

The Real-Time IP Network

Moving IP Networks Beyond Best Effort to Deliver Real-Time Applications

Abstract

The Internet has experienced tremendous growth over the past decade – fueled primarily by best effort applications such as e-mail, web browsing and file transfer. Despite this impressive growth, IP networks have had mixed success in supporting newer Real-Time applications and converged legacy services. Limited primarily to commodity best-effort services, carriers have struggled to earn profits from their investments in IP networks. In today's competitive telecommunications environment, it is vital for carriers to move beyond best effort services to higher-margin Real-Time services such as voice, business IP, VPNs and gaming. This paper identifies the specific reliability and stability limitations of traditional best effort networks, and outlines the steps required to upgrade an IP network to achieve the performance required for profitable Real-Time services.

Author

G. Hudson Gilmer
hgilmer@avici.com

Best Effort Networks – Growth without Profit

Since the early 1990s, carriers have invested heavily to build out and rapidly expand IP networks. IP was heralded as the source of future growth, services and profits. Despite tremendous traffic growth and the near ubiquity of IP in residential and business networks, the promise of profits has not been fulfilled. In 2003, despite the fact that IP networks represent over 50% of worldwide carrier network traffic, they account for less than 10% of revenues and profits.¹

As long as IP services continue to be dominated by commodity best effort services, price will be the deciding factor for carriers and carrier margins will be low.

The good news for carriers is that a second wave of IP growth is emerging in which customers are using IP for more than simply Best Effort applications.

Real-Time Apps: The 2nd Wave of IP Growth

Real-Time IP applications such as voice, business IP, video, and interactive gaming are the next wave of growth for IP networks, aided by the continued rollout of broadband IP access. These bandwidth-intensive applications are referred to as Real-Time because, unlike best effort applications, they must be transported through a network with minimal delay or latency.

These Real-Time applications have evolved from experimental “shareware” to profitable fee-based services over the past 4-6 years. Voice over the Internet began as a loose band of technology visionaries in projects like VON (Voice Over the Net), and has evolved into a billion dollar market embraced by Tier-1 carriers. Similarly, audio file sharing began with shareware such as Napster and Kazaa, and is evolving into a legal fee-based business model with sites such as iTunes and Pressplay. Another example is the rapid growth of MPOGs (Multi-Player Online Games), which began with web-based games such as Everquest, and now

claim millions of online users on platforms such as Microsoft’s X-Box, Nintendo’s Gamecube and Sony’s Playstation.

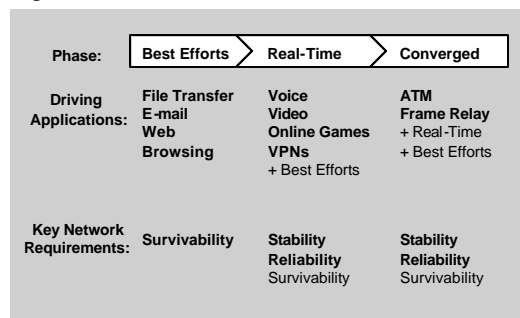
In the early shareware/best effort stage of these applications, users were accustomed to frequent disconnections and poor quality. However, as these customers migrate to fee-based Real-Time applications, their expectations of service quality and reliability increase.

This new phase of Real-Time IP Service growth presents compelling carrier opportunities for service differentiation and premium pricing, but demand higher levels of reliability and stability than traditional IP networks were designed to deliver.

Convergence to IP: Driven by Dollars

Another driver for upgrading the performance of IP networks is convergence. In the late 1990s, convergence of “everything over IP” was advocated for a variety of reasons including technology, strategic and industry hype. Today, there is only one factor driving the migration of legacy networks and services over IP – the business case.

Figure 1: Evolution of IP Networks:



The pressure on carriers to achieve cost savings and revenue growth from convergence is immense. The ubiquity, scale, growth and efficiency of IP networks make it an attractive platform for supporting these emerging applications.

Carriers are pragmatic about the timing and speed of convergence, and are unable to write off large investments in existing legacy networks, however virtually all carriers have

¹ “Worldwide Telecommunications Services Forecast and Analysis, 2002-2007”, IDC

articulated and begun a strategy of migration towards IP. As with Real-Time applications, however, traditional best-effort IP applications do not achieve the stability and reliability levels required for ATM, Frame and voice networks.

Designed for Survivability , Not Stability

The precursor to the Internet – ARPAnet, was developed with the primary design goal of survivability. The Internet’s survivable design has enabled rapid growth of best effort services, but now limits its ability to support high-margin Real-Time and converged services.

Traditional layer-2 networks are susceptible to link and node failures, and therefore are designed with highly reliable (99.999% availability) elements and often employ SONET APS link protection. In contrast, IP networks are highly dynamic, re-calculating routing paths for the entire network (re-converging) in response to any outage/change in the network. This characteristic reduces the need for high reliability in individual network elements (routers) or links, because the network can simply re-converge and find an alternative route around outages.

This dynamic, unstable network model is perfectly suited to best effort applications such as e-mail, web browsing and non-critical data, which are not time-sensitive and tolerant to limited packet loss.

In contrast, applications such as voice, video, and ATM are time-sensitive and rely on a highly stable network. In order for IP networks to support such higher-margin applications, three key improvements are required: improved router reliability, improved network stability, and faster convergence when disruptions do occur. By addressing these three basic requirements, carriers can achieve the reliability and stability on their IP networks required to support new Real-Time services and converge legacy networks.

The following sections examine the causes of instability in today’s best effort IP networks, and identify four steps carriers can take to build a Real-Time IP network capable of converging

legacy services and supporting profitable new services.


Step 1: 99.999% Router Availability

The first area to focus on improving IP networks is the reliability of routers themselves. In contrast to traditional Central Office equipment such as voice and ATM switches, IP routers have not historically been designed for 99.999% availability. Recent data from RHK shows that a typical large IP network achieves between 99.95 and 99.99% availability.² This corresponds to 50-260 minutes of downtime - 10-50 times higher downtime than the “five nines” availability benchmark of legacy data networks.

To build a router that delivers 99.999% system availability requires a comprehensive focus on all dimensions of the routing platform, and a full understanding of the network environment into which the router is placed. Features such as non-stop routing (for control plane protection), robust field-hardened software, fully redundant hardware, and rigorous testing/QA procedures all contribute to achieving carrier-class system availability.

Many vendors claim to achieve 99.999% availability, however many of these claims apply only to hardware availability, despite the fact that software failures are a larger cause of IP Network downtime (see Figure 2). In addition, many availability claims are predictions based on models, rather than measurements of performance in operational networks.

Figure 2: Distribution of IP Network Outages



² “The Coming Era of Absolute Availability”, May, 2003; Shing Yin & Ken Twist, RHK

Step 2: Local Link Protection

A study by the University of Michigan³ observed that 32% of outages in a large regional IP network were attributed to link failures. Due to the connectionless nature of IP networks, the impact of link failure is not limited to traffic traversing the failed link. Rerouting can cause congestion elsewhere in the network, and protocol convergence can lead to routing errors and instability network-wide. Such traffic loss and network instability is unacceptable for networks carrying Real-Time and converged services.

In order to prevent traffic loss and network instability, links must be locally protected from failure. Many link protection schemes have been developed with the common objective of locally protecting the link in less than 50 milliseconds. Fail-over times of less than 50 ms are imperceptible to the data and control planes and avoid triggering convergence of the layer-3 protocols. The result is that local link protection schemes with sub 50 ms fail-over times do not impact even Real-Time services running over them.

These local protection schemes each have benefits and tradeoffs. SONET APS is considered the gold standard of local link protection, but with 1:1 active/protect it is too expensive to deploy broadly. MPLS Fast Reroute offers effective protection and the flexibility to share backup links, but adds complexity due to the need to provision protection at each hop. Link Aggregation provides 1:N local protection of Ethernet links, but does not extend to SONET/SDH links.

A new class of local link protection schemes has emerged, designed to offer speed, link upgradeability, cost effectiveness, and simplicity. For example, Avici's Composite Links™ enables up to 64 physical POS links to be grouped into a single logical link. In the event any member link fails, traffic is redistributed across the surviving links in less than 45 milliseconds. This not only provides

³ C. Labovitz, A. Ahuja, F. Jahanian; "Experimental Study of Internet Stability and Wide-Area Backbone Failures", Merit Network, Inc. and University of Michigan

cost-effective 1:N protection, it also provides a non-disruptive mechanism for provisioning additional bandwidth.

Step 3: Eliminate Scheduled Downtime

Another significant cause of (planned) downtime in carrier IP networks is Network Administration. Network Administration includes routine operational changes including software upgrades, hardware upgrades, link upgrades and configuration changes.

Traditional "best effort" routing platforms do not support in-service hardware or software upgrades, forcing carriers to schedule frequent maintenance windows for routine network administration tasks and upgrades.

For example, installation of a simple software patch on a traditional core router requires a full reboot and results in 10 minutes or more of router downtime, assuming no complications. Even installation or removal of "hot swappable" line cards on traditional routers can impact all traffic for one or more seconds. Some of the most severe outages result from configuration changes which contain errors, however traditional routers offer no mechanism for users to make protected configuration changes with the ability to revert to the previous configuration.

Customers using best effort applications such as e-mail are unlikely to notice the impact of scheduled downtime. However, with the emergence of round-the-clock applications such as global VPNs, interactive gaming, and voice, customers are more sensitive to any network disturbance, and are beginning to demand availability that SLAs apply 7x24, effectively forcing an end to the practice of scheduled downtime.

To achieve zero planned downtime, routing platforms must be designed to accommodate network administration activities without service impact. The following checklist contains key features required to eliminate planned downtime:

1. *Hitless in-service software upgrades*
2. *POS link aggregation mechanism for non-disruptive link expansion/upgrade*

3. Support for protected configuration changes (revert to previous config)
4. Hot swappable line-cards with no traffic impact (or less than 50 ms impact)
5. Modular hardware scalability with in-service expansion

New routing platforms are beginning to emerge which offer in-service hardware scalability and software upgrades. By migrating to such platforms, carriers can virtually eliminate hardware and software upgrades as a source of disruption in the network.

Step 4: Rapid Route Table Convergence


The strategies outlined above are designed to dramatically reduce the number and frequency of disruptions to an IP network. Even with such precautions, however, no network is completely immune from disruptions such as human error or topology changes.

Planned or unplanned network events in IP networks produce topology changes that require route forwarding tables to converge upon the most up to date routing information for the relevant protocols. These protocol convergence events occur on a near-daily basis – even in well-managed networks. While these convergence events or “microbursts” typically only last for up to 20 seconds, they can cause significant service and financial impact to carriers.

The first order impact is that packets are lost. While 10-20 seconds of lost packets may be an inconvenience for users of best effort applications, it results in disconnection for Real-Time and layer-2 services. RHK estimated that each microburst drops over 6,000 online gaming players and nearly 100 VoIP calls, with a cumulative (direct and indirect) cost of over \$200,000 per month.⁴ These costs will only increase in future years as the number of Real-Time applications and users continues to increase.

⁴ “The Coming Era of Absolute Availability”, May, 2003; Shing Yin & Ken Twist, RHK

Figure 3: Microburst Costs per year



Microbursts not only have a localized impact on services running over the impacted network element, but can also cause more widespread network damage and instability, including:

- Transient routing loops where packets bounce through the network until their time-to-live expires or another path is established. Routing loops consume resources, eating up both bandwidth and packet storage.
- Poor network performance during microbursts, as packets are forced to take non-optimal routes until the network converges.
- Loss of routing protocol sessions, which time out and/or tear down as a result of failure to properly exchange update information. These failures can occur in as little as 3 seconds and trigger vast quantities of control plane traffic as protocol sessions attempt to re-establish themselves and re-learn routes from peer routers.

In order to minimize the direct and indirect impacts of microbursts on Real-Time services, convergence times must be improved from tens of seconds to second and eventually millisecond timeframes. Improving protocol convergence times from current levels to single second levels can eliminate up to 95% of the service impacts and their associated costs (see Appendix 1), however such improvements are not easy to achieve.

An order-of-magnitude improvement in protocol convergence time is difficult to achieve on traditional routers because of the need for extremely fast route processors, optimized convergence algorithms, distributed line card processing, and configurable timers.

Further improvement of protocol convergence times to millisecond levels can dramatically simplify network design by eliminating the need

The Real-Time IP Network

for local link protection to be widely deployed/provisioned.

Because core routers handle the bulk of traffic redirection in large networks and sit on the greatest number of paths, significant improvements in protocol convergence times at the core improves overall network performance regardless of where the microburst occurs.

The Real-Time IP Network Business Case

Now that the 4-step roadmap to the Real-Time Network has been completed, let's examine what a converged Real-Time network can mean to a carrier's bottom line.

One of the major challenges for IP Service Providers has been the commoditization and pricing pressures inherent in selling undifferentiated best effort IP services. With the Real-Time network, carriers have the opportunity to sell differentiated high-margin services designed for the stringent performance needs of video, VPN, voice and gaming customers. The result is higher revenues through premium margins and lower customer churn.

Figure 4: Traditional vs. Real-Time IP Network



On the cost side of the equation, moving to the Real-Time network delivers compelling cost savings by eliminating redundant network elements, significantly reducing service-impacting outages, and enabling convergence of networks.

- **Eliminating Redundancy** - By deploying routers with 99.999% system availability, inherent redundancy and

in-service upgrades, carriers can remove the redundant routers in their networks and decrease and eliminate up to 50% of their router capital costs and operational costs. Deploying cost-effective local link protection mechanisms, such as Composite Links, enables carriers to reducing the need for vast amounts of spare capacity network-wide (typically up to 50%). This enables carriers to increase network utilization without increased costs.

- **Reduced Outages & Downtime** - The largest cost reduction impact is eliminating the costs of network outages and service failure. With costs of network downtime conservatively estimated at \$24K per minute, a traditional IP network can easily exceed \$5M per year. (see Appendix 2) Migrating to a high-availability Real-Time network can reduce downtime and associated costs by a factor of 10. Direct and indirect cost savings include reduced SLA penalties, increased customer retention and reduced customer service/repair costs.

Carriers who recognize the need to upgrade from a best effort network to the Real-Time network can achieve a competitive advantage through higher revenues, lower costs, and greater profitability.

The Real-Time IP Network


Appendix 1 - Cost of Microbursts

Users impacted by microbursts per month

Service	Growth/yr	2000	2001	2002	2003	2004	2005
Interactive Gaming	53%	5,584	8,544	13,072	20,000	30,600	46,818
VoIP	72%	3,930	6,760	11,628	20,000	34,400	59,168


Microburst Financial Impact on Carriers:

	Cost per User Impacted	2000	2001	2002	2003	2004	2005
Interactive Gaming	\$ 10.00	\$670,096	\$1,025,247	\$1,568,627	\$2,400,000	\$3,672,000	\$5,618,160
VoIP	\$ 14.00	\$660,319	\$1,135,749	\$1,953,488	\$3,360,000	\$5,779,200	\$9,940,224
Total		\$1,330,415	\$2,160,996	\$3,522,116	\$5,760,000	\$9,451,200	\$15,558,384



Comparative Microburst Costs

	2000	2001	2002	2003	2004	2005	Avg Protocol Convergence Time (secs)
Traditional Ntwk	\$1,330,415	\$2,160,996	\$3,522,116	\$5,760,000	\$9,451,200	\$15,558,384	20
Avici Router w. RAPID	\$66,521	\$108,050	\$176,106	\$288,000	\$472,560	\$777,919	1



Sources:

VOIP Growth Rate - Insight Research Corporation, October 23, 2002

http://www.insight-corp.com/pr/10_23_02.asp

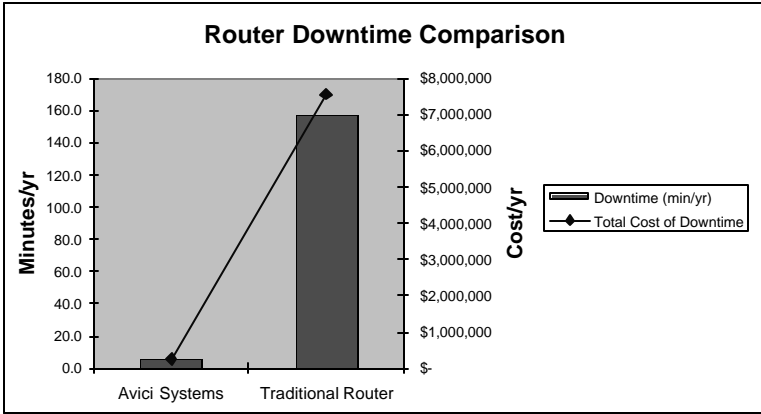
Interactive Gaming Growth Rate - Executive Summary Consulting, Inc., "The State of Massive Multi-Player Online Games 2002"

<http://www.zona.net/news/2002mmogreport.html>

Users Impacted and Cost of Impact, 2003 - RHK "The Coming Era of Absolute Availability" , 2003

Appendix 2 - Cost of IP Network Downtime

	System Availability	Downtime (min/yr)	% of Service Impacting Outages	Downtime Cost/min**	# of Routers in Network	Total Cost of Downtime
Avici TSR/SSR/QSR	99.999%	5.3	10%	\$ 23,966	20	\$ 252,098
Traditional Core Router	99.970%	157.8	10%	\$ 23,966	20	\$ 7,562,942



Sources:

Traditional Core Router Availability - RHK, "The Coming Era of Absolute Availability", 2003
 Downtime Cost per minute - Network Strategy Partners, "Reliable IP Nodes", 2002 (average availability used)