MULTAPPLIED NETWORKS

Benefit from our Hard-Learned Lessons:
Evaluating Bandwidth Optimization Technologies

This whitepaper outlines the existing technologies we examined before we developed our **BONDED INTERNET™** service. We held ourselves to some pretty important tenets to ensure customers would benefit from our work:

1. Keep it simple
2. Make it cost-effective
3. Maximize performance
4. Maintain carrier autonomy

We invite you to read on and contact us should you wish to explore our **BONDED INTERNET™** solution for your business.

Multapplied Networks Inc.

1127 West 15th Street
North Vancouver, BC
V7P 1M7

General: 866-578-6957
Sales: 778-945-1026

www.multapplied.net
sales@multapplied.net

A Brief Introduction

When we originally set out to explore bandwidth optimization technologies - and ultimately created our **BONDED INTERNET™** advanced bonding solution - we weren't looking to change the Service Provider industry.

However, as we explored the speed, cost, performance, and scalability limitations of other technologies, we were convinced that we were going to have to build something new to satisfy our requirements.

Since its creation, Service Providers around the world have been using our **BONDED INTERNET™** solution to deliver fast, highly reliable, and secure access networks for both their on-net and off-net customers.

You may be familiar with some of the technologies outlined in this paper, and therefore understand their various limitations. We believe that if you're looking to improve bandwidth to your store, branch office, restaurant, or head office, our solution remains the most cost-effective, secure, and feature-rich networking technology in existence.

We invite you to read through these pages with a scrutinizing mind. We know you'll have your doubts, so we invite you to register for a demonstration by visiting us at:

www.multapplied.net/about-us/contact-us/

We didn't set out to change the networking world, but we're glad we did.

Come see why we're so proud of **BONDED INTERNET™.**

Channel Bonding and MLPPP

We began by exploring Channel and MLPPP Bonding - the most commonly used methods for increasing bandwidth. These methods frequently require ISP cooperation and have limitations related to chosen access technologies and bi-directional bandwidth management.

ADSL Channel Bonding

We considered implementing a form of ADSL channel bonding (frequently referred to as G.Bond or ADSL2+), but there were too many limitations for us to be able to justify using it as the base technology for our *Bonded Internet™.*

ADSL Channel Bonding is technically identified under the G.998 specification, and works at the DSLAM level, with ATM time division multiplexing. This bonding is able to reliably provide a multiple of the slowest common denominator of a group of ADSL ports.

The trouble we found was two-fold. Firstly, the DSLAM vendors were not all implementing this functionality, which meant that (secondly) not all Service Providers would be able to offer it to their customers. If customers were tied to a vendor, the service would not be universal, nor would it be able to offer any kind of redundancy.

Below, we've outlined some of the drawbacks we encountered with ADSL Channel Bonding:

- All links must go through a feature-rich DSLAM which limits availability
- Unable to bond other types of connections such as Cable, Fibre, T1, Broadband Wireless
- Unable to bond more than two connections
- No data compression
- No redundancy, QoS or Compression
- Very long installation times
- Often requires ports to be adjacent on the DSLAM which can be difficult in highly subscribed areas
- No ability to add functionality in the CPE or DSLAM

Multilink PPP

PPP (Point-to-Point Protocol) has been around since the 1980s. In the early days, PPP allowed for IP to be tunneled over serial connections such as dialup and ISDN. An extension was added in the 1990s called Multilink PPP, which allowed multiple circuits (each with their own PPP session) to be bundled together.

This same old technology is being used today to combine ADSLs and T1 circuits for additional speed. It has allowed additional life to be given to these aging technologies. (More on those on the next page)

A few of the disadvantages of MLPPP:

- All circuits must be of the same type and speed
- Only works at Layer 2, which means that other layer 3 (IP) connections can not be bonded
- Poor centralized management and monitoring
- Poor handling of unhealthy links
- Configuring CPEs requires Cisco or advanced networking expertise

Next, we explored the two most common network redundancy technologies and evaluated them against our criteria. Load Balancing, a common LAN-based solution that decreases link-load by managing traffic; and BGP, a less-common and more onerous redundancy solution.

Layer 4 TCP/UDP Load balancing

We live in the era of multi-WAN load balancing. There are many examples of these 'boxes', including Peplink, Elfiq, and other routers. The major drawback, related to bandwidth, is that a multi-WAN load balancer does not bond connections together.

When a new connection is established through the device (for example, a web page download), the device assigns the connection to one of the links and the connection stays on that link until it is completed (your web page has been displayed). With enough simultaneous traffic, the device can make full use of each line, but any single transfer will not be faster than the speed of the single fastest line.

Load Balancing partially relieves link congestion, but does not solve bandwidth issues. Here are some of the consequences of load balancers:

There are a couple of small advantages:

- Economic - one-time purchase of hardware required (frequently between $3K and $7K)
- Can be sufficient with a large number of users that generate enough traffic to saturate any one link
- No subscription required (except for annual support, in most cases)

We found, however, that there were too many limitations for load balancing to be considered for our purposes:

- The speed of any session is limited to the speed of the link over which that session is transmitted
- Inbound traffic balancing is very rudimentary and often requires DNS hacks to implement
- Since a connection is bound to a single link, if that link fails then all connections on the link will fail and must be re-established over another link.
- No data compression
- Unable to offer inbound QoS, because they have no control over the download side of the link
- Failover is not guaranteed to work, depending on how an ISP link fails

Border Gateway Protocol (BGP)

Border Gateway Protocol is the most complex and most difficult Internet routing protocol to configure. For some, it is considered the only real option for network redundancy. For end-users to take advantage of the benefits of BGP, there are many requirements on the LAN/WAN side. BGP also requires the consent of your ISP(s) as well as a network administrator capable of handling the routing tables, autonomous systems (AS), and peering relationships. For this reason, BGP is a poor choice for both minimizing costs and maintaining simplicity in the ever-changing network environment.

BGP Multi-path extension can be used to insert multiple routes into the routing table. With Cisco routers, the default CEF (Cisco Express Forwarding) behaviour is to forward packets on a per-destination basis. This means that if the client uses only a single IP with a large flow, that there will be no download performance increase by having multiple connections. Essentially, then, BGP does not increase bandwidth.

CEF can be reconfigured for per-packet load-balancing. However, while providing a costly load-balancing solution, the main disadvantage of Cisco per-packet load balancing is a lack of packet out-of-order recovery and handling of connections that have different latency and jitter characteristics. This means that individual sessions are still subjected to the speed of the slowest link available.

We've mentioned all of the costly requirements for BGP, but we should also mention that BGP load balancing typically requires an expensive, dedicated direct layer 2 connection between the CPE and service provider router. BGP load balancing can be combined with GRE tunnels that are layered on top of commodity broadband connections, but that has its own limitations that affect MTU size and so forth.

Configuring Tunnels

To maintain simplicity and increase bandwidth, we then considered and attempted link aggregation by setting up multiple VPN tunnels to encapsulate traffic. We then used per-packet load-balancing to split traffic between the tunnels. We thought we'd succeeded...we hadn't.

Generic Routing Encapsulation (GRE)

Generic Routing Encapsulation (GRE) is a protocol that encapsulates packets (adds headers and/or footers) in order to route other protocols over IP networks. It was designed to carry Layer 3 protocols over an IP network. It creates a private point-to-point connection similar to that of a Virtual Private Network (VPN).

GRE tunnels can be used in conjunction with IP load balancing solutions such as per-destination or per-packet. GRE tunnels are, essentially, IP-in-IP tunnels and are limited to the bandwidth of the link over which they've been built. Additionally, GRE tunnels do not have bandwidth management features, and so can't allocate any specific bandwidth to any particular tunnel.

For that reason, GRE tunnels alone will not allow you to solve bandwidth issues you may have. The other drawbacks we encountered and that you should consider are as follows:

- Performance with a low number of connections is poor. High performance requires many concurrent sessions.
- Poor handling of DHCP and PPPoE connections - static IPs are required.
- Poor handling of differing MTU (Maximum Transmission Unit) sizes, causing many applications to break.
- No security (requires PPTP, IPSEC or other encryption)
- No handling of out-of-order packets
- Complicated to configure - requires advanced networking and Cisco switch/router experience.

SSL tunnels via OpenVPN

Another approach is to configure multiple SSL tunnels using the popular OpenVPN software. We worked through the pains associated with this option, but ultimately found that there were severe drawbacks to using individual tunnels for each Internet connection in the bond.

First, we experienced a lot of frustrations around combining links of varying jitter and/or latency. Load-balancing by way of OpenVPN would allow us to optimize the bandwidth we needed, but we were still limited in our selection of usable links.

There are multiple problems with combining multiple OpenVPN tunnels:

- Requires links that are nearly identical in speed as well as latency
- Achieving >90% consistent aggregate performance is very difficult especially with heterogeneous connections
- Poor handling of DHCP and PPPoE connections - static IPs are required.
- Routing changes were complicated and difficult
- No centralized management and monitoring
- No handling of out-of-order packets

We know tunnels

An insider secret - Versions prior to 4.0 of the Multapplied solution used OpenVPN with per-packet load balancing behind the scenes. Those versions caused a ton of frustration with the performance and reliability problems, especially when bonding connections from different carriers. In version 4, we built a brand new solution using our own tunneling mechanism - combined with all the knowledge we had gained - that solved all of the problems we encountered with individual tunnels.

Traditional Bonding

Bonded technologies seem to have been, until recently, the ISPs' solution of choice for delivering network bandwidth beyond the speeds of any single connection. The trouble with these second-generation solutions is that there is no redundancy built into them. It didn't take us long to abandon these options.

Bonded ADSL

There are many ISPs today providing Bonded ADSL - a technology that combines multiple circuits using Multilink PPP - some go beyond MLPPP but require the use of an expensive Cisco router or a linux server.

Depending on the provider, Bonded ADSL can come in a variety of forms, so it is important for the end-user to be aware of the significant limitations each can have. For example, some providers of Bonded ADSL think it's perfectly fine to bond the download speeds, but not the upload, leaving crucial outbound traffic at a significant disadvantage.

Bonded ADSL can be economically priced, but has several technical disadvantages which have direct business impacts.

There are many service providers out there offering Bonded ADSL. While there are some clear advantages, you should be aware of some of the limitations:

- Aggregate speed is a multiple of the slowest connection
- All lines must be through the same provider
- The ISP is a single point of failure
- No ability to have a "failover only" line such as a 4G wireless modem or backup cable connection
- Generally limited to only 4 connections
- No compression
- No end-to-end QoS (Quality of Service)

Bonded T1

Bonded T1 (Frequently seen on websites as NxT1 where N represents the number of T1s used in the bond) uses Multilink PPP to combine multiple T1 data circuits. Thus in a lot of ways, it is similar to Bonded ADSL.

An advantage is that T1s can be available in areas that ADSL is not, such as in remote areas. T1s are symmetrical with the same 1.54 Mbps speed for both up and download. Thus, they do have a higher upload speed than ADSL.

In fact, T1s generally offer enough outbound bandwidth to satisfy the requirements of many SMB businesses. But as more companies adopt hosted, or cloud-based applications, the need for bandwidth (download especially) is surpassing the capabilities of bonded T1s. Before considering Bonded T1s, consider the following:

- Speed only increases by a multiple of 1.54 Mbps with each bonded link
- Other link types such as DSL, Cable, or Fixed Wireless cannot be added to the bond.
- The T1s must all be provided by the same ISP
- The bonding is done through the Multilink PPP protocol, so the same disadvantages mentioned for MLPPP apply here.
- Costs are high - typically $300-400 per circuit, plus extra for the bonding service and bandwidth
- Installation time can be lengthy - weeks to months
- Generally involves lengthy multi-year contracts

Multapplied's BONDED INTERNET™

When we finally decided that there was no bandwidth-improving service that met our criteria, we decided to build our own solution: BONDED INTERNET™.

We believe BONDED INTERNET™ offers the performance our customers need, the ISP autonomy they want, the cost-effectiveness they demand, and is simple enough even our marketing guy can configure and manage it.

Multapplied's Bonded Internet™ Solution

We're very pleased with how our Bonded Internet™ turned out. We spent a lot of time and effort to ensure it met our top criteria, and delivered the kind of network performance our customers need.

You've read a lot about the benefits and limitations of the other common networking technologies - and you may be using one of them today. We believe we're ahead of the pack, having created a solution that:

- Can bond any type of connection, in both directions
- Bonds connections from any service provider
- Can maximize the true aggregate of all the link speeds, and not a multiple of the slowest link
- Handles links with different latencies and speeds
- Includes compression to accelerate packet delivery
- Fails-over on network outages in 0.3 seconds
- Offers Quality of Service past the LAN and to the front-door of the Internet
- Is cost-efficient for universal adoptability
- Has centralized management and monitoring
- Does not require DNS hacks
- Does not require technical skills to deploy or maintain

If you're interested in a network solution that maximizes the bandwidth you have, ensures proper redundancy metrics are met, is quick to deliver, and is easily managed, we believe you've found it in **BONDED INTERNET™** - our advanced bonding solution.

A Typical BONDED INTERNET™ Deployment

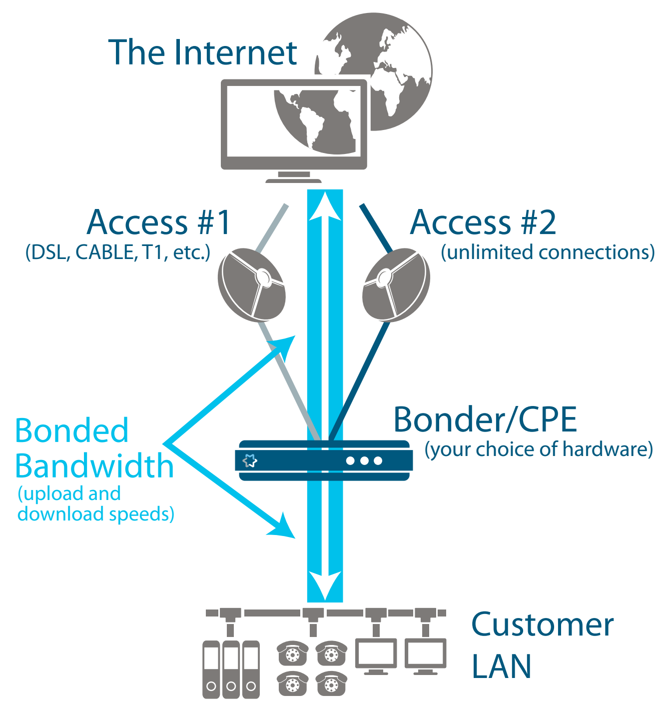

Come see BONDED INTERNET™ for yourself.

www.multapplied.net | 866-578-6957