ELFIQ Build Smarter Networks

Buildind Smarter WAN Connectivity

A guide to next-generation Internet and private WAN management for decision makers

Table of Contents

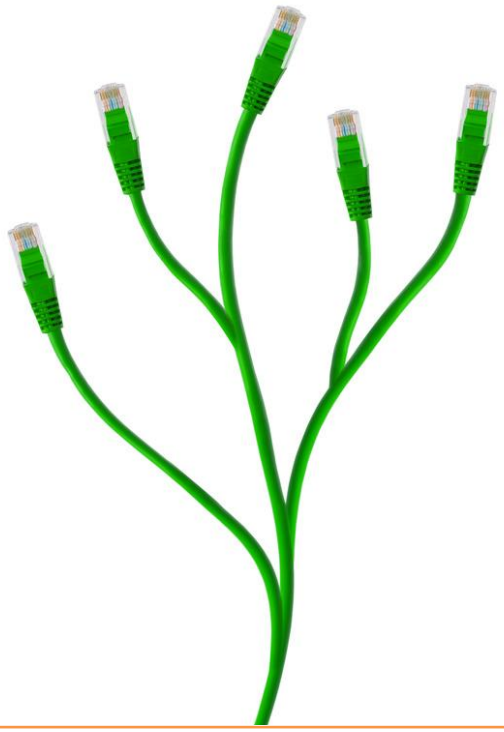

1. Executive Summary

The relentless pressure to reduce costs, increase network performance, improve productivity, and meet regulatory requirements and business continuity needs are a constant battle for the information technology and telecommunications groups of any organizations. Technologies have emerged to help solve these requirements but in many cases they have proven to be inefficient, difficult, unsecured or expensive.

With the Layer-2 link balancing technologies that have entered the market place in the past few years, these objectives now can be met and exceeded, while reducing the complexity of telecommunications infrastructures. Wide Area Networks (WANs) can now enjoy the same levels of performance and reliability as Local Area Networks (LANs). This document explores how organizations can benefit from Layer-2 link management and its applications to build smarter WAN implementations.

2. A quick history of link balancing, BGP, needs and bandwidth requirements

Since the mid-90's, businesses been relying heavily on the Internet and its many diverse uses. With this new business medium came some new challenges, especially in the field of bandwidth management, where organizations need to ensure the availability of Internet access for their users, applications and business transactions.

The challenge is the result of a single point of failure strategy still widely followed today, which involves the use of a single provider and link to the Internet. This strategy puts businesses and their branch offices at risk of being a victim of downtime due for a variety of reasons such as damaged lines, link saturation, equipment defects and logical errors (routing, provisioning errors).

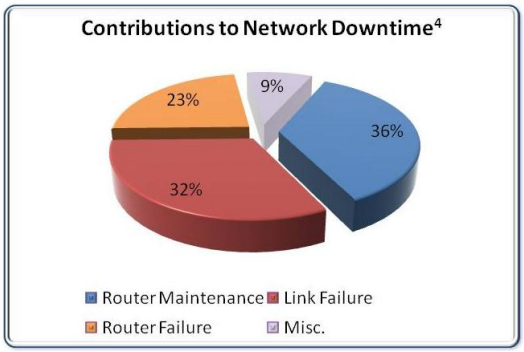

Organizations need to assume that their links will fail for some period of time every year, **where 99.5% reliability results in nearly 2 days of downtime**. This amount of downtime may lead to an array of consequences depending of the type of organization, from mild productivity decreases to significant losses of business related to the unavailability of systems required to perform transactions. E-commerce now represents overall transactions in the billions of dollars each year, and many businesses are dependent on this revenue stream.

A growing concern is link saturation, where no bandwidth is available due to overuse of the existing link. More and more users access technologies and services which consume a lot of link capacity, the favorites being YouTube™, SlingBox™, Skype™, BitTorrent™ and other file-sharing tools. It has become fairly easy to overwhelm a carrier link and prevent normal business from occurring.

New regulations and requirements implemented for many regions of the globe in many industries, along with business continuity and disaster recovery are now key projects being deployed to meet demand, regulations and compliance. Many compliance projects have sections dealing with business continuity and access to vital information for key personnel and auditors, making connectivity reliability a boardroom-level discussion to ensure that organizations meet requirements and are not penalized for failing audits.

Link management is a technique which has been in use for a number of years but costs and complexity have been the principal hurdles to bringing this increase in productivity to most organizations around the world. Border Gateway Protocol[1] (BGP) programming is the traditional method of achieving this result, but it is a costly approach. One aspect of BGP that users must keep in mind is that it was not designed to do link balancing but to perform Internet connectivity.

The first source of cost is the routing equipment, where units designed for the small and medium-sized market are not able to support BGP while the top-tier models are. The second source of cost is related to skills, where organizations need to locate an individual able to configure and maintain BGP devices on top of reconfiguring existing devices and firewalls.

These two items put link management out of reach for a lot of organizations and-/-or their branch offices, before the other considerations regarding BGP such as ISP cooperation, expensive reserved address blocks, and performance/latency issues. ISPs now offer managed BGP services but such services have a high price tag attached to them and can create other types of issues, related mainly to responsibility for problem resolution. BGP also has limited link management abilities forcing a specific link per traffic type and not having the ability to dynamically redirect traffic.

Security with BGP communications needs to be proactively managed, since BGP-based routers may be vulnerable to Internet-based attacks. With such a potential risk, organizations using this approach must proactively maintain the routers with the latest version of their respective firmware to ensure vulnerabilities have been addressed. Should BGP be outsourced, customers must monitor their suppliers' routers to ensure that the units are being updated.

Average hard downtime per month [5]	
Average hard downtime per month	1.7 outages
Average duration per hard downtime	67 minutes
Average total hard downtime per year	23 hours
Average percent of employees affected	28% of employees

With the growing popularity of multi-homed network connections, the BGP routing table used by the required equipment is growing at an alarming speed, causing a slowdown in information processing. Once its size reaches the 244,000 mark, experts[3] anticipate that some older equipment will no longer be able to support such a large number of routes, causing failures and new equipment acquisitions that results in surging infrastructure costs.

New solutions have been introduced in the marketplace, including Elfiq Networks' link balancers based on Layer-2 of the OSI model[2], which enables organizations to deploy more IP links of any type in record time without any interruptions to businesses' operations.

Another balancing technique was developed, using Layer-4 of the OSI (transport) model to enhance link balancing and management. Though product vendors are able to design products rapidly under

this model and provide a significant improvement over the BGP model, it opened up a set of considerations where the benefits may be offset by the drawbacks:

- Layer-4 products act as a proxy; they terminate sessions to manage them causing a degradation in performance
- The device by nature is visible to Internet traffic, making it vulnerable to traditional attacks, and causing maintenance overhead in maintenance and a new security risk in the infrastructure that which be monitored
- A temporary reconfiguration of the network infrastructure is required to integrate and configure the device

3. Layer-2 link management and its benefits

The OSI model is a layered approach to network communications, where Layer-1 is the physical layer and Layer-7 is the application layer. Many products on the market today have opted for an easy way to approach link balancing, using Layer-4 (transport layer). This method enables rapid development of solutions but has its drawbacks in functionality and performance.

The most efficient way to handle link balancing is by performing network operations on Layer-2, enabling a series of key benefits including:

- Wire-speed transmissions of data packets
- Easy link addition and configuration
- Support for any type of IP link (public or private)
- Secure network operation, the unit is transparent to the Internet making it invisible to attacks
- Removal of reconfiguration of existing equipment, including routers and firewalls
- Universal compatibility with all equipment (firewalls, routers, switches)
- No need for ISP involvement or cooperation for configuration and management
- No need for BGP programming or BGP capable equipment
- Advanced traffic handling such as link resilience and geographic per-site policies

Layer-2 link balancing overcomes the performance limitations of Layer-4 implementations and the complexity of BGP-based approaches, enabling organizations to rapidly and efficiently deploy link management into main offices and branch offices alike.

4. Traditional link considerations and alternative technologies for redundant communications

The most widely available links Internet service providers will offer are the ones related to their copper wire and fiber optic networks deployed in almost every city and owned by the telephone companies. The most commonly used link technologies include digital subscriber lines (DSL), T1, T3 and DS3 links. These links enable rapid data transmission and availability is widespread.

The challenge while building a redundant strategy is to locate viable alternatives to the traditional

links. Should an organization rely solely on a single technology, failure of the links will result in complete communications failure, where it happens regularly that repairs to a region's infrastructure accidentally damage wires and will take several hours to complete.

To alleviate this single point of failure, it is recommended to diversify the link technologies used to ensure the availability of Internet access. Following are some suggested technologies:

- Cable modem technology is widely used in most cities and the physical network is usually separate from the telephony networks. These links provide significant performance but no symmetrical data transmission (same performance for outgoing and incoming traffic).
- Emerging xDSL technologies such as ADSL2+, where throughput can reach 24Mbps, are becoming more common in many regions as carriers upgrade their equipment. These links provide increased performance while keeping costs relatively low.
- Microwave point to point carriers provide a method of data transmission based on "line of sight" and independent from traditional wires altogether between the customer site and their points of presence. These carriers can offer symmetrical links and bandwidth up to 100Mbps, making this technology very attractive. Some vendors also make private WAN links available on their networks, providing an ideal complement to existing public and private links.
- 3G Mobile networks offer a viable alternative when building a failover strategy or complementing existing links. CDMA carriers provide EV-DO-based access, with data transmission rates up to 3.1Mbps, and HSDPA can reach up to 14.4Mbps with planned increases in the near future.
- Satellite links based on VSAT technology offer a worthwhile alternative in many regions where other technologies are not available. These links provide performance up to 4Mbps. The downside of this technology is its dependence on weather conditions and response time.
- Utilities fiber networks: A lot of utilities organizations (electricity, natural gas) are diversifying their businesses by enabling data transmission on their networks. Some natural gas organizations have deployed fiber optic networks in their gas conduits, as have electricity-based utilities with their power distribution grids. This network is an alternative to the traditional telco-based network that, when used in conjunction with it, enbles a strategic pairing of reliability and performance.

By combining diverse technologies on Layer-2 link balancing equipment, organizations can build highly efficient and redundant telecommunications infrastructure implementations to enable business uptime, and grow performance for users and applications alike.

5. Efficient strategies to keep businesses running

With growing dependence on IP links to conduct business, regulations and compliance requiring business continuity and sudden demand on existing infrastructures, improved telecom strategies are critical to achieving objectives. To alleviate these pressures, multiple scenarios can be employed, such as:

- Link addition: to meet bandwidth demand or add redundancy at a site, combining links from different technologies and providers will enable uptime and performance.

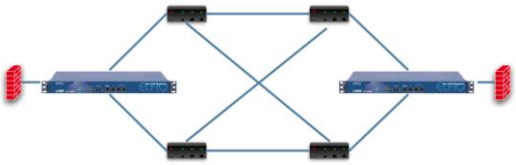

- Link resilience: when implementing point to point or site to site applications and services like VPN and VoIP, using multiple links at each site to enhance performance and ensure uptime is a key architecture item. Planning for a link failure at one site while all other links are operational can keep business and communications running while repairs are made

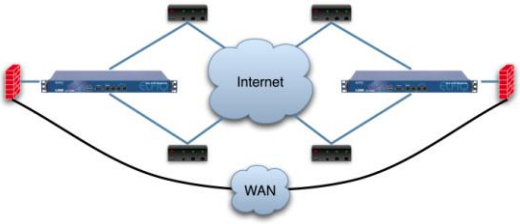

- Multi-site traffic distribution: having the ability to manage incoming and outgoing network traffic to keep business active can effectively maintain services running without users being affected. Should incoming traffic be unable to reach its destination, services like user VPN access, web mail, CRM and others would not be available and cause a loss in productivity. For e-commerce sites, transactions would not occur and revenue would be lost.

In the case of outgoing traffic, employees' ability to conduct business would be significantly reduced and applications would not be able to exchange data with redundancy sites. Geographic balancing between sites lets organizations to keep business running as usual and meet regulatory requirements for business continuity.

6. Cost reduction strategies

A key benefit to link management is the cost savings related to a robust telecom strategy. Costs can be reduced through a number of activities and options, providing direct and indirect savings to organizations.

Proposed cost-reduction strategies:

- Low-cost link aggregation: combining multiple links such as DSL to replace or complement more expensive lines such as T1s and DS3s can save thousands every year.

For example, a T1 line will start at US$350 per month (November 2007, Tier-1 ISP, USA) providing a 1.5Mbps symmetrical link. Replacing this link with 3 DSL links (3Mbps downstream, 768Kbps

upstream, November 2007, Tier-1 ISP, USA) at US$59.99 per month (US$179.97 per month for all 3) will provide a much faster connection to the Internet.

Not only will the organization save over US$2000 per year, it will have 3 links in case of router failure. To strengthen this proposition, a low-cost cable modem link will provide faster transfer rates at about the same price and diversify the link strategy in case of wire failure. ADSL2+ can also be added to this portfolio, where downstream traffic can reach 24Mbps and upstream can reach 3Mbps.

- Adding links to existing single-link strategies will ensure business uptime, directly impacting productivity and application availability.
- Use of alternative technologies: adding alternative vendor links that provide more bandwidth for a smaller monthly rate saves organizations thousands of dollars while providing increased bandwidth.
- BGP-based implementations: changing existing BGP link management by adding Layer-2 link management will enable the use of low-cost equipment, as well as reduce the complexity, knowledge and manpower required to manage the telecommunications infrastructure.
- Failover for branch offices: adding a second link to an existing link for branch offices or retail point of sale systems can be done easily and fairly inexpensively. A wireless high speed Internet link with a mobile carrier can ensure business continuity should the first link fail for only $59.99 per month for a link reaching up 1.4 Mbps downstream and 800 Kbps upstream (November 2007, major US wireless carrier).
- Testing carrier reliability and performance: without modifying or canceling existing agreements and without adding complexity, businesses can effectively test alternative carriers and technologies without business interruptions and performance degradation. Should a link not meet the established criteria, it can be removed in seconds, and users will not feel the difference.
- Carrier migrations: using a link balancer will enable organizations to plan, test and easily migrate primary links when selecting a new primary supplier.

7. Link management considerations

When planning link management for an organization, a list of key items must be considered based on organizational needs and should include:

- Bandwidth management through user-defined configuration for both traffic balancing and failover using multiple algorithms
- The ability to use all links simultaneously
- The ability to add any type of IP link without reconfiguring existing assets and migrating IP addresses
- The number of managed IP addresses for additional links.
- Deployment strategies with regard to infrastructure modifications and migrations
- Quality of service (QOS) for marking, queuing and packet prioritization
- Incoming traffic balancing and management with a strong DNS strategy
- Server availability verification prior to allowing traffic flow to the servers
- Outgoing traffic management

- Ability to obtain quality metrics from each link to verify ISP performance and SLAs
- Multipath linking of multiple sites to bring resilience to point to point applications such as IPSec VPN tunnels and VoIP multisite implementations
- Inbound and outbound geographic link management to prevent downtime and build a robust business continuity / disaster recovery strategy with multiple sites

8. Conclusion

Link management and balancing is a mature technology component that any organization relying on Internet and telephony needs to take a critical look at adopting. With faster bandwidth available at lower costs, it has become very easy with the help of Layer-2 link balancers to bring a new level of performance, resilience, uptime and reliability to telecommunications.

The Layer-2 approach to link balancing is a smarter way to handle this challenge as it provides the ease of use, rapid deployment, security and line-speed performance that were previously missing from the marketplace, and it is becoming a serious contender to take over from the BGP approach that has been the most commonly used method. Layer-2 balancing also provides a new level of security and performance to the other approach of Layer-4 balancing.

Produced by Elfiq Networks

Elfiq Networks is a technology leader and innovator in the field of WAN link management and balancing. With hundreds of successful installations in over 115 countries, Elfiq's Link Balancer products help organizations of any type and size perform more competitively every day with the ability to use multiple Internet and private links easily and securely.

For more information on Elfiq Networks' products and technologies, please contact:

Elfiq Networks
1155 University, #712
Montreal, Quebec, H3A 3A7
Canada
Telephone: 888-GO-ELFIQ / 514-667-0611
Internet: www.elfiq.com
Email: info@elfiq.com

References:

1: Wikipedia (http://en.wikipedia.org/wiki/Border_Gateway_Protocol)
2: Wikipedia (http://en.wikipedia.org/wiki/OSI_model)
3: Network World (http://www.networkworld.com/news/2007/092707-rethinking-internet-routing.html?page=1)
4: Hussain, Iftekhar. Fault-Tolerant IP and MPLS Networks. Indianapolis: Cisco Press, 2005.
5: The Costs of Downtime: North American Medium Businesses 2006, Infonetics Research, March 2006.

September 2011