Business Benefits of Modular Software in Carrier Ethernet Routers

NETWORK STRATEGY PARTNERS, LLC
MANAGEMENT CONSULTANTS TO THE NETWORKING INDUSTRY
www.nspllc.com
September, 2005

NETWORK
STRATEGY
PARTNERS

Network Strategy Partners, LLC (NSP) — management consultants to the networking industry — helps service providers, enterprises, and equipment vendors around the globe make strategic decisions, mitigate risk and affect change through custom consulting engagements. NSP's consulting includes go-to-market strategies, development of new service offers, pricing and bundling as well as infrastructure consulting. NSP's consultants are respected thought-leaders in the networking industry and influence its direction though confidential engagements for industry leaders and through public appearances and trade magazine articles. These interactions assure NSP's clients that they will be among the first to know the latest industry concepts and emerging technology trends. Each consulting engagement is uniquely structured—no forced methodologies or canned reports are employed. NSP's consultants' collective experience is derived from leading firms across a broad spectrum of professional disciplines including management consulting, engineering, marketing, financial analysis, and IT management. Contact NSP at www.nspllc.com.

Table of Contents

Introduction

The service providers' adoption of MPLS/VPLS-based Ethernet as their network edge infrastructure is taking place in a manner very similar to that of ATM 10-15 years ago. This is due to Ethernet's ability to run different types of traffic (ATM, MPLS, IP, SONET, etc.), deliver the reliability expected in a carrier environment, and scale to support growing bandwidth needs and millions of subscribers. At the same time, the cost of an Ethernet-based infrastructure is substantially lower than a traditional edge solution both in terms of the initial capital expenditure and the ongoing operational costs. These qualities all make Ethernet a natural choice for the edge and aggregation infrastructure of residential triple play services.

The carrier Ethernet edge router is the key to the delivery of today's services, which include residential triple play, VoIP, layer 2 VPNs including VPLS, and emerging wireless services such as WiMAX, EV-DO and 3G. It links the entire carrier edge to the core network and controls service definitions and QoS for multiple services that are delivered to thousands of customers over a converged network infrastructure. As such, the Ethernet edge router must meet the standards of true, carrier-grade availability. It must also be able to provide very rapid service restoration to meet the exacting requirements imposed by the real-time nature of voice and video services. The reliability of the Ethernet edge router, consequently, has a vital impact on the carrier's business case.

This report analyzes the relationship between the reliability of edge routers and carriers' ability to attract and retain customers by providing them with IPTV that doesn't pixelize, clear phone calls, and always-available high-speed VPN and data connections. We refer to these reliability expectations as the end-users' Quality of Experience.

The report defines the availability requirements for Ethernet routers, calculates the availability of routers using traditional monolithic software architecture and that of routers using modular software, and compares the business case impacts of the two software architectures.

Key Findings

- Ethernet edge routers must deliver availability of five 9s to meet consumer Quality of Experience expectations for voice and video service. Modular router software meets this requirement with 0.999995 availability while traditional router software falls a full order of magnitude (0.99992 availability) below the quality expectation.

- Traditional router software's failure to meet availability expectations can potentially dampen user demand by as much as 21 percent for voice and 22 percent for video services. This has a ripple effect on the triple play wherein the success of bundles is dependent on the take up rates of individual services.

- The greater availability delivered by routers using modular software can potentially increase the average revenue per user (ARPU) by upwards of 20 percent.

- Modular software's delivery of five 9s availability directly reduces churn by eliminating service quality as a point of consumer dissatisfaction and, indirectly, by supporting higher take-up rates for rich service bundles that, in and of themselves, have substantially lower churn levels. Low churn, in turn, assures successful triple play projects by driving lifetime customer account values above the breakeven threshold.

- Modular software's ability to eliminate nearly all major downtime impacts can eliminate millions of dollars per year in network operations cost for a typical carrier edge network

Service Availability and Quality of Experience

Quality of Experience connotes the satisfaction end-users take in the quality and availability of their multimedia experiences. End users judge Quality of Experience by comparing a new voice, video or Internet service to that which they have deemed the standard. This standard, in turn, dictates the reliability requirement of the underlying network infrastructure.

Plain Old Telephone Service (POTS) is the standard against which consumers compare all other services. POTS' overall reliability is primarily determined by the reliability requirement for the Class 5 central office switch. The reliability requirement for the Class 5 Switch is three minutes downtime per year or 0.999994 availability. This has been built into telephone circuit switches for more than fifty years.

In today's MPLS and IP-based triple-play networks, the carrier Ethernet edge router (in combination with a softswitch) plays a role that is analogous to that of the Class 5 switch in a circuit switched network, providing switching and service management functions to tens or hundreds of thousands of customers. These customers don't care over which type of network their voice service comes just as long as their Quality of Experience is not impaired. Since their Quality of Experience for voice is based on the five 9s reliability of the Class 5 circuit switch, and since these customers don't care whether the network is circuit or packet, the edge router in a packet network must also provide five 9s reliability.

This requirement for edge routers with five 9s reliability is equally, if not more, important for the delivery of video services. Customers have the same Quality of Service expectation for video services regardless of the delivery vehicle just as they do for voice services. IP-based video introduces an additional network element, the edge router, into the network that does not exist in legacy Cable networks. The edge router's reliability consequently must contribute minimal downtime risk for IP-based video to provide Quality of Experience comparable to that of legacy video services. Testing has shown that consumers are *ten times* more sensitive to packet loss for video services than they are for voice services.[1] Even the shortest router outages are particularly damaging to TV viewers' Quality of Experience, causing jitter, freezing, pixelization and macro blocking among other things. What's more, since 37 percent of people in the U.S. watch more than two hours of TV per day[2] chances are many customers will witness any given incident of service impairment.

Internet service, while not currently achieving the availability of voice and video services, also faces increasing consumer Quality of Experience expectations. Table 1 compares the availability of telephone circuit switches, IP-VPN service and the measured performance of routers on the Internet.

---

[1] Ciena citing a customer lab test as reported by *Heavy Reading*

[2] "Middletown Media Studies: Media Multitasking … and how much people really use the media"; R. A. Papper, M. E. Holmes, M. N. Popovich; *The International Digital Media & Arts Association Journal* Vol. 1 No. 1; Spring 2004

Table 1: Service Availability

Service/Network Element	Availability	Downtime per Year
Telephone Circuit Switch	0.999994	3 min
IP-VPN (SLA) [3]	0.999	526 min
Internet (Network Component) [4]	0.99	5,256 min

Table 1 shows that the network component of web browsing transactions on the Internet currently delivers much lower reliability than IP-VPN service and circuit switches. However, the rise of telecommuting and home businesses, for example, is making reliable Internet access more important. A teleworker's ability to link into his company's VPN is essential to his Quality of Experience.

Entertainment services, also, are driving the need for greater reliability. With online gaming, for example, the game players' Quality of Experience is dependent upon the network providing sub-second response to game controller commands. The increasing popularity of Internet-based video downloads and streaming video will put further upward pressure on consumers' Internet Quality of Experience expectations.

Router Availability Requirements for Triple Play Services

Premium service offerings—such as those delivered by FTTP access networks and Ethernet routers—must provide a high Quality of Experience to induce consumers to switch from their existing service offerings. Both top line and bottom line business cases are impacted if these reliability expectations are not met. Router availability that is below expectations will cause take-up rates for triple play services and rich service bundles to be low, ARPU to drop, and churn to increase. Low availability also increases the work and cost of the network operations organization and degrades the value of the service for Ecommerce and content providers.

Given the above, the edge router must meet three types of reliability requirements to support successful triple play services:

- Service availability must be at least as good as that of existing services, which constitute the consumers' point of reference.

- Router availability must be sufficient to meet reliability expectations for individual services while delivering service bundles over IP— within a household, each individual service must meet expectations while multiple services are simultaneously in use.

[3] See http://global.mci.com/terms/us/products/vpn/dedicated/

[4] "Measuring End-User Availability on the Web: Practical Experience"; Matthew Merzbacher and Dan Patterson; *International Performance and Dependability Symposium*; Washington, DC; June 2002

- Router outages must be quickly resolved to achieve good Quality of Experience for real-time services such as video and voice.

Edge router availability must meet or exceed five 9s availability.

Comparison of the Availability of Modular versus Traditional Router Software

The following sections compare the availability delivered by modular versus traditional router software. The analysis focuses on software rather than hardware or network reliability because software provides the functionality used to deliver the very high availability needed to meet consumers' expectations. For example, networks are usually designed with diverse network links to provide an alternate transmission path in case the primary link fails. It is the router software, however, that determines whether the alternative link can be activated in time to avoid a service-affecting outage. Similarly, routers are usually equipped with redundant processors, but it is the router software that determines whether the standby processor must be rebooted before restoring service or can recover the system state from the primary processor.

Traditional router operating systems are monolithic in nature. Traditional routers, and by extension many popular Ethernet switches as well, have an operating system in which all system functions are closely intertwined and not well partitioned. The benefit of this monolithic system is that it can help performance, especially with non-optimized software code. The drawback is that a failure of one function will crash the entire operating system, regardless of how small the function is.

By comparison, a modular operating system breaks down different router functions into self-contained processes. Each process operates independently within its own protected address space. Typically, each major function will have its own process. For instance, routing protocols like OSPF would be separate from system configuration, which would be separate from system monitoring, and so on. This independence is coupled with the ability to restart any of these processes if it fails without needing to restart all other running processes. With functions very modular in nature, it is also easy to make changes to a single function or to add new functions without forcing changes to all other functions in the system. As a result, it is easier to develop new functions and to fix bugs within existing systems.

Modular router software that has distributed control planes is designed to both eliminate router outages and reduce the downtime impact of an outage. This type of software design is specifically critical for smaller systems that do not support hardware redundancy. Among the key features of modular software are:

- The ability to isolate software failures
- Automatic Process Restart, which permits restoration of software functionality without service impacts
- Hitless Protection System for VPLS/MPLS deployments, which permits the restart of protocols such as LDP, RSVP, OSPF and ISIS without affecting service delivery

- MPLS Fast Reroute, which provides a vehicle for re-routing MPLS traffic around a network fault analogous to that used by SONET/SDH and approaching SONET/SDH's 50 ms restoration time

Table 2 provides a detailed list of modular software's design features and describes how they reduce downtime impacts for particular outage categories.

Table 2: Modular Software Features for Reducing Router Downtime

Feature	Benefit
Automatic Process Restart (APR)	Restart Process without Hard Reset
Hitless Protection System (HPS) for VPLS/MPLS	Graceful protocol re-start for LDP, RSVP, OSPF, ISIS—eliminates a Hard Reset
MPLS Fast Reroute	Reduces Link Failure impact—restoration approaching 50 ms
Loop Detection & MAC move detection	Helps avoid routing loops—eliminates a software configuration outage
Traffic Management	Keep links from being overloaded by distributing traffic across multiple links—reduces software control plane failures
No single process brings down Line or Main CPU	Keeps processors running—reduces control and data plane software outages
Monitor daemons restart unhealthy processes	Improves software performance—reduces control and data plane software outages
Hot standby control module	Improves hardware & software availability—reduces software & hardware/control outages
Hitless upgrade of software modules	Eliminates Hard Reset for software upgrade—eliminates software update outages
Complete separation of data and control plane	Reduces Hard Reset events—eliminates simultaneous control plane & data plane failures
Exception data packets handled by Line Card CPU	Helps minimize DOS effects—reduces software control plane crashes
Main CPU does not receive data packets	Helps minimize DOS effects—reduces software control plane crashes
Out of Band communications between line and central CPU	Helps minimize DOS effects—reduces software control plane crashes
All CPU bound packets are rate limited to prevent DOS	Helps minimize DOS effects—reduces software control plane crashes
Virtual memory OS on all CPU	Protects entire system from coming down—reduces software control plane crashes
Standby CPU can recover state from Primary or Line CPU	Reduces reset times—reduces severity of all outages

Probability Tree for Calculation of Router Downtime

In this section, we construct a probability model that measures the downtime impact associated with ten outage categories for modular and traditional router software. The model's findings are used to compute router availability and assess the effect of improved availability on the triple play business case. The ten outage categories include planned outages, such as those incurred to make hardware and software configuration changes, and unplanned outages, such as those caused by a hardware failure in the control system or a link failure. Table 3 defines the ten outage categories.

Table 3: Carrier Edge Outage Categories

1. *Software Update/Configuration Induced Outage*: Outage caused by any upgrade or configuration of the router OS. Includes patches, new software releases, bug fixes, routing/policy changes, etc.
2. *Hardware Upgrade/Configuration Induced Outage*: Outage caused by the addition, replacement or reconfiguration of hardware not due to failure. Includes capacity upgrades, performance enhancements, or adds/moves/changes, etc.
3. *Software Failure – Control Plane*: A failure of any of the routing control protocols that reside in the control plane. Includes unicast, multicast, and MPLS protocols, etc.
4. *Software Failure – Data Plane*: A failure of any software controlling the data plane. Includes any microcode or software that controls IP forwarding ASICS, processors, or other data plane hardware.
5. *Software Failure – Other*: A failure of any other aspect of a router OS that leads to downtime. Includes management failures, OSS integration bugs, router hangs, etc.
6. *Hardware Failure – Control/System*: A failure of any of the common components in a router. Including router processors, switch fabrics, clock schedulers/timers, fans, power supplies, etc.
7. *Hardware Failure – I/O Card*: A failure of a particular line card, either from core or edge uplinks/downlinks.
8. *Link Failure*: A failure of edge-core, core-core, core-peer network links. Includes fiber cut, transmission equipment failure, etc.
9. *Power Outage*: A failure of the router or node due to loss of electrical power
10. *Other/Unknown*: Failures that were undeterminable or too small a percentage to categorize.

The probability model employs an engineering assessment of each software design's approach to resolving an outage. A planned software upgrade using traditional software, for example, nearly always requires a hard reset (re-boot) of the router. Modular software's in-service software-upgrade capability, on the other hand, generally permits software to be upgraded without affecting service. The probability model, therefore, is built by comparing each software design's features and approach to resolving each of the ten outage categories.

The availability delivered by each software design is calculated by using a tree of conditional probabilities. Each path from the base of the tree to the up-most branch describes one possible outcome of a chain of events that begins with the occurrence of an outage as follows:

Outage Category → Router Resolution → Resolution Outcome → Downtime Impact

The event sequence begins with a prediction that an outage of category X has occurred where the relative frequency of the outages is derived from academic studies of router downtime and interviews with network operations personnel.[5] Next, engineering analysis is used to determine whether a hard reset, soft reset or no reset will be needed to resolve the outage. This is followed by an estimate of whether the action taken will, in fact, resolve the outage. Finally, the downtime impact (major, minor, imperceptible) incurred as a result of this particular sequence is determined. There are ten outage categories, three possible router resolution actions, two possible resolution outcomes and three possible downtime impacts. One hundred eighty possible outcomes, therefore, can occur each time an outage event takes place. Two probability trees are constructed, one for each software design, each of which models the 180 possible outcomes.

Exhibit A, which illustrates a portion of the probability tree, depicts the full sequence of events for a category 5 outage that can be resolved by a Soft Reset using traditional software. The numbers associated with each branch of the tree are the conditional probabilities that this outcome will occur given the event sequence charted on the tree. So, for example, given that an outage has occurred there is a 0.05 probability that the outage will be category 5, software failure – other.

[5] "Experimental Study of Internet Stability and Wide-Area Backbone Failure", University of Michigan

Exhibit A

Probability Tree for Traditional Software (Partial)

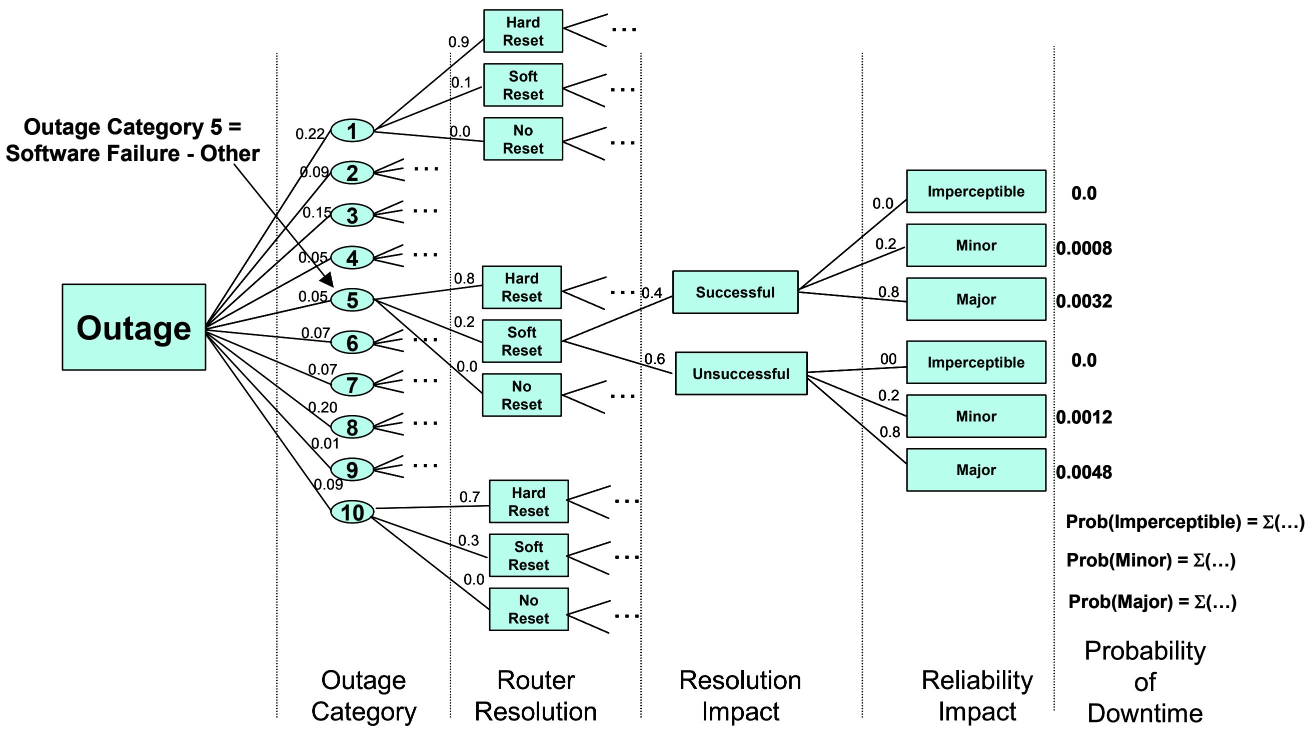

The probability of a particular outcome is determined by tracing a path from the left to the right and then multiplying the probabilities together. For example, for traditional software, the probability that an outage will be Category 5, will require a soft reset to resolve it, will be resolved successfully, and will have major downtime impact is:

Prob[Software-other (0.05) X Soft Reset (0.2) X Successful (0.4) X Major (0.8)] = 0.0032

Once these calculations are completed for each of the 180 possible outcomes, the probability of imperceptible, minor and major downtime impacts are calculated by summing the individual results (shown as the last column in Exhibit B) for each of the three possible levels of downtime impact. The same type of calculation is done for modular software. However, the conditional probabilities change to reflect the way particular modular software features address each outage type. In the example discussed above, modular software requires neither a soft nor hard reset to resolve a Category 5 outage, so the downtime impact is imperceptible in all cases.

Comparison of Downtime Impacts

Table 4 compares the downtime impacts for the ten outage categories for traditional versus modular and distributed router software.

Table 4: Downtime probability per outage category

Outages			Probability Tree Results by Outage Category						
No.	Outage Category	Prob.	Traditional			Modular & Distributed			
			Imp.	Min.	Maj.		Imp.	Min.	Maj.
1	Software Upgrade/Config	**22.0%**	0.000	0.004	0.216		0.198	0.022	0.000
2	Hardware Upgrade/Config	**9.0%**	0.073	0.009	0.009		0.090	0.000	0.000
3	Software Failure - Control Plane	**15.0%**	0.000	0.006	0.144		0.143	0.008	0.000
4	Software Failure - Data Plane	**5.0%**	0.000	0.002	0.048		0.050	0.000	0.000
5	Software Failure - Other	**5.0%**	0.000	0.002	0.048		0.050	0.000	0.000
6	Hardware Failure - Control/System	**7.0%**	0.000	0.000	0.070		0.070	0.000	0.000
7	Hardware Failure - I/O Card	**7.0%**	0.056	0.014	0.000		0.056	0.014	0.000
8	Link Failure	**20.0%**	0.000	0.000	0.200		0.000	0.200	0.000
9	Power Outage	**1.0%**	0.000	0.000	0.010		0.000	0.010	0.000
10	Other/Unknown	**9.0%**	0.000	0.005	0.085		0.081	0.005	0.005
	Total	**100.0%**	**0.129**	**0.042**	**0.829**		**0.738**	**0.258**	**0.005**

Table 4 shows that modular software has a dramatic effect on major downtime impacts—they are reduced from a probability of 0.829 for traditional software to only 0.005 for modular software. Correspondingly, imperceptible downtime impacts are infrequent (0.129) for traditional software and frequent (0.738) for modular software. These reductions in downtime are achieved in each of the ten outage categories. Notably, the most common outage category, software upgrade/configuration, entirely eliminates the major downtime impacts nearly always incurred with traditional software by application of modular software's hitless upgrade feature. The second most common outage category, link failure, also sees an improvement of the downtime impact from major to minor through application of modular software's MPLS Fast Reroute feature.

Table 5 restates these findings in terms of the average downtime impact per outage measured in seconds.

Table 5: Average Downtime per Outage

Downtime Impact	Average Downtime (Seconds)	% of Outage Traditional	% of Outage Modular
Imperceptible	0.05	12.9%	73.8%
Minor	5.00	4.2%	25.8%
Major	30.00	82.9%	0.5%
Average Downtime (Seconds)		**25.09**	**1.46**
Downtime Reduction			**94%**

The average downtime per outage is reduced from 25.09 seconds with traditional software to 1.46 seconds with modular software—a 94 percent reduction in downtime. Therefore, if 100 outages occur each year traditional software has 0.99992 availability (42 minutes of downtime per year) while modular software has 0.999995 availability (3 minutes of downtime per year.)

Impact of Router Availability on the Triple Play Business Case

The impact of service availability on end-users' Quality of Experience makes it a critical component in the success or failure of a triple play deployment. Due to its greater reliability, an edge router with modular software and a distributed control plane can make the difference between a successful or unsuccessful triple play deployment. The triple play business case is predicated on the increased user take-up rates, ARPU, and customer loyalty that result from the bundling of services and superior performance. This, in turn, is dependent upon customer satisfaction and Quality of Experience as discussed in the "Router Availability Requirements for Triple Play" section of this report.

For telecom providers, the key strategy is to add IP video to the service bundle. This has the dual advantage of creating a new revenue source and making the voice, data, and video bundle so attractive that existing voice customers choose to keep their voice service.

The business case includes the following elements:

- Service bundles produce more revenue than standalone service offerings
- The fixed cost nature of the triple play access network requires high service take-up rates and high ARPU to recover construction costs
- Service bundles and high Quality of Experience reduce churn
  - Cost to acquire customer must be recovered
  - Additional sales to existing customers are less expensive than sales to new customers
  - Service stickiness involves greater perceived value that supports pricing premiums
- Low operating expense is needed to achieve profitable growth and recover construction costs

Service Bundles and Take-up Rates

Edge router availability directly impacts take-up rates for new and bundled services[6], which allow service providers to grow their overall revenues and increase ARPU.

Bundling allows customers to buy multiple services for less than the cost of buying services *a la carte*. The ARPU increase comes from selling more services overall as a result of the bundles, which provide an incentive to pay a small increment for a third, fourth or fifth service[7]. This is all about perceived bang for the buck, whether or not the additional services are fully utilized by the customer or not. Customers will chose to pay *a la carte*, however, if the additional services are unreliable.

Though high availability alone does not provide an incentive to subscribe to a service, low availability will deter customers from adopting the new offerings. The triple play requirements section of this report established a router availability requirement of 0.99999 (five 9s) for voice and video services. As previously shown, modular software meets that requirement while traditional software, calculated at 0.99992, does not.

---

[6] For examples of typical service bundles see RCN's "Essentials" bundles - http://www.rcn.com/bundles/

[7] Many carriers confirm this result. For example, UBS Warburg found that Cox Communications' bundled subscribers have ARPU of $130 compared to the cable industry average of $62. "Broadband Applications Fueling Customer Demand: Assessment of Market Drivers and Opportunities"; Allison Cerra, Adam Green; *Broadband Services, Applications, and Networks: Enabling Technologies and Business Models*; IEC; Chicago; 2004

Exhibit B compares individual service take-up rate response to router availability.

Exhibit B

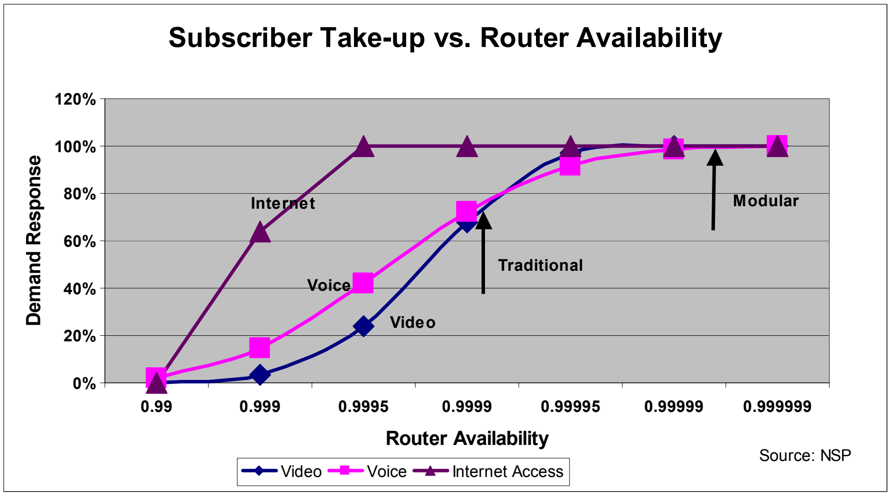

Subscriber take-up rate versus router reliability is modeled as a function of the gap by which actual router reliability either exceeds or falls short of customer expectations. For example, the availability requirement for a router supporting voice service is 0.999994—see "Router Requirements" section. Therefore, if the router delivers five 9s or better reliability the take-up rate is nearly 100% of the potential. However, if router reliability is three 9s or lower the take-up rate is near zero.

As shown in Exhibit B, traditional software's 0.99992 (42 minutes of downtime/year) availability is more than adequate for consumer acceptance of Internet-Access. Hence, the demand response is 100 percent of potential. However, this same level of reliability is nearly an order of magnitude below that required for voice and video acceptance, so only 79 percent of the voice and 78 percent of the video potential is realized.

The effect of router availability takes on increased importance as service providers create and price various combinations of service bundles. Table 6 shows the effect that improved router availability has on the take-up rate and corresponding ARPU for the three-service bundle relative to the other offerings. The percentages of demand response are derived from Exhibit B. Pricing is based on average U.S. rates. Modular software greatly increases the uptake of the three-service bundle and the corresponding ARPU derived there from.

Table 6: Take Rates and ARPU for Service Bundles[8]

	Video + Phone + Internet	Video + Internet	Phone + Internet	Video + Phone	ARPU
Bundle Price/Month	$ 97	$ 85	$ 57	$ 60	
Traditional	20%	10%	65%	5%	**$ 68**
Modular	60%	20%	15%	5%	**$ 87**
				Percent Increase	**28%**

Churn

Improving router availability so that it exceeds the requirements for voice and video services has a dramatic effect in reducing churn and consequently improving the lifetime value of a customer account. Improved availability reduces churn in two ways. First, it eliminates Quality of Experience concerns caused by router outages and poor router performance. Since end-users, as a result of their wireline voice experience, have come to expect five 9s reliability for phone calls, VoIP services and the routers supporting them, must be equally reliable. Consumers' likely reaction to poor quality can be gauged by examining the mobile telephone market where Quality of Experience remains an issue. For example, Harris Interactive, the market research firm, found in a recent survey (1/26/05) that 22 percent of mobile telephone users would switch carriers to get better reliability. The second way that improved router availability reduces churn is by increasing the take-up rate and richness of service bundles because bundles have a dramatic effect on customer loyalty. Here are several findings on the value of service bundles.

- "When a local or long-distance customer adds just one additional service—DSL, wireless or dialup—churn decreases by about 45 percent" – BellSouth
- "Customer churn in two-product households is 18 percent lower than in one-product households." "Customers with three products—cable, high speed Internet and telephone—have a 48 percent lower churn rate." – Cox Communications
- "Bundling drives retention and higher average monthly revenues. With SBC's bundled services customers typically spend more than twice as much as unbundled customers." – SBC[9]
- "Our research shows that satisfaction levels and, consequently, client 'stickiness' rise when companies purchase bundled products and services from the same carrier—by

[8] The service mix is modeled by extending the concept shown in Exhibit B to include the more severe reliability expectations imposed on a triple play service bundle. (Take-up rates are not inhibited when router reliability exceeds customer requirements while take-up rates are decreased where customer requirements are not met.) The percentage mix of service bundles (absent router reliability demand suppression) is based on the annual reports of service providers' revenue mix and NSP's studies on behalf of its service provider clients.

[9] *The Quarter in Focus,* SBC 4/21/04

nearly 40 index points for long distance and nearly 60 index points for local service." – J. D. Power and Associates

- "Our experience indicates that our bundled customers are significantly less likely to churn, and we experience less price erosion when we effectively combine our offerings." – General Communication, Inc.[10]

Exhibit C illustrates the importance of reducing churn for the triple play business case, showing theoretical lifetime net revenue for a customer account versus that of the account lifetime.

Exhibit C

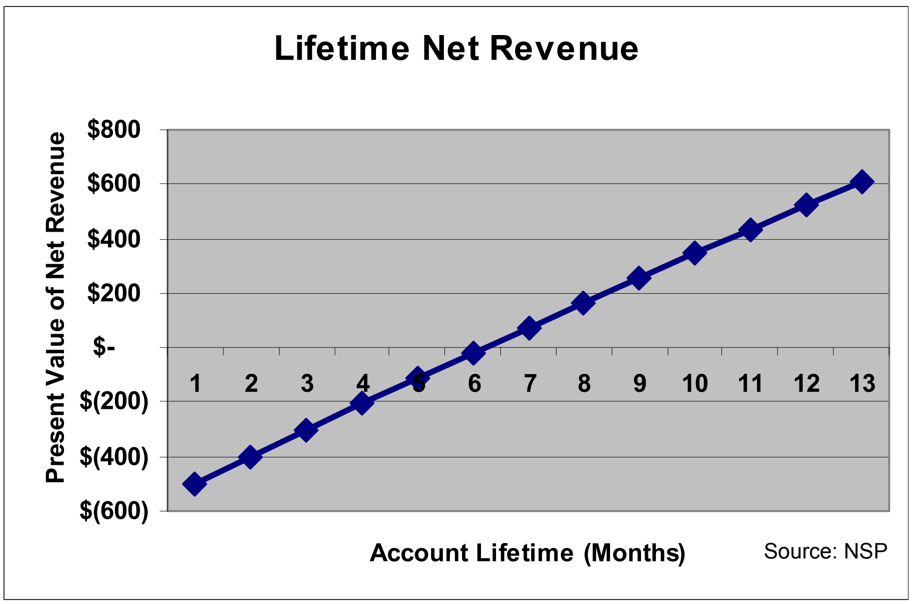

The highly simplified example above assumes account acquisition costs of $500, ARPU of $100/Mo., and a 15 percent annual discount rate. By these calculations, customer accounts of less than six months duration make a negative contribution to net revenue. Hence, customer accounts only begin making contributions to offset triple play construction costs beginning in the seventh month. If the provider's initial start-up costs for the triple play infrastructure are approximately $1,200 per subscriber, each customer must be retained for at least 18 months for the provider to recoup its investment. (Actual time-to-breakeven will likely be even longer since this example does not address ongoing operations costs such as equipment maintenance, billing expense and interconnection charges.) High churn, consequently, destroys the triple play business case.

[10] 10-K; General Communication, Inc.; Feb. 24, 2004

Operations

Cost efficient operations are needed in addition to high take-up rates and ARPU for a successful triple play business case. Router outages have a direct impact on network operations expense.

The "Comparison of Downtime Impacts" section of this report shows that under traditional software 82.9 percent of outages result in a major downtime impact while only 0.5 percent of outages incur a major downtime impact under modular software. Each major outage triggers a great deal of work within the network operations organization because the outages will re-occur if root causes are not found.

Each major outage triggers a series of actions in the NOC, field service organization and second and third-level support for the network management process— Fault, Configuration, Accounting, Performance, and Security Management (FCAPS).

Table **7** details the work processes that are triggered by each outage that causes major downtime impact.

Table 7

Operating Activity Analysis of Outage Related Work[11]

Network Management Processes	Hours per Major Outage		
	Engineer	Tech I	Tech II
Labor Rate ($/Hr.)	$ 119	$ 120	$ 121
Capacity Planning	2		
Network Administration	4		
Service Provisioning & Configuration	4		
NOC Support	8	40	40
Field Support	4	9	
Total Hours	**22**	**49**	**40**
Cost	$ 2,618	$ 5,880	$ 4,840
Total Cost per Outage			**$ 13,338**

Table **7** estimates the cost of each such major downtime event to be $13,338. The cost may be incurred in two ways. First, it may cause overtime charges and additional consultant and vendor fees. Often, however, the outages take priority over proactive work such as capacity planning, performance monitoring, security planning and routine maintenance. These proactive network

[11] These estimates are derived from workflow analyses conducted by Network Strategy Partners with our clients' network operations organizations.

management activities are the hallmarks of an operationally mature organization. If most staff time is consumed by the reactive tasks associated with diagnosing and correcting outages, then the organization devolves into an operationally immature organization characterized by poor performance, high cost, poor morale, high-staff turnover and more outages.

Modular software nearly eliminates major downtime impacts caused by outages. This in turn greatly reduces outage related operating expenses. For example, if 100 outages occur in one year then modular software will incur a major downtime impact once every two years on average (0.5% probability of a major downtime impact per outage event) whereas traditional software will incur 83 major downtime impacts per year (82.9% probability of a major downtime impact per outage event.) At $13,338 per major downtime outage events, this implies $6,669 per year in outage related expenses for modular software compared to $1,107,054 per year for traditional software.

Conclusion

Ethernet routers used to provide triple play services must meet availability requirements comparable to those of central office circuit switches (0.999994) to satisfy consumers' expectations for voice and video services. Failure to meet this availability level will impair many aspects of the triple play business case.

Routers using modular software are estimated to achieve availability of 0.99999—close to the required level, while routers using traditional software are estimated to achieve 0.99992 availability—an order of magnitude below the availability requirement. Modular software achieves the required reliability through a set of features that resolve most (73.8 percent) outages with imperceptible downtime. In contrast, traditional software incurs a major downtime impact for 82.9 percent of all outages.

This improved availability associated with modular software helps assure the success of the triple play business case by shortening the payback time of a triple play project, increasing ARPU, decreasing churn and helping maintain a more operationally mature and efficient organization.

Appendix

The exhibit (below) provides an example of how delay impacts Quality of Experience, using the Mean Opinion Score (MOS) to rate user satisfaction with a telephone call. Scoring is done on a five-point scale where scores of four or better indicate a satisfactory rating. As illustrated below, end-to-end delay is a primary contributor to a low MOS score.

Delay Impairment of Reference Connection[12]

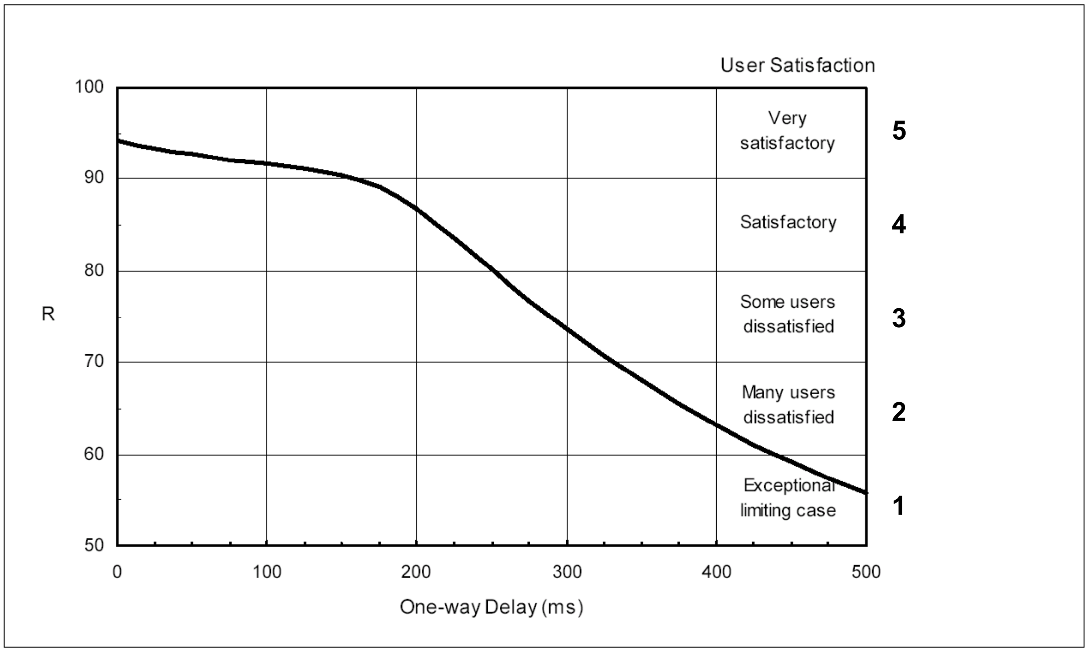

[12] Voice Quality Recommendations for IP Telephony, TIA/EIA TSB116, March 2001