



## Solutions to the Infrastructure Gap:

# The Good, The Bad, and The Ugly of First-Generation Network Technologies.

With the increase in demand for Internet applications, organizations are struggling to find the performance they require in their networks without drastically increasing costs. This paper looks at the First-Generation technologies used to bridge this gap, and points out their limitations.

This paper provides the reader with a high-level overview of the benefits and detriments of Bonding, Load-Balancing, and MPLS network technologies and looks to the future with a new technology called 'Autonomous Network Aggregation'.


# Solutions to the Infrastructure Gap: The Good, The Bad, The Ugly

It's no revelation that the computer is the most widely adopted technology in recent history. However, the Internet is an anomaly by way of its decentralized nature: it has no outright owner, and there exists no manager to regulate the recent overwhelming increase in traffic.

Long ago, the Internet achieved its critical mass and its growth (adoption rate and usage rate) is continuing at an exponential rate. Recently, web-savvy pundits have proposed that the Internet will be accessed by 6 Billion people in 2012 – all consuming data on more than 2 Billion Internet-enabled devices. Demand for the Internet, and the information and entertainment contained therein, is skyrocketing.

Commonly in business, high demand for products comes with high prices. However, and it is possibly due to the lack of a single “owner”, the Internet has seen the opposite. As demand rises, prices drop. The users are regulating the costs and as competition increases (Deregulation for some, the sheer ease with which entrants can ‘get in the game’ for others), the price of “supply” is dropping significantly. We believe that the demand is coming from increases in usability, stickiness, and entertainment-value of online content.

There is, however, required infrastructure in order to meet the accessibility requirements of users. And so, as prices drop, so do profit margins for those developing that infrastructure.



We believe it works like this: Service Providers make the Internet accessible for the End User. The End User goes online and seeks out content. Developers create rich (and bandwidth-laden) content for the end user, who then demands improved infrastructure from the Service Provider.

The cycle, while entertaining, is fierce and is affecting everyone. Incumbent Service providers are increasingly unable to deliver the types of pipes required by users who want to access online content. The service providers that have the infrastructure networks are now consuming market share from tier-2 providers because the Tier-2 providers haven't got their own Ethernet networks. ISPs used to peer (allow traffic from other ISPs to flow across their network), but with online content getting more bandwidth-laden (Netflix being a prime example) peering is becoming less and less viable as profit margins sink.

Author  
Cameron Couch

The prohibitive costs associated with creating Ethernet networks leaves many infrastructure gaps – and the task falls to governments to open up bandwidth capabilities to their populace.

In April of 2009, American Congress initiated a plan to improve America's broadband accessibility. Likening today's broadband challenge to the electricity challenge of the early 20th century, the government charged the Federal Communications Commission (FCC) to develop a National Broadband Plan to ensure every American has "access to broadband capacity." Broadband, it stated, "is a foundation for economic growth, job creation, global competitiveness and a better way of life."

The Plan, however, isn't set to be fully realized this month, nor this year, nor in the next five years. The National Broadband Plan is set to realize the delivery of "high-speed for everyone" by 2020. In order to address the gap between infrastructure availability and the bandwidth requirements of today, a number of solutions have come into the market that answered the question: "How do I get enough speed out of what is available to me today?"

This paper will go into each of the three major technologies that help accelerate communication networks and enable network convergence (another big reason why organizations and individuals are looking for more bandwidth) and will provide the good, the bad, and ugly on each one.

### Bonding (aka Channel Bonding)

Bonding is the oldest of the convergence technologies. It involves combining two cloud connections to give the user one large pipe, thereby increasing throughput ability. One plus One equals Two. (Sort Of).

#### *The Good*

Bonding was the first technology to truly bridge the infrastructure gap. Organizations the world over adopted the technology as it was a quick way to solve lag in the network. With Bonding, a company looking to get increased bandwidth does not have to move to an expensive Ethernet connection, but can simply bond two low-speed DSL connections to achieve their required bandwidth – almost.

#### *The Bad*

The reason Bonding doesn't necessarily thrive in the market place is that it doesn't truly aggregate the speeds of each connection. Instead, the Bonded solution only doubles the lowest-speed connection (Or, the denominational connection). Bonding a 3Mbit connection with a 5Mbit connection, yields just a 6Mbit connection (2 X 3Mbits). For certain, one extra Megabit is a major increase in upload or download speed, but the technology developers stop short of offering a truly beneficial solution to companies by limiting the aggregated speeds. If a company is paying for a 6Meg connection, it would behoove the supplier to offer the 6Megs. And if the organization is paying for

a 6Meg and a 3Meg connection, the supplier should be able to offer 9Megs.

The traditional Bonding solution also requires the use of two similar connections (ie. Two DSLs, or Two Cable connections). This leaves the network with no diversity in access-products and, therefore, neither speeds nor costs are not optimized.

### *The Ugly*

The major problem with Bonded solutions stems from the technology's dependence on single-carrier connections. The main reason why single-carrier connections inhibit a network is because there is no redundancy or reliability in the configuration. Network technologies should allow for the use of diverse carriers or suppliers so that the organization relying on the solution is not limited by the up-time of one supplier. Instead, the system can be held up by two or more connections from different carriers, and gain an unparalleled reliability in the process.

The lack of redundancy from diverse carrier connections also prohibits Bonding solutions from offering any acceptable level of Quality of Service. Because the technology traditionally must comprise same-type connections from the same carrier/supplier, if the carrier has quality issues or goes down, the entire network is affected.

### Load Balancing

Load Balancing is a technology that was developed to reduce latency for critical applications in the network. This technology allows for a single or group of applications to run along a specific connection, while another group of applications runs along another connection.

### *The Good.*

Load Balancing was the first technology to improve in some areas on Bonding. Where Bonding increases the overall bandwidth of the network, the issue still remains that some applications are inherently more important than others, and therefore need to have priority in the pipe. Load Balancing helps take care of the prioritization of these critical applications.

There are two major benefits to Load Balancing technology. The first benefit comes in the form of segregation. With Load Balancing, customers can select which applications go down which connections. In a simple Load-Balanced set-up, "Application A" will travel down "Connection A" and "Application B" will travel down "Connection B". This segregation ensures that traffic down "Connection A" is smooth and unaffected. Regardless of what load is placed on "Connection B", all applications on "Connection A" will operate without jitter or latency.

The other major benefit of Load Balancing technology is "Diverse Carrier Support" (ie. that it allows companies to use connections from differing carri-

ers). In our earlier network example, “Connection A”, then, could be supplied by “Carrier A” and “Connection B” could be supplied by “Carrier B”. The diversity gives a bit of added reliability to the network, as the organization as a whole is not limited by the up-time of one carrier.

### *The Bad.*

The issue with Load Balancing, is that the connection speeds are not aggregated at all. The increased performance of a given application is not due to any increase in overall bandwidth, but in the available bandwidth of the pipe along which the application travels. Load Balancing does not so much increase speeds across the network, as it does decrease the amount of traffic that affects critical applications.

### *The Ugly.*

As much as “Diverse Carrier Support” is a benefit of Load Balancing, it must be mentioned that much of the redundancy from Diverse Carrier Support in Load Balancing is actually *false*. The problem with segregating traffic is that there is no redundancy for the applications on the connection through which they travel. For example: If “Application A” is running down “Connection A”, and “Connection A” goes down, so, too, does “Application A”. There is no redundancy for any single application. Sure, the network is still operational, but only those applications on “Connection B” will be operable.

There is no way to truly prioritize applications in a Load-Balanced network. It is simply a guessing game as to which of the carriers is deemed “most reliable” at the time of configuration. There is ultimately no redundancy for any application – which is the real benefit of “Diverse Carrier Support”.

## MPLS - Multi-Protocol Label Switching

MPLS is a packet-prioritization technique that was developed to ensure application quality and priority across a network. It essentially allows for information ‘packets’ to travel across a network safely and quickly. The data is encapsulated with “keys” on either end in order to rank the data packet and give it priority across the network.

### *The Good*

Without a doubt, the best part of MPLS is the reliability. Supported by dedicated accesses, an MPLS network rivals all of the aforementioned networks. The process by which data packets travel across the network is efficient and allows for enough time-savings to make the system a unique balance of speed and reliability. What’s more, an MPLS network allows for application prioritization and Quality of Service – both of which stem from the reliability and are the major benefits to users of this technology. Without Quality of Service, a network can limit an organization’s ability to do business.

### *The Bad*

Unfortunately, an MPLS network with Quality of Service is going to cost the organization sometimes crippling amounts of money. In 2007, 35% of registered businesses in the United States reported having between five and nineteen employees. To that market, an MPLS network is cost-prohibitive. So while the technology is first-class, its cost is not.

### *The Ugly*

For what it does, an MPLS network is a very usable technology. Deterrents come from price, as well as what we call “wasted bandwidth”. MPLS networks work on symmetrical connections which means an organization will get the same in upload speed as it gets in download speed. Regular broadband connections are asymmetric (that is, there is higher download speeds than upload speeds). With MPLS, asymmetrical connections are not supported – and therefore, the entire network is limited to the upload capability of the connection being used to support it. A 1.5Mbit MPLS will have 1.5Mbits of download speed and 1.5Mbits of upload speed. If the connection used is an ADSL that has 6Mbits down and 1Mbit up, the network will only support traffic at 1Mbit in either direction. Essentially, the organization is wasting 5Mbits of bandwidth.

The problem with MPLS rings truest when the organization is converging voice and data networks over the MPLS. A 1.5Mbit MPLS will generally be cut into two pieces – 768Kbps for voice, and 768Kbps for all other traffic. A T1-based MPLS will offer enough bandwidth to support voice, but the remaining 768Kbps is not sufficient for the other applications used by the organization. And so the user is required to increase bandwidth beyond the T1, adding more unnecessary costs to their network.

### *The Future is here.*

In order to solve the Infrastructure gap and the increasing demand for more bandwidth, a solution has emerged called Autonomous Network Aggregation. This new technology takes the goods from each of the aforementioned solutions and combines them to offer a truly next-generation solution to organizations.

### *Aggregation*

This new Aggregation technology offers the combined speeds of each of the connections involved. If an organization combines two DSLs using this technology, they receive the sum of the connections as a “super connection”. That is, if the organization aggregates a 6M connection with a 3M connection, the result is a 9M network. More speed can be acquired by adding a compression technology that compresses data packets where possible.

### *Diverse Carrier Support*

Similar to Load Balancing, the Aggregation solution offers increased reliability through the ability to combine connections from different carriers. In this case, however, the resulting “super connection” is a single pipe, so that applications can travel across either of the networks if one is not performing suitably.

### *Prioritization*

From MPLS, the Aggregation solution steals the concept of prioritization. With this solution an organization can rank any number of applications to ensure that the ‘packets’ from a high-priority application travel across the network along the smoothest and fastest route possible.

### *Quality of Service*

Finally, the Aggregated solution improved on the QoS offered by an MPLS network – by making it support both the upload and the download speeds of an asymmetrical connection. That is, if a company has an application that requires more download-bandwidth than upload, the system can provide Quality of Service on all of that application’s bandwidth. It can offer the QoS on all of the bandwidth, in any direction, for all of the applications.

The Internet’s stickiness is being readily realized by almost everyone who uses it. As adoption grows, so too will bandwidth requirements. We’re not in the business of predicting bandwidth usage, but we can see already that the current infrastructure in the market is being consumed by business and residential users alike.

The “First-Generation” solutions to the infrastructure gap each have drawbacks when we consider bandwidth and/or reliability requirements. As more organizations move critical applications to the Cloud, the gaps in infrastructure are going to become more daunting and more injurious to the success of those organizations. The current thought is that companies are required to increase their bandwidth by way of dedicated, expensive, Ethernet connections to help them move successfully into the next phase of Internet usage.

We think current thought is wrong. We believe the future of communication networks is enhanced reliability from diverse carrier support, application prioritization, and Quality of Service. We also believe that the speeds that are available to an organization can be optimized in such a way as to deliver enough bandwidth for that organization to realize true efficiency in its network. Today.

#### **TELoIP Inc.**

487 Adelaide Street West  
Suite 200  
Toronto, Ontario, Canada  
M5V 1T4

*TELoIP is a Toronto-based technology firm serving the most forward-thinking organizations in the North American communications market. TElOIP creates and sells network convergence solutions based on a technology that enables voice, data, and video to work simultaneously over the same infrastructure. Their patent-pending Ana™ system aggregates and accelerates any connection, from any carrier - combining the speeds of each connection and providing customers with reliability for their critical applications. Bandwidth requirements are growing at exponential rates, TElOIP's technology helps companies scale up quickly and reliably to meet tomorrow's growing bandwidth requirements. For more information, please visit [www.teloip.com](http://www.teloip.com)*