

White Paper



VPNs: Maximizing Uptime and Performance

Table of Contents

1. Executive Summary.....	3
2. Overview of the Problem	3
3. Supporting Mobile Users and Small Branch Locations.....	5
4. Supporting Site to Site VPNs.....	7
5. Complex Scenarios.....	9
6. Conclusion.....	10

1. Executive Summary

Since the 1990's, virtual private networks (VPNs) have revolutionized how organizations have been interconnected, replacing frame relay systems, dial-up systems and other similar technologies with lower-cost and widely available Internet connections for both offices and mobile employees.

With this new connectivity paradigm, organizations have exponentially grown their usage of this technology into a corporate dependence, creating new problems for the existing approach. Bandwidth has become a double-edged sword that must be handled as such if the long-term use of VPN systems is to be successful.

2. Overview of the Problem

In 2011, VPNs are found everywhere from small businesses to global enterprises, even in the homes of the technologically adept. Low-cost publicly available carrier links such as xDSL and cable modems have enabled mobile users and small branch offices to employ VPNs, while larger corporate offices now have easier access to larger carrier links in their regions, making VPNs a universal business tool.

VPNs are facing two difficult issues, the first being the new information technology reality found in most organizations, where bandwidth use has increased due to a few factors:

- more mobile and remote users employing remote services as organizations enable this service and grow their remote employee rate to reduce real estate costs;
- the explosion of media-rich applications and Web sites in recent years;
- a centralized computing model based on fewer datacenters, virtualization and storage area networks (SANs);
- the use of mobile devices such as smart phones.




This issue directly concerns VPNs and their ability to serve users properly, not to mention overall organizational productivity. This document does not address WAN acceleration, but components discussed are compatible with products providing such capabilities, such as Citrix's Branch Repeater.


The second issue is the availability of the ISP circuits used at both ends of the VPN tunnel. Should a carrier be unavailable, the VPN cannot operate, and if the location to which users connect ceases operations, the entire VPN infrastructure collapses for the duration of the outage.

Organizations have to face the fact that their carrier services may cease functioning normally as research indicates:

- **Elfiq Networks customer survey:² Technology failure rate experienced within 12 months (2009):**



- **Elfiq Networks customer survey: Number of carrier failures within 12 months (2009):**



- **Infonetics Research monthly downtime expectations (2006):¹**

Average hard downtime per month	1.7 outages
Average duration per hard downtime	67 minutes
Average hard downtime per year	23 hours
Average percentage of employees affected	28%


3. Supporting Mobile Users and Small Branch Locations

Mobile users and small branch offices need to remain in contact more than ever via VPN, and solutions exist enabling organizations to reduce the risks to and saturation of their existing carrier link(s). One of these tools is a link balancer, a network-based appliance supporting the use of multiple ISP carriers concurrently and seamlessly. The vast majority of organizations today do not have any redundancy in place, and a few rely on dual-WAN routers to provide a failover process for their services. The challenges of this approach include:

- **IP address management:** Each carrier used in this scheme has its own IP address, making it difficult to manage inbound VPN traffic.
- **Cost management:** In this scenario, the organization is billed for two carrier links while using only one, thus failing to maximize available bandwidth.
- **Complexity management:** Generally with dual-WAN routers, firewall rules have to be duplicated to meet the needs of the secondary carrier link.

To better manage this process of connecting client and mobile VPN access, link balancers can be employed to boost bandwidth resilience and performance. These devices provide the ability to use multiple concurrent ISP carrier links in an existing network, compensating for the two core issues of link availability and link saturation.

This category of device is carrier-independent, so any IP carrier link can be used, from a datacenter-class OC3 (or greater) to a low-cost xDSL or cable modem. Wireless carriers, such as 3G, WiMAX or wireless point to point circuits, can also be added to an existing network. For more information on carrier options and avoiding points of failure, please visit www.elfiq.com/ressources. Link balancers should be able to support VPN access exploiting any commonly used technologies (SSL, IPSec, PPTP).



Internet traffic can be easily compared to the traffic we all experience going to work, where we depend on many factors that determine the overall trip time and enjoyment of the experience. Should the commute happen on a single-lane road, there is no opportunity to go around an obstacle, be it a damaged road segment, a stalled vehicle, a slower vehicle causing congestion or the weather. Managing Internet access is very similar to managing road traffic and has been a challenge for organizations because of the ever-expanding repertoire of Internet-based services users connect to.

The use of multiple concurrent ISP circuits means a strong bandwidth management strategy can be implemented to improve performance and uptime. Adding incremental symmetrical ISP links from multiple carriers using different carrier technologies connecting to different central offices (when possible) will offer the best case scenario for managing VPN clients connecting to a VPN server.

These devices make the concept of “multi-lane highway” for bandwidth possible, allowing more traffic to travel over more lanes. The existing carrier can be significantly complemented with high-download capacity ISP links such as fiber, DSL and cable modems.



A key capability link balancers bring to organizations is the management of inbound services across multiple carrier links. This is very beneficial for all types of VPNs since multiple users and offices connecting to a site can connect through multiple ISPs to share the load and provide continual access in case of carrier failure. The easiest way to provide a single point of connection is to use a fully qualified domain name (FQDN), such as “vpn.company.com”, and DNS entries will point to each ISP carrier’s IP address for the VPN service (for more information on inbound balancing, please visit www.elfiq.com/idns).

Two common cases where a link balancer is employed at a VPN server location are:

- **Mobile users:** Laptop users can connect to vpn.company.com through multiple carriers, and depending on the policies, link status, link saturation and selected balancing algorithm, traffic will be directed to the optimal option at the time of connection.
- **Office to office VPNs:** The same principles apply as with mobile users, the difference being that the VPN’s preferred “NAT Traversal” configuration must be activated for this type of service to be used optimally with a link balancer. Using a link balancer at each end will greatly enhance resilience as well.



A key factor to consider when employing multiple carriers to manage VPN traffic at a single site is that VPN sessions need to be persistent. This means that a single session cannot be balanced across multiple carriers dynamically. Should a session change carriers, it will be reset, and the traffic passing through the VPN tunnel being changed from one carrier to another will be lost.

To increase the overall performance of WAN services such as VPNs, a few key items are becoming commonly implemented with link balancers:


- Lower-cost circuits such as xDSL and cable modems are being used to offload asymmetrical activities such as Web surfing to maximize telecommunication budgets.

- **Traffic segmentation:** By redirecting traffic (including VPN) onto more appropriate links such as cable modems for Web surfing and fiber for VPNs and VoIP, organizations can gain more bandwidth without renegotiating the symmetrical carrier contracts.
- **Quality of Service (QoS):** QoS will permit minimum and/or maximum bandwidth allocation for specified services. This will guarantee that critical traffic will be prioritized while lower-priority traffic will not impinge on the critical traffic's bandwidth. This can be set across multiple carriers.
- **Adaptive configurations:** Based on set conditions such as link failure, link saturation or time of day, the configuration can be dynamically altered to deal with the current situation. This is very useful, especially when carrier links are down, for efficiently controlling the failover process and guaranteeing services such as VPN access.


4. Supporting Site to Site VPNs

When interconnecting multiple offices, where the branch office may require more than one ISP carrier for efficient traffic handling (large number of users, limited availability of ISP carrier services), using a link balancer at each site becomes a viable option. Having a link balancer installed at each site then makes it possible to combine the links to improve traffic performance and reliability.


The fundamentals are the same as those discussed in section 3, where link saturation and uptime are the key issues plaguing VPN services. Before a link balancer is installed at each site, the VPN service diagram would look like this:



The green line connecting both sites is the VPN connection established without the use of any incremental carriers (note that the following diagrams use two sites with two carriers to simplify the discussion). Using the link balancer's capabilities with multiple concurrent carriers means a new paradigm can be implemented that raises throughput capacity and overcomes any saturation issues experienced:




The additional lines represent incremental paths used to transfer traffic from left to right and right to left based on the selected balancing algorithm, policies and link status, effectively improving the performance of the VPN tunnel and providing complete resilience between sites. With advanced balancing algorithms and the right selection of carrier links, a literal combination of carrier links can be achieved for greater performance:



The large green line represents the combined capacity of all carrier links utilized. This represents a significant performance gain for any organization, and any carrier link, from low-cost xDSL links to fiber links, can be used under the right conditions.

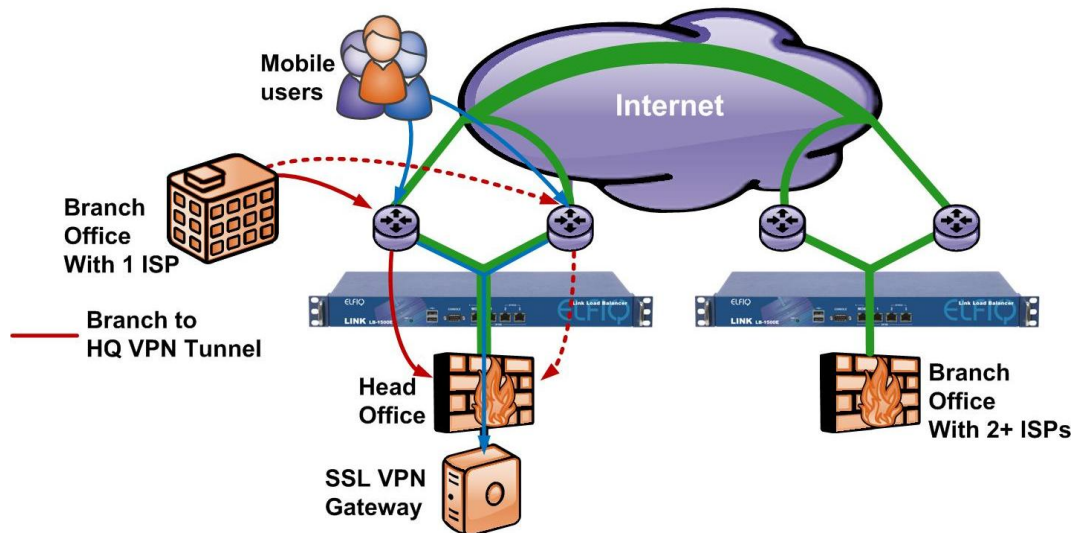
Should a carrier fail in this type of implementation, the VPN tunnel is not terminated but in fact continues operating normally. The link balancer handling the failed link will pass the traffic to an alternative link to ensure continuity.



This is key in deployments where a significant number of users (50 or more) are using the VPN and/or critical services are connected through VPN such as voice over IP (VoIP) sessions and email access. **Since the VPN sessions are not interrupted, users are not affected by the downtime normally associated with carrier link failures, thus significantly improving overall organizational productivity.**

5. Complex Scenarios

Many organizations experience more complex scenarios which must be addressed and cannot be documented here. When complexity increases, all topics discussed in the previous section can be combined to create a connectivity project addressing all types of VPN scenarios. The following diagram provides an example:



Mobile clients are using all available carriers through selective balancing, the branch office has a failover process available should the carrier link it is supposed to connect to fail, and the larger branch office is connected through a resilient link balancer-managed tunnel. This example outlines one of each type, but complex networks meshing VPN deployments are possible with a larger number of sites/users.

6. Conclusion

Building a strong connectivity and bandwidth management strategy will overcome the commonly experienced VPN issues. Saturation and carrier availability can be resolved through the use of link balancers to manage multiple concurrent carriers in order to meet an organization's bandwidth and high availability requirements.

This approach considers the uptime requirement, locally available carrier links, cost control and optimal performance, combined with maximum uptime. As organizations continue to use more VPN services, link balancers will be able to complement and increase their efficiency.

Produced by Elfiq Networks

Elfiq Networks is a technology leader and innovator in the field of WAN link management and balancing. With thousands of successful installations in over 93 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



May 2011

© Copyright 2011, Elfiq Networks (Elfiq Inc.). The contents of this document are protected by copyright. Any modification of this document, in any shape or form, is prohibited. Any redistribution, publication or derivation of the contents of this document without written authorization from Elfiq is also prohibited. All rights reserved. All names, trademarks and copyrights are the property of their respective owners.

® Elfiq registered in the U.S. Patent and Trademark Office and Canada