ECESSA®

Everything You Need to Know About Network Failover

Technology Brief

Overview

The Internet has become so pervasive and integral for conducting business and communicating with customers, partners and employees, that network performance, high-availability and uptime are now required for running the day-to-day operations of an organization. Network downtime not only costs money and impacts productivity, it can also adversely affect a company's reputation among customers and partners. For many companies, their entire business strategy depends on how well the network performs.

Many events can cause a network or site to go down, such as natural disasters, security attacks, a backhoe cutting a fiber optic cable or failing network infrastructure. Organizations must plan for, and allocate the budget to implement appropriate network infrastructure to ensure their data centers and remote offices have the protection, resiliency and redundancy they need in anticipation of disasters.

The cost of downtime is often considerable, especially when factoring in lost productivity. A recent IHS report indicates that costs of downtime fall into three categories: lost revenue resulting from the outage (17%), lost employee productivity (73%) and cost associated with fixing issues (5%). North American enterprises lose $700 billion in revenue per year due to network and communication technology downtime. The more distributed a company's network is, the more likely it is to suffer service-provider interruptions. According to the survey, retailers are affected the most, with service providers accounting for more than 30 percent of their downtime costs. Another cause of downtime is human error, which accounts for about one-fifth of costs. For financial institutions, this percentage jumps to nearly one-third.

Failover

Failover within a communications network is the process of transferring tasks from a failed component to a similar redundant component to avoid disruption and maintain operations. Automated failover is the ability to automatically reroute data from a failed component such as a server or network connection, to a functioning component and is essential for mission-critical systems.

Entire network architecture, or individual network components, may be configured as a cold spare (manual failover, requiring human intervention), hot spare or High Availability (HA) pair (automatic failover, with a short delay) or fully redundant (synchronized disaster recovery site). The three critical elements requiring failover configuration are power, network connectivity and server capacity.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

Device Failover

In a failover situation involving a device, such as a firewall, router, WAN controller, server load balancer, disk drive, web server, etc., data is transferred to the same type of redundant component to ensure there is limited interruption in data flow and operation. If a primary component becomes unavailable because of either failure or scheduled downtime, the secondary component serves as a backup and takes over for its failed counterpart; this configuration is referred to as a High Availability pair, or HA pair.

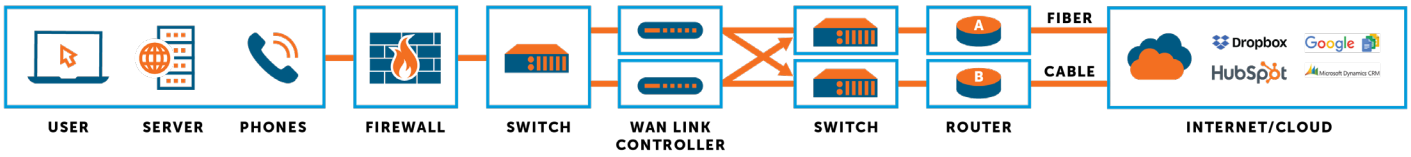

The capability to switch automatically to a redundant or standby system or network upon failure happens without human intervention (see Failover Hierarchy for other types of failover). Automated failover is essential in servers, systems or networks requiring continuous availability and a high degree of reliability — those responsible for mission-critical processes and data.

Failover Hierarchy

As mentioned earlier, some failover types are only partially automatic by intention and require manual intervention. This is called "manual failover"— failover is achieved with human intervention once approval is given and change control processes are followed. When hardware is a "cold spare," failover is time consuming and is susceptible to human error.

In contrast, when hardware is a "hot spare" or "High Availability (HA) pair," the failover mechanisms run in the background, so the transfer takes place automatically. The data on both systems is automatically synchronized. To the user, failover resembles a very fast automatic service reboot. However, the current transaction, or communication session, may be lost because it was not possible to synchronize the data prior to failure, or during the actual failover event. The system quickly recovers and normal system operations are restored without human intervention or approval.

The most reliable failover scenario, used for complete business continuity, is the use of a "Disaster Recovery (DR) site," where both network systems are permanently running in parallel — data on both systems is 100% synchronized at all times. Users will not be aware of any failures. This level of failover protection usually requires a corresponding modification to the client. To run both with systems in complete synchronicity, the connections to the client must be mirrored 100%. This normally requires clients that have connections with two or more servers at the same time, and can communicate with all of them. A normal web browser cannot do this.

Most enterprises implement a hybrid strategy for network resiliency that include some of each failover mechanism mentioned above. For example, a business may have a cold spare for a core Local Area Network (LAN) switch, a HA-pair for a firewall and WAN controller, with cold spares of most equipment at a disaster recovery site for manual failover if the need arises. Choosing the proper failover strategy depends on the cost of down-time and acceptable loss of productivity for that business and its customers.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

The Critical Role of SD-WAN

The convergence of voice, data and video over a single IP network is making the network infrastructure one of the most critical elements in operational success. These Voice over IP (VoIP), video and data services are increasingly integrated with business-critical applications such as email, Unified Communications as a Service (UCaaS), Customer Relationship Management (CRM), cloud-based storage, etc.

Therefore, all forms of communication with customers, suppliers and employees are inextricably tied to network operation. If the network fails, access to critical information can be lost or compromised, with potentially calamitous results. For example, an airport risks massive delays that impact passengers or patients' health may be compromised with a major medical center experiencing application delivery delays.

Software-Defined Wide Area Networking (SD-WAN) is a technology targeted at reducing and eliminating network connectivity down time and providing resiliency and failover for networks. Proper deployment of this technology can eliminate outages and provide for business continuity, regardless of the business size, geographic location, or software applications being used.

Who Needs SD-WAN?

- Small to Medium-sized Businesses (SMBs) need inbound and outbound load balancing and SD-WAN services for an increasing assortment of critical-business traffic, from VoIP to email. For example, the local corner store that does online banking and bill-pay over the Internet, or a manufacturing company that needs email, web services, hosted ERP and ecommerce applications available all the time.
- Companies with a central headquarters, branch offices and remote employees need secure and reliable data communications. They need reliable performance and high-availability of their VPN data, including the ability of the VPN connection to automatically failover if a WAN link goes down.
- Web hosting companies, MSPs, TSPs, MSSPs, ASPs and small ISPs need incoming link balancing and failover to ensure their services are reliable, with extra bandwidth and redundancy available to their servers. Their mission-critical applications need to be available 24/7/365. If a WAN link goes down, the failover process must be smooth and transparent to users.
- Many companies are deploying VoIP and Virtual Desktop Infrastructure (VDI) to cut expenses and enhance productivity. These companies now need Quality of Service (QoS) levels and traffic-shaping for guaranteed bandwidth to critical services and applications such as VoIP and VDI.
- Companies that have ERP, CRM, UCaaS or other software hosted in the Cloud or accessed over the Internet.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

The following are other critical elements that comprise a resilient network environment:

Power

With power failures being one of the most common reasons for network and systems failures, all critical network components at either the primary data center, call center or SD-WAN site must be connected to a power source that has very high-availability — 99.999% in the case of a data center.

A LAN that provides critical services such as a hospital or bank should be equipped with Uninterruptible Power Supplies (UPS) for each component of its distribution and access portions. These should be connected to emergency power sources to maintain internal communication. The WAN routers, switches, firewalls, etc. need the same form of protection to provide continuous communication and interconnection to external sites and other public networks.

Large data centers and critical operations, such as call centers, must rely on multiple electric power companies to provide utility power to their locations. The power is brought into the critical site from different geographical locations. If power is interrupted by an accident that severs electric lines at one location, the other utility can continue to provide uninterrupted power.

Emergency power generators may be used instead of alternate utilities. These generators, together with UPS equipment, can provide a continuous stream of electrical power for days if necessary, while utility power is being restored.

Network Redundancy

As referenced above, the level of network redundancy should be determined by the cost of down-time and acceptable loss of productivity for a business and its customers. Determining these values is a business analysis; what is the cost of a lost sale? What is the cost of employee productivity? Based on these answers, a comprehensive failover, resiliency and redundancy plan can be created. This plan should include considerations for redundancy of network devices such as switches, routers, gateways, firewalls, WAN controllers and ISP connections. There should also be consideration given to redundant components such as power supplies, CPUs and circuit cards for the network switches and routers. Refer to the Failover Hierarchy section for examples of each type of redundancy and a hybrid strategy.

Load Balancing

Many companies deploy SD-WAN solutions to merge WAN link load balancing and failover to cost-effectively ensure network performance and eliminate downtime for business-critical, time-sensitive applications. SD-WAN enables redundant WAN and ISP access and can provide both outbound and inbound WAN/ISP load balancing and failover. Load balancing with an SD-WAN device provides the ability to use all available bandwidth, from different carriers and technologies, in an active/active fashion creating more usable bandwidth while providing resiliency against outages with automatic failover.

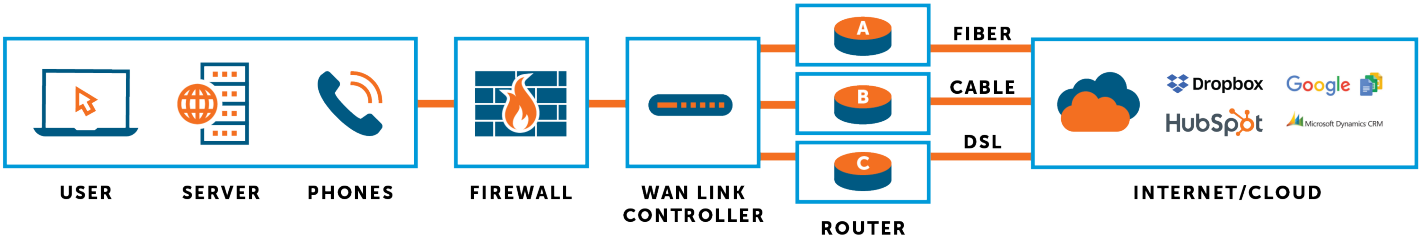

Consideration must be given to the redundancy and diversity of WAN links in conjunction with automated failover. Redundancy can be achieved by providing multiple links and multiple types of links for a single site and between multiple sites. For example, if the WAN network utilizes MPLS, it might be prudent to provide different links such as broadband, so that if a carrier's entire service goes down, the organization can have a backup strategy, which may include cable, satellite, cellular, microwave or fixed wireless service.

Link route diversity is a best practice whereby two or more links travel different routes to your locations. Another best practice is to use carrier diversity — multiple carriers — to provide Internet access diversity and redundancy.

WAN Bandwidth Capacity

Capacity factors must be assessed to avoid failures or congestion caused by unanticipated high traffic volumes when a failover occurs. The first factor is to assess the peak capacity coming from the primary link that failed, often the default link. The second factor is to assess the peak capacity of the secondary link where traffic will be rerouted to. The size of the WAN links should allow for both peak capacities, plus an additional 25-40% accommodating new peak traffic volumes.

Bandwidth should be ample enough to provide ISP failover and redundancy. If one link were to fail, you will still need enough bandwidth for users to be productive. Intelligent link load balancing monitors bandwidth availability throughout the network and priority-assigns traffic to the link with the greatest available bandwidth to guarantee time-sensitive traffic (i.e. voice, video and other critical applications) receive the bandwidth required for smooth performance. Most SD-WAN solutions also include WAN optimization features, like dynamic Quality of Service (QoS), that allow businesses to set bandwidth rules to prioritize business critical applications and contain non-critical traffic to avoid contention during a failover event.

In addition to the availability of WAN links, there is a need for a mechanism to connect users to available servers. If a server the user is connected to suddenly becomes unavailable, the load balancer redirects the request to one of the other replicated servers. This action causes the loss of the original session-to-credential mapping where the user is new to this substitute server, and is normally forced to login again.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

Packet-Level Bonding with SD-WAN

There are several levels of network resiliency, performance and control available with SD-WAN — from automatic failover and session load balancing with multiple ISP links, to packet-level duplication and aggregation (previously referred to as channel bonding) – to meet the needs of any business. Some businesses have mission-critical applications that cannot be allowed to fail, such as VoIP traffic for a call center, financial transactions at a bank, access to medical records at a hospital or clinic or flight operations at an airline. In instances like these, deploying an SD-WAN solution with packet level control is the only way to guarantee that no packets are ever lost or corrupted, maintaining application continuity at all times.

To achieve this level of performance and resiliency, SD-WAN devices are installed at both a local and remote site and direct traffic over the various WAN connections, combining (or bonding) bandwidth at the packet level. Each site connected by such a bonded link is assigned a unique identifier that allows it to be differentiated from other sites.

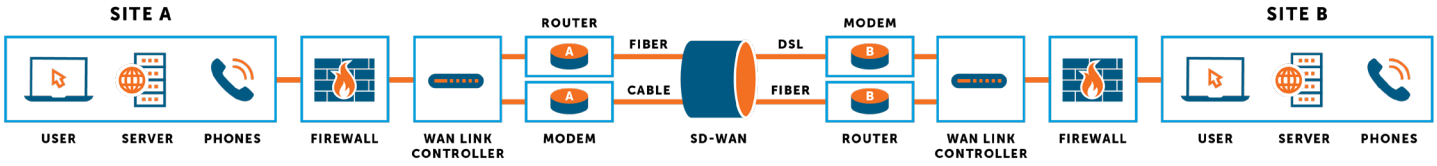

Each site is also configured with addressing information for both the local and remote end of the bonded link. This allows the SD-WAN device at each end to identify traffic that should be sent across the bonded link and direct it to the specified IP addresses on the WAN link(s) of the remote site. When the device identifies such outgoing traffic, it is disassembled at the packet level into separate streams of data, then encapsulated for transmission through the virtualized WAN (think virtual tunnels) and sent over all available WAN links. Since each encapsulated packet contains addressing information for a specific remote location, data is easily reassembled at that location. This finite level of packet control allows for more routing options, like true packet level duplication or aggregation. It also allows for encryption flexibility, allowing some traffic to be encrypted with AES-128 or AES-256 bit security per tunnel.

The benefit of this encapsulation approach over discrete VPNs is in the simplicity of management. When dealing with a large, multi-site network that needs mesh connectivity, the VPN combinations can number into the hundreds. With packet-level bonding with SD-WAN, each endpoint is a single IP address, making fewer discrete connections necessary. Also, with SD-WAN, the devices automatically maintain these encapsulated tunnels, with encryption, so that no IT administrator is required manage discrete VPNs.

For more information about SD-WAN functionality, please see Ecessa's technology briefs at www.ecessa.com/technology-briefs/.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

Summary

The Internet has become essential for conducting business and communicating with customers, partners and employees. Network performance, reliability and uptime are essential for running the day-to-day operations of many organizations. Network downtime not only costs money and loss of productivity, it can also adversely affect a company's reputation among customers and partners. Many events can cause a network or site to go down, such as natural disasters, security attacks, human errors and infrastructure failures. When evaluating how to avoid network failures, it is important to evaluate the many options available to ensure high-availability, network uptime and optimal network performance. It is also critical to examine solutions that will not only help avoid network failures, but are also operationally cost-effective.

Key SD-WAN Deployment Objectives:

- Improve network performance and eliminate downtime for business-critical, time-sensitive applications
- Globally manage WAN resources
- Provide redundant hardware failover and monitoring capabilities for mission-critical applications to eliminate all potential single points of WAN link failure
- Establish reliable network connections
- Ensure inbound and outbound traffic management over best performing WAN link
- Provide WAN link load balancing for ample bandwidth for critical applications
- Cost-effectively increase scalability and throughput of WAN connectivity
- Failover to secondary data center if all links at primary data center are down
- For enterprises connecting many locations, coordinate SD-WAN among all locations to provide uninterrupted access to data center, Internet and cloud-hosted resources for reliable performance of applications such as VPN, VoIP, UCaaS, VDI, CRM, Office 365 and more.

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®

Ecessa SD-WAN Solutions

Ecessa's phased deployment process ensures integration success and empowers network administrators to configure, manage and monitor their networks confidently with our cloud-based management tools. By offering a scalable range of both physical and virtual solutions that allow customers of all sizes and industries to take a progressive approach to SD-WAN at their own pace, Ecessa provides every business — at every evolutionary stage and IT budget — with a smart, sustainable path forward toward a wider, more perfect network.

Ecessa's Premises-Based Solutions:

- Contain their intelligence on an enterprise-grade router, so Internet connectivity issues won't have an impact on performance
- Do not require businesses to change IP addresses — use what you have
- Are highly configurable, allowing them to seamlessly fit into the most complex networks
- Work well in highly regulated industries (such as energy, banking, finance and healthcare) where end-to-end control of data is a priority and compliance is required
- Are carrier agnostic, leveraging any connection type: public (broadband, cable, XDSL) or private (MPLS), wired or wireless (satellite, cellular, line of site, etc.)
- Support a wide variety of bandwidths, ranging from 75 Mbps up to 20 Gbps
- Can be purchased as a monthly recurring charge or as a one-time charge to fit financial requirements
- Are software upgradeable, so the feature set can grow with your needs

7250 Metro Blvd #200 · Edina, MN 55439
Phone (763) 694-9949 · Toll Free (800) 669-6242
www.ecessa.com

ECESSA®