

Business Continuity and Disaster Recovery — the WAN is a Strategic Component



Worry-Proof Internet

2800 Campus Drive · Suite 140 · Plymouth, MN 55441
Phone (763) 694-9949 · Toll Free (800) 669-6242

www.ecessa.com

Business Continuity and Disaster Recovery — the WAN is a Strategic Component

Business Continuity is the plan and preventative mechanisms that help businesses remain up and running in the event of a disaster. Business Continuity is not something implemented at the time of a disaster. Rather, Business Continuity represents the activities performed daily to maintain service, consistency, and recoverability. Business Continuity is the activity performed by an organization to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions. Incidents may include local events such as building fires, regional events such as earthquakes, hurricanes, storms and floods or national events. From a network technology standpoint, Business Continuity can be affected by network connectivity failing, or a WAN link being congested or bottlenecked, which may limit business functions.

The foundation of Business Continuity is the policies, guidelines, standards, and procedures implemented by an organization. All system design, implementation, support, and maintenance must be based on this foundation in order to have any hope of achieving Business Continuity and disaster recovery.

Overview

IP-based Wide Area Networks (WANs) which comprise the Internet have become the default communication method for organizations in all vertical markets to conduct business through transaction-based applications and communicating with customers, vendors, partners and remote employees. While WAN infrastructure has become a ubiquitous communications method, it also presents many challenges associated with ensuring Business Continuity. While Internet Service Providers (ISP) continue to improve upon their ability to deliver consistent service, ISP outages continue to be a problem. As long as natural disasters, equipment failures, human error and security threats remain, ISP outages will continue to be an issue. Organizations that rely upon ISP connectivity need to take proactive measures to ensure the resiliency of their business, including their WAN infrastructure. Addressing secure and reliable WAN (Internet) access going out from the LAN, and coming into the LAN is critical for today's Business Continuity and disaster recovery planning.

The benefits of a sound disaster recovery plan are evident. During major hurricanes, thanks to IT facilities located outside of the disaster areas, businesses located in the Southeastern part of USA were able to recover quickly. The case for a solid Business Continuity and disaster recovery plan is made even stronger when you consider the diversity of today's workforce, with the increase in remote branch offices, road-warriors and telecommuters. The Internet, and the web in particular, have made it possible through the use of diverse technologies to enable secure access to business applications via VPN, or in conjunction with wireless networks. However, it is important to note that the resiliency of this approach depends upon the flexibility of the WAN infrastructure. Also it is important to note that if remote access to the main headquarters, remote offices or a disaster recovery site is removed, any amount of Business Continuity and disaster recovery preparedness at those sites will be of no use.

Business Transactions over the WAN Pose Major Challenges to Ensuring Business Continuity

Business Continuity planning covers many aspects of an organization's ability to avoid major business disruption from a disaster. For IT risk mitigation, the protection and prevention of lost data has been a principal concern. However, in recent years, Internet infrastructure and the systems that support it are taking a more prominent role in their ability to automate business processes and communication across an organization, their customers, and partners.

One of the easiest and most cost-effective ways of dealing with WAN reliability issues is multi-homing, using WAN Optimization Services with WAN link load balancing and failover. An organization will use a multi-homed network to bundle two or more WAN links and/or service providers to connect their LAN to the Internet. If they have multiple sites they can also use this technique to interconnect between sites to ensure reliability and optimum performance for critical applications.

This paper describes how small-to-medium sized enterprises can use advanced, yet affordable WAN Optimization Services to leverage the benefits of multiple sites, while maintaining high-performance and

reliability for applications delivered across all sites. It will also show you how to use WAN Optimization Services to manage, load balance, and deliver failover of multiple WAN and ISP links to ensure site-to-site network connectivity, and deliver applications reliability. Creating a network that supports disaster recovery plans to keep users connected to critical applications is the key to leveraging the full power of the Internet for Business Continuity and transaction completion in the face of a disaster.

WAN Infrastructure is Vital for Ensuring Business Continuity and Disaster Recovery

A major emphasis for Business Continuity is the protection of IT systems that enable and support critical business processes, applications and data. For virtually any organization today, computer networks are tightly integrated into business processes. For any organization that uses the network to conduct business and communicate with customers, partners and external employees beyond its LAN, the WAN has assumed an ever-increasing role in supporting the automation of business applications such as order fulfillment and communications using email and VoIP.

Many organizations with appropriate budgets are adding fault-tolerant sites (also called disaster recovery sites), and assigning the back-up hosting of their business applications to these sites with greater protection against disasters. Disaster recovery sites deploy replication technologies such as clustering to provide continuous service delivery to users, despite a major incident forcing downtime at the main facility.

The Internet (WAN) Challenge

Organizations that have increasing dependence upon WAN networks understand the importance of addressing Business Continuity and disaster recovery planning for IT infrastructure. While much of an organization's IT infrastructure may be internally owned and operated, and the organization has control over its Business Continuity, the external WAN is outsourced to an ISP or Telco. Utilizing an ISP for WAN connectivity places the Business Continuity in the hands of a third-party, leaving the organization vulnerable, with much less control over Business Continuity issues.

Business Continuity

It is very common for companies to use back-up applications to ensure that business transactions are completed smoothly. For example, when an order-placement system goes down, email can be used as a back-up to ensure that the order is fulfilled. When email goes down, the telephone can be used to place the orders. Businesses have an inherent back-up communication method they use to keep business running as smoothly as possible. When a WAN link or ISP link has an outage, there is nothing inherent in the WAN to backup the link. So, when the WAN link goes down, there is the potential for all of the applications and network services to be unavailable, which can result in a major disruption in business.

Below are examples of how businesses rely on critical applications being delivered over the WAN (Internet) for their ongoing operations.

- Employee productivity - Remote employees with web browsers accessing corporate applications with VPN or wireless network using handheld devices to access email
- Product sales - Resellers and agents purchasing, quoting and placing orders
- Order tracking - Shipping fulfillment and tracking through online services from FedEx, UPS, and others

Removing Risk, While Reducing the Cost of WAN Failures

Before making a decision on which solutions are appropriate, and how much budget should be allocated to mitigate the risk of an ISP or WAN link outage, it is important to understand what the likelihood of an ISP or WAN link outage would be, and the damaging impact upon a business.

ISP outages are very common, and will continue to occur as long as natural disasters, system failures, human error, security threats and service provider disputes continue. In a 2005 report, organizations incurred an average of 19 hours of ISP outages, and another 9 hours of WAN service degradation. This same report calculated that over the course of that year, these outages cost small-to-medium sized enterprises an average of nearly \$200,000 in lost revenue and productivity. Other events illustrate this point, and demonstrate the wide variety of causes for

WAN service disruptions.

- In December of 2004, Global Access Points, an ISP for Notre Dame University accidentally disconnected cables in Chicago that left the university without connectivity for 4 hours
- In September of 2005, during the Katrina hurricane disaster, most of the telecom facilities in the Gulf areas were wiped out
- In October of 2005, Cogent Communications and Level-3 had a business dispute they failed to solve. As a result, Level-3 disconnected its peering connection with Cogent. This dispute left customers of one provider unable to connect to customers of the other provider, and caused severe Internet slow-downs around the globe.

Email Outages

Recently, a survey was undertaken to understand the frequency, severity, and cause of email outages in North American corporations that use Microsoft Exchange, Lotus Notes, and Novell Groupwise. The survey results showed that enterprise email systems are prone to a variety of potential failures including SAN (Storage Area Network) failures, incorrect configuration, losses in network access, database corruption, and viruses.

Survey results showed that within a 12-month period, there is a 75-percent likelihood of an unplanned email outage within any given company. The length of email outages in the companies surveyed ranged from 2 minutes to 120 hours with the average email outage being just over 32 hours long. The largest concentration of outages (29%) was between 4 - 24 hours in duration. More than 43 percent of the outages lasted longer than 24 hours - a length of time that can lead to significant business disruption.

The majority of email outages were caused by unplanned events, most of which were due to technological failures (i.e. server hardware) which accounted for 35%, 19% were due to network connectivity losses, averaging 27.4 hours; 16% were due to SAN failures, and 16% were due to database corruption. While natural disasters accounted for only 14% of unplanned email outages, the average downtime due to such disasters was over 60 hours.

Service Level Agreements

ISPs Service Level Agreements (SLAs) only guarantee nominal financial reimbursement, and do not guarantee that a company's sites will remain up and running during a disaster. Therefore, it remains in the hands of the organizations themselves to protect their vital assets, and to ensure the resiliency of their WAN infrastructure. To ensure Business Continuity, the risk and cost from ISP link outages must be addressed. ISPs that offer SLAs to appease organizations will not suffice for the following reasons:

- SLAs don't guarantee up-time (they only stipulate reimbursement if up-time requirements are not met)
- SLAs usually only reimburse customers for the cost of the prorated connectivity that was lost - typically a fraction of the cost of the total damages

Dealing With the Risk of WAN Outages

One of the easiest and most cost-effective ways of dealing with WAN reliability issues is multi-homing (or link load balancing) by using WAN Optimization Services that are capable of doing link aggregation, load balancing and failover.

The value proposition that these services offer is that, while little control can be maintained over the continuity of service from a single ISP over a single WAN link, diversifying ISPs and WAN links provisioned over varied physical and logical paths can greatly reduce downtime. Establishing back-up links is not a new idea. However, what is new is the ability for an organization to deploy affordable WAN Optimization Services that provide organizations with the flexibility to easily provision low-cost bandwidth links to meet their specific needs, while providing WAN link redundancy with automatic failover. WAN Optimization Services can...

- Immediately detect ISP and WAN link failures, and automatically failover to an available link – with the transition being virtually transparent to users

- Utilize multiple and diverse link types such as T1s combined with broadband, wireless, and others to create a cost-effective yet resilient network
- Provide simultaneous utilization of all available links and available bandwidth (via link load-balancing), so that connectivity costs are not wasted on a underutilized back-up line
- Work independent of ISP peering relationships - ISP peering relationships and cooperation is not needed, and there are no problems with supporting different IP address spaces issued by different ISPs

WAN Optimization Services Compared to BGP

As discussed earlier, reducing costs is a critical part of planning a solid Business Continuity solution, which should be in line with the expected risks and resulting impact. However, the WAN is getting more complex, and more critical applications are running over the WAN every day. Organizations need solutions that address their specific Business Continuity needs, while at the same time, not compromising the reliability and performance of critical applications delivered over the WAN.

ISPs and large enterprises have multi-homed for years using Border Gateway Protocol (BGP) to connect to multiple Internet backbones, but BGP has many restrictions. For one, it requires that ISPs cooperate with each other and set up “peering” agreements between routers, but because of the performance impact to their networks, many are not willing to do so. BGP also requires expensive routers, designated address blocks and an Address Space Number (ASN), which are sometimes not available to small businesses. BGP requires that gateway hosts exchange dynamic routing tables, which must be constantly synchronized which can lead to delays of up to 30 minutes in changing the traffic direction.

WAN Optimization Services use Network Address Translation (NAT) to unify traffic coming from and going to different destination IP addresses on the Internet. They can be configured with at least one routable IP address for each router/WAN link that is connected to the network.

The biggest benefit of WAN Optimization Services resides in their ability to conduct outgoing and incoming load balancing and failover without defining BGP routing tables or utilizing any of the underlying complicated routing techniques. The ability to offer this functionality without the expensive or complicated networks and equipment necessary to use BGP is what makes them affordable, especially for small and medium-sized enterprises.

Get it right – Choosing the WAN Optimization Service

Outbound Load Balancing and Failover - WAN Optimization Services should provide outgoing load balancing and failover at the TCP/UDP session layer on a per-session basis. The user defines weights (bandwidth capacity) based on the bandwidth of each WAN link. When a session is generated from the LAN, the micro-appliance computes which link has the most available bandwidth and routes traffic from that session over that particular WAN link. The micro-appliance typically allows the selection of two link load balancing algorithms:

- Symmetrical round robin - routes sessions to all links in a round robin manner.
- Intelligent (weighted) load balancing - computes a ratio between the weight (bandwidth capacity) of the different WAN links, and then routes sessions accordingly. That is, the faster the link, the more sessions that will be sent over that link, in order to make the most efficient use of all the bandwidth available.

Inbound Load Balancing and Failover – Inbound Load Balancing and Failover is accomplished by the micro-appliance acting as the authoritative DNS server for the domain(s). The micro-appliance resolves requests from the DNS caching servers to IP addresses on all available WAN links, with the caching servers in turn handing off the addresses in a round robin format. In this manner, all externally initiated sessions are load balanced over all available links. Since the micro-appliance is resident at the domain site and is able to directly monitor the link status, failed links are removed from the DNS tables immediately upon failure. By setting the host name record Time-to-Live (TTL) to a short period, the DNS caching servers will flush their address tables when necessary and will update them from the micro-appliance regularly, and thus be informed when a link fails.

Cost-effective - A quality WAN Optimization Service should deliver easy and affordable WAN/ISP link aggregation, inbound and outbound load balancing, failover and optionally, WAN Virtualization. You may use two, three or however many WAN links and ISPs you need. Leverage low-cost links, eliminate link congestion and bottlenecks, and use the service’s traffic shaping and application prioritization features to guarantee minimum

bandwidth to specific applications.

You can take advantage of the cost of a consumer ADSL link, and get business connectivity at that price. Not only can you get the flexible capacity – you can also buy cost-effective links from multiple ISPs, so that if one link goes down, you can automatically switch over to the other links.

Through bundling (aggregating) multiple, diverse Internet links from one or more ISPs, the WAN Optimization Services reduce the need to purchase multiple and expensive high-speed links. This enables you to increase bandwidth by using cost-effective links without compromising up-time. In addition to managing scalability and redundancy, the service cost-effectively utilizes all available WAN bandwidth through intelligent link load balancing, traffic shaping and application prioritization. WAN Optimization Services should provide controls for how bandwidth is used to support applications and connectivity. This allows you to take advantage of the most cost-effective ISP rates, while ensuring appropriate levels of bandwidth are available for specific applications.

WAN Optimization Services should allow you to choose the WAN link performance/cost ratio that best fits your needs; provides you with complete service provider independence; and eliminates the complexity of network protocols such as border gateway protocol (BGP). The device's inbound and outbound bandwidth aggregation capability combines two or more Internet connections and provides Internet-based applications with access to the total available combined pool of bandwidth. Bandwidth aggregation supports link load balancing to route Internet sessions from congested links, to links with more available bandwidth. It also provides automatic failover of Internet sessions from failed links to functional connections to eliminate a single-point-of-failure. For example, if you have a T1 line (1.5 Mbps), and need additional bandwidth, you would typically have to upgrade to a T3 line (45 Mbps). However, this may be significantly more bandwidth than you require, and will be a significant increase in cost.

With a WAN Optimization Service, this same scenario can be accomplished with two 768 Kbps DSL links that can be combined for a total aggregated bandwidth equivalent to a T1 - at a fraction of the cost. You can also add additional lower speed links such as xDSL, cable, wireless, and others, with a relatively small increase in cost that can more closely match your needs. In addition to receiving more cost-effective bandwidth, you are dramatically increasing the reliability of your WAN network due to the new levels of redundancy through the aggregation of multiple Internet links.

WAN/ISP link aggregation is independent of WAN technologies, and is fully compatible with xDSL, cable, wireless, T1, E1, T3, E3, satellite, fiber, etc. This flexibility allows you to mix and match connectivity to best fit your needs.

Redundant Internet access - Redundant Internet access is the ability to switch traffic among multiple Internet connections through a technique called multi-homing, which more and more small and medium-sized companies are finding they need. When one link goes down, WAN and ISP failover automatically switches your Internet traffic to an appropriately functioning link. Additionally, bandwidth management enables guaranteed bandwidth to your most important application needs.

WAN and ISP failover - When a micro-appliance detects a link failure it should automatically update the DNS record for your domain so that the server requests are sent to the IP address of your alternate server or server cluster. WAN Optimization Services should also provide for device failover through its active/passive failover capability. This eliminates the chance of the micro-appliance being a single-point-of-failure.

Site redundancy - Many businesses need to redirect Internet traffic to a disaster recovery site should a catastrophe disrupt a main site. WAN Optimization Services have in effect, reduced the cost to ensure that site failover and fallback occur automatically, and reliably, making this functionality practical and affordable even for the smallest businesses.

Traffic Shaping and Application Prioritization – Traffic shaping and application prioritization is the ability to prioritize network traffic to ensure that adequate bandwidth is always available to specific bandwidth-intensive applications, especially during periods of congestion. Rules determine bandwidth minimums and maximums for specific types of traffic and use load balancing and automatic failover to direct this traffic to links with sufficient bandwidth.

Availability – Micro-appliances could also be configured in a high-availability mode with one WAN link controller acting as the primary, and a second WAN link controller as a standby.

Performance - Performance of applications over the WAN directly affects response time. This includes not only total average transaction time, but assures that users located at performance-challenged sites (such as branch

offices) still receive the acceptable level of performance. Performance is an important criterion for any piece of networking equipment, but it is critical for WAN Optimization Services, because datacenters are central points of aggregation. As such, the micro-appliance should support extremely high volumes of traffic transmitted to and from sites. A simple definition of performance is how many bits-per-second the device can support. While this is extremely important, in the case of a micro-appliance, other key measures of performance include how many WAN links, number of supported concurrent sessions, and number of domain names and hosts within each domain.

Security - More advanced WAN Optimization Services include built-in firewalling and security features to provide added security and lower the cost of support, maintenance and overall infrastructure complexity through device consolidation.

WAN Virtualization - WAN Optimization Services can support WAN Virtualization to bond multiple network links into a single high-bandwidth channel to ensure high-availability for applications. If one link goes down or degrades in performance, traffic is automatically directed to the best working links without interruption. WAN Virtualization is a form of load balancing which allows for stateful failover of traffic to the best performing links to ensure critical applications avoid problems that occur when they are stopped on one link and restarted over another link. WAN Virtualization ensures that critical applications avoid failures, and are never adversely affected, even after brief disruptions.

Summary

For IT organizations, a solid Business Continuity plan should focus not only on data and application protection, but also on WAN and Internet resiliency that are necessary to ensure continuous access to data and applications delivered over the Internet in the event of a disruption or disaster. ISP outages are a reality today, and they will certainly continue into the future, and the integrity of a business cannot be left only in the hands of third-party ISPs.

Multi-homing, or WAN/ISP link load balancing, is an effective approach to dealing with WAN reliability and performance issues. WAN Optimization Services provide a quick return on investment, compared to multi-homing approaches such as BGP. Selecting the appropriate WAN Optimization Service that addresses your Business Continuity needs and relevant business applications, without compromising their performance when delivered over the network is critical.