at 15 locations on four continents.

According to documents released by Cisco, the goal for CAPNet is "five nines or better." The company notes that even a brief outage can disrupt the workflow of thousands of employees. CAPNet utilizes Gigabit Ethernet, OC12, OC3 and DS3 circuits to connect sites in major cities in Europe, North America, Asia and Australia. Each site has at least two diversely routed circuits on redundant hardware to maximize fault tolerance. CAPNet bandwidth ranges between 45Mbps to 622Mbps, depending on the location.

This is Cisco's take on what is needed to achieve high availability and its benefits to the organization:

"To achieve five nines, enterprises must assess and improve network design, operations, management and support. The payoff is big — a highly available network that decreases operating costs, increases employee productivity, streamlines supply chain activities and provides the infrastructure necessary for applications such as videoconferencing and IP communications. In many respects, the road to high availability is paved with good operational practices and common sense. The adoption of best practices can often increase operational overhead in the short term but, if implemented properly, will reduce overhead significantly in the long term.

"Focus on the areas of network design, configuration, monitoring and alerting, on-call support procedures and documentation. Without focusing on these areas, operational excellence is difficult to achieve and even more difficult to maintain. When starting a high-availability program, look for opportunities to implement a few easy improvements first. Remember, however, that you can't get to five nines using technology and standardization alone. Many failures occur because of weak operating, management and support procedures.

"The benefits of high availability far outweigh the costs. Highly available networks improve company image, decrease operating costs, increase employee and vendor productivity, and support modern IP communication applications, such as videoconferencing and voice over IP."²

Planning for a Resilient Network

For many businesses, planning for a resilient network should start in the office of the CEO or CIO, either literally or at least figuratively. The reason for this is that building resiliency into the network is an investment that is often offset against the threat of lost sales, lost opportunities or, just as important, losses to the organization's reputation and brand. It is the business decision-makers who have to define where the organization can least afford to be vulnerable, and it is up to IT to inform them of just how vulnerable the network might be and just how much those vulnerabilities could cost the organization. For IT and network professionals, it is important to understand which applications are most critical and can least afford downtime, and to then find out the level of uptime required. For example, companies that do business over the web or generate advertising revenues on the web will lose money every second their networks are down. These would be typical applications for a high-availability, resilient network.

Once the IT department understands which applications are most critical, it then has to assess where they are vulnerable. At this point, it is important to bring on a third-party partner, such as CDW, that has experience in building multivendor resilient networks. A partner can help with the next stages of planning, which includes performing an accurate assessment of the current environment and a gap analysis to determine if the infrastructure, sites and production environment can scale to include a new, resilient infrastructure. A typical planning process would include:

- **Conducting a Network Audit:** An audit provides an inventory of all the hardware and software components within the organization and identifies missing product patches, multiple software versions, product end-of-life or end-of-support deadlines and any factors that may indicate possible security risks or network performance issues. You also need to consider current applications and data on the network, such as VoIP, e-mail, SQL, Internet and video-on-demand.
- **Performing a Network Assessment:** An assessment maps all network devices, links and protocols, and provides a detailed view of the complete LAN and WAN design. It can identify critical network issues and provide a remediation to address them. An assessment looks at current network topology, including network devices, physical and logical links, external connections, frame types, routed and routing protocols, application-specific themes and IP

²How CAPNet Achieved High Availability, Cisco Systems Inc., 2008

addressing schemes. It should also include a traffic and network utilization analysis.

- **Developing a Strategy Plan:** A strategy plan considers all of the organization's operational needs to ensure the network has the flexibility to support future growth. You need to evaluate timelines, availability expectations and prioritization of each project component.

Considerations for High Availability

The primary reason for high-availability, resilient networks is to maintain uptime; loss of service, slow or unresponsive service results in lost revenue and/or lost productivity. Resilient networks provide other benefits as well, including the ability to maximize bandwidth utilization and increase levels of management, measuring and reporting. In addition, the network design could account for failover sites to enhance disaster recovery and business continuity.

In designing networks for high availability — both WANs and LANs — IT and network professionals may incorporate a variety of components depending on the types of services they are providing on the network, such as data, voice and video. Designing the network is not just about choosing the right hardware and software. It's also about deploying a methodology that configures the right features and functions to provide uninterrupted access to applications and services. Among the key considerations are creating multiple active paths for data; redundant network components (internal or virtual to the device, maybe the same physical hardware); deploying load-balancing software; and utilizing network management, monitoring and measuring tools to identify, diagnose and fix points of failure. Another important point is to look at your sources of power and determine how to limit your vulnerabilities: Do you have a generator? Do you deploy UPS devices? Are you using features such as Power over Ethernet (PoE)?

WAN Considerations: Most organizations have a primary path for their WAN, which could be frame relay or, more likely, Multiprotocol Label Switching (MPLS). MPLS is the direction most organizations are taking because it allows for Layer 3 routing protocols to be dynamic, which lends itself to high availability. MPLS also allows for quality-of-service (QoS) settings to be honored when prioritizing traffic during times of congestion. There are several options for designing WANs for creating multiple paths. You can create a second private link for an alternative provider to split traffic, and you can also split

traffic with a second router. Some organizations, in fact, insist on having two or more carriers so if one carrier's network goes down, there is still a backup path.

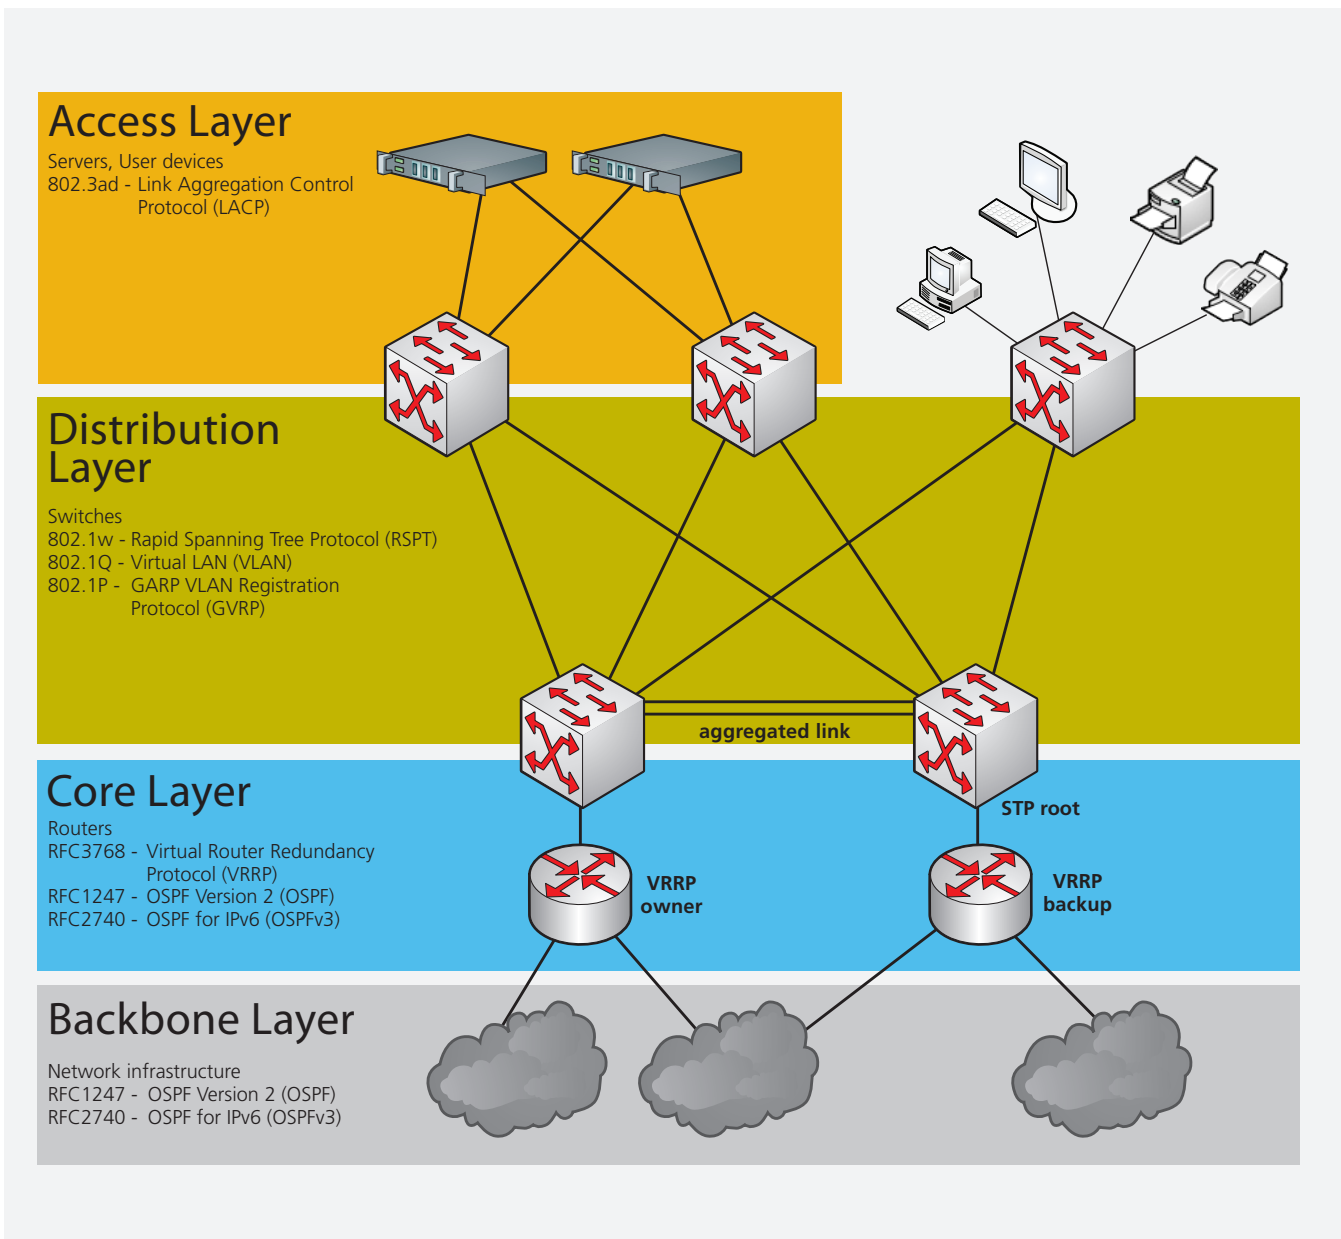
In general, a WAN application delivery solution should include three components, with the first providing the most considerable impact:

1. **WAN accelerators:** These network devices use application acceleration, WAN optimization, web caching, and QoS and traffic shaping to deliver LAN-like performance over the WAN. Network hardware vendors, including Cisco and Riverbed, offer WAN accelerators purpose-built for branch and remote offices of all sizes.
2. **Load-balancing software:** This optimization software intelligently prioritizes and accelerates the delivery of critical applications hosted on multiple servers. It monitors and distributes traffic according to the organization's prioritization of the type of data, IP address and destination.
3. **Application delivery controllers:** These network devices offload server processing of web-based applications from web servers and may include load-balancing software plus advanced routing capabilities, system health monitoring and other tools for optimizing the network.

LAN Considerations: Many organizations are making their LANs more highly available by building in resiliency among the core components — for example, building in redundant supervisors, power supplies, fan trays and other methods at the router, switch and chassis levels. Routing protocols provide for an inherent level of resiliency because they are designed to route around problem areas, although there is a time lag involved in establishing a new path after a failure. In addition, resiliency can be achieved by using virtual devices within the physical devices at the core level. The goals in network design include providing constant availability, enabling upgrades on the fly and eliminating single points of failure. Current trends in LAN design are bringing redundancy and resiliency to both the network core and the edge.

Throughout the network, components such as routers, switches and chassis should include multiple points through which data can travel to the network to users. Load-balancing software is critical because it monitors and distributes traffic on the network based on defined prioritization of the type of data, IP

Basic Components in Resilient Network Design



address and destination. You also want to incorporate management and monitoring software that will be able to identify and diagnose failed network connections from a central location, and then repair and/or replace them without disruption to data availability.

One of the significant benefits of having a resilient network is the ability to deliver overall performance enhancements and higher levels of availability to support your organization's com-

pute resources, particularly as technologies such as virtualization and cloud computing become more widely used. Many organizations are striving for high availability and redundancy in their data centers to enhance business continuity, reduce costs and improve ROI. Computing technologies such as virtualization and software-oriented architectures, for example, require high-availability networks to enable the exchange of data regardless of operating systems or programming languages.

Planning for High Availability

As more and more IT departments build redundancy into their data centers, they are finding that it is equally important to look at their network infrastructure for redundancy and high availability. A recent report by Realtime Publishers states that high availability is not about trying to fit business needs into a solution, but is about having the solution fit the needs of the business. The report suggests using an availability scale, with availability characterized as follows:

- **Unprotected:** IT services that are not part of any high-availability solution; these may use some kind of redundancy, but by and large, they are unprotected
- **Reliable:** Application and services or hot-swappable components that can eventually be recovered; these do not impact business
- **Recoverable:** Redundant infrastructure components that have some automatic recoverability built in; some downtime is acceptable
- **Highly Available:** Mission-critical IT services that drive the business; demand for availability is high, so these need high redundancy at the hardware, software and communications levels
- **Always (Continuously) Available:** Few, if any, of these IT services exist in most organizations, as they require absolute zero acceptable downtime; these are triple, even quadruple, redundant

Among the high-availability technologies recommended in the report are hot-swappable components, snapshots, replication and failover, high-availability clusters, synchronized systems and virtualization.³

Designing a Resilient Network

The network design must include specifications for availability, reliability, security, scalability and performance. Often, network engineers recommend designing a resilient network in modules. This allows an organization to provide the highest degree of resiliency by segmenting traffic and preventing a single point of failure. The basic components of a resilient network are the access layer, distribution layer, core layer and backbone layer, as seen in the diagram on the previous page.⁴

In this design, the access layer provides connectivity to end-user devices, such as servers, computers, laptops and printers. The distribution layer provides Layer 2 distribution through the switched network. Available in this layer are advanced features such as QoS, gateway redundancy and other key

features. Distribution can be achieved using any Layer 2 loop-free protocol, such as the Spanning Tree Protocol (STP), described below. The distribution layer connects to the core layer by utilizing multiple high-speed Layer 3 links in order to take advantage of high convergence of advanced routing protocols. Core switches and routers are placed in the core layer. This ensures Layer 3 availability of the gateway IP address when one of the routers goes down. In most cases, the routing and switching capability will be integrated into one device. The core layer provides access to the backbone layer and high-speed switching to other modules in the infrastructure. The backbone layer is a common part of the routed network. To distribute routing information, a dynamic routing protocol such as OSPF is used. The routing protocols support resilience and topology changes by default.

In designing resilient networks, it is critical to eliminate single points of failure, which means having redundant links to critical servers and network devices. In the network design shown in the accompanying diagram, all critical devices are used twice to avoid having a single point of failure. Therefore, any single device can be turned off without any significant disruption to the connected applications and users. Redundant links can create problems, however. For example, in Layer 2 switched environments, redundant links can cause switches to flood packets throughout the network, effectively halting the switching of production traffic. STP is a Layer 2 protocol designed to prevent such flooding by placing one of the redundant links in a blocking state. Although STP prevents Layer 2 loops, it is slow to converge. STP improvements, such as Rapid STP, help to decrease the convergence time.

At Layer 3, advanced routing protocols enable the highest level of network resilience when utilizing redundant links. Not only can advanced protocols load-balance traffic over redundant links, but they can also converge in a matter of seconds in the event of a primary link failure. Aggregate redundant links at Layers 2 and 3 are a common best practice to increase resiliency. Technologies such as EtherChannel combine switched or routed links into one logical link, effectively doubling the bandwidth on the link and minimizing convergence. Since the switch or router sees aggregated links as a single link, if one of the links fails, traffic continues to flow through the others.

Managing and Growing the Network

Another important consideration in building resilient networks is design scalability. The network has to be ready to take on

³The Essential Series: Making High Availability Pay for Itself, Realtime Publishers, January 2010

⁴Recommended Resilient Campus Network Design, Best Practices Document, CESNET, March 2010

new users, applications and services and be responsive to the needs of the business. One thing you know for sure: Business demands on data networks are changing constantly. VoIP has introduced voice to traditional networks, and new applications for streaming video, delivering high-definition TV and other high-bandwidth requirements mean that IT departments have to manage their networks for constant growth. While demands on the network could be simple now, that could change very quickly. The design of the network has to be capable of adapting, without losing the benefits of high availability.

When starting out and managing the network, it is important to understand what you want to achieve and make sure that goals are communicated clearly. The objective of the network, in terms of resilience, will affect the final design. What do you want in terms of acceptable downtime, the types of applications you are supporting — such as voice, video, HD TV, streaming media, HVAC and security (badge readers and cameras) — and the types of applications you want to enable across the network? Another important consideration is the stakeholders — not only management, but also end users. Often, the person working in the accounting department, for example, can communicate his or her needs better than anyone else.

CDW: Your Partner for Resilient Networks

Building a resilient network is a significant technical challenge because it typically involves multiple products from multiple vendors. It also demands a wide range of options for building additional transmission. CDW offers a wealth of experience, expertise and technical talent in helping organizations build the resilient network solution that is right for them. CDW also offers the opportunity for organizations to operate their network as a managed service, utilizing CDW Managed Network Services (MNS) to provide monitoring and management support for WANs and LANs. CDW can design, engineer, implement and provide ongoing network management support 24x7x365. Further, CDW will help you conduct a full assessment of where you are and where you want to go. The CDW approach includes:

- An initial discovery session to understand goals, requirements and budget
- An assessment review of your existing environment and definition of project requirements
- Detailed vendor evaluations, recommendations, future environment design and proof of concept

- Procurement, configuration and deployment of the solution
- Ongoing product lifecycle support

CDW's focus on networking insures that your network is properly designed and configured for your needs now and into the future.

About CDW

CDW is a leading provider of technology solutions for business, government, healthcare and education. Ranked No. 41 on *Forbes'* list of America's Largest Private Companies, CDW features dedicated account managers who help customers choose the right technology products and services to meet their needs. The company's technology specialists offer expertise in designing customized solutions, while its technology engineers and solution architects can assist customers with the implementation and long-term management of those solutions. Areas of focus include notebooks, desktops, printers, servers and storage, unified communications, security, wireless, power and cooling, networking, software licensing and mobility solutions.

CDW was founded in 1984 and as of March 31, 2010, employed approximately 6,150 coworkers. In the 12 months trailing March 31, 2010, the company generated sales of \$7.6 billion. Intently focused on responding to customers' technology needs with a sense of urgency, CDW helps customers achieve their goals by providing the right technology products and services they need — when they need them.