June 4, 2001

Transforming Core IP Architectures

Prepared for:

PLURIS

Colorado Office
1600 Stout Street
Suite 1510
Denver, CO 80202
303.446.0001

Oklahoma Office
1307 South Boulder Avenue
Suite 120
Tulsa, OK 74119
918.382.0007

Texas Office
2605 Sagebrush Drive
Suite 220
Flower Mound, TX 75028
972.899.0086

www.telechoice.com

SW437001

The Strategic Catalyst
TeleChoice
for the Telecom Industry

Table of Contents

Transforming Core IP Architectures

Current core routing platforms leave service providers with limited solutions for supporting IP-based services. Service providers often contend with these platforms' lack of scalability, sub-optimal performance, low levels of reliability, complex network architectures, complicated network management, high power and space consumption, expensive deployments, and rapid technology churn. These shortcomings make it nearly impossible for service providers to efficiently and cost effectively keep pace with the phenomenal growth in IP traffic.

Fortunately, new technology innovations in IP core routers have purpose-built architectures that address these service providers' most crucial networking challenges.

Forces Affecting Service Providers and Their Core Networks

At the highest levels, forces such as economics, technology, competition, and usage demands put pressure on service providers and largely influence their behaviors, as shown in Figure 1.

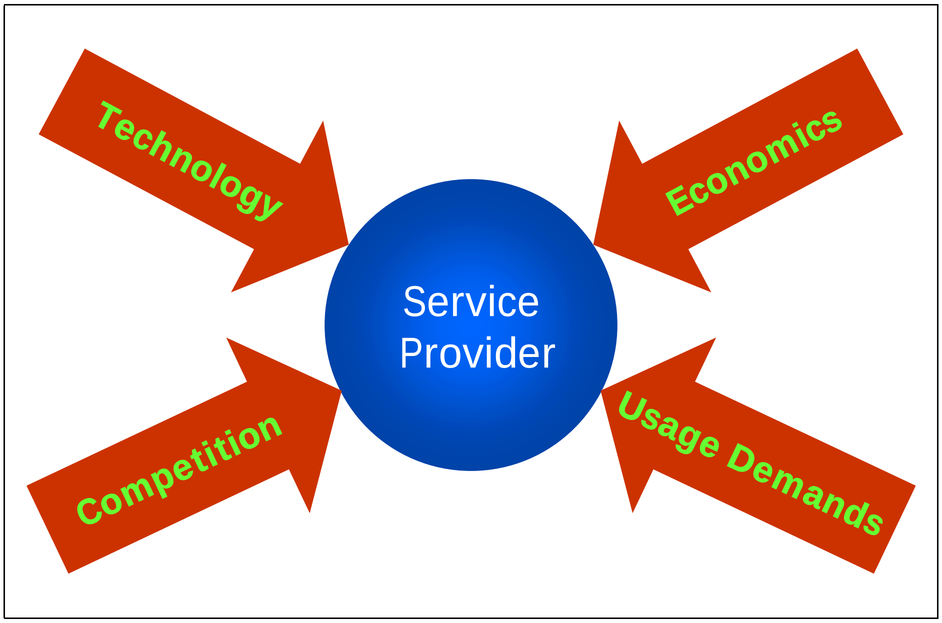

Figure 1. Forces Affecting Service Providers and Their Core Networks

- **Economics.** In the past, the market quickly rewarded service providers who deployed new technology because of their promise to achieve aggressive revenue projections regardless of profitability. Today's economy and market focus primarily on sustained profitability, putting pressure on service providers to show meaningful reductions in capital and operational expenses and to quickly generate new and incremental revenues.

- **Technology.** Innovations in technology are rapidly evolving, particularly in networking platforms. However, many of these innovations provide incremental technology improvements rather than order-of-magnitude improvements or technology disruptions. As a result, service providers find themselves introducing new platforms and technologies into their networks every 12 to 18 months. Most service providers get stuck in an endless cycle of learning, evaluating, integrating, deploying, and supporting new platforms.

- **Competition.** The hypercompetitive environment forces service providers to make decisions and act quickly to capture market share in a new market, build barriers to entry for other competitors, protect their customer base, or stay ahead of their competitors. Therefore, service providers have no choice but to deploy the best platforms available to them at any given time, even though these solutions have quick churn cycles and require frequent forklift upgrades.

- **Usage Demands.** The development of new applications, new means of communication, and new ways to conduct business have significantly increased the demands placed upon service provider networks. The virtualization of how people and machines interact with one another demands increased network traffic and reliance. As a result, service providers feel more pressure to deliver the most reliable, most scalable, and best-performing services. The impact of network failures today compared to 10 to 15 years ago is magnified because of the number of applications and high volumes of traffic the network now supports. Today’s companies have near-zero tolerance for network failures. Although service providers feel these pressures in all aspects of their business, they feel more immediate effects on the core IP network and platforms deployed.

Current Backbone Architectures Supporting IP

The most common architectures supporting IP traffic are (1) IP over ATM over optical and (2) IP over optical. Examining these alternatives provides a baseline for understanding service providers’ existing IP backbone architecture challenges.

IP over ATM over Optical

Many service providers use ATM for interconnecting their core IP routers. ATM’s inherent connection-oriented paths make the IP routers appear one hop away from each other even when the traffic traverses several ATM switches. ATM’s faster interfaces, traffic management, aggregation, quality of service, and restoration features help service providers improve the performance, reliability, and availability of the network.

ATM switching has long supported and augmented the IP network; however, current IP traffic growth rates pose questions regarding ATM’s ability to scale and support large volumes of IP traffic. In addition, ATM’s fixed cell length is inherently inefficient for carrying variable length IP packets. Fortunately, innovations in routing technology have begun to address the efficiency and scalability issues that plague ATM, hence allowing IP routing to become a viable technology in the core of the network.

IP over Optical

Most service providers build tiered IP networks where each layer of the network consists of function-specific routers built to reside at a specific network layer. Some routers are built and better suited for the premises, access, edge, or core layers. Pure IP networks became possible as IP evolved to provide the same functions as ATM, offering the same or better performance levels.

Pure IP networks primarily rely on the routers to perform a mix of transmission and service-enablement functions. In the premises, access, and edge layers, routers typically have higher degrees of service intelligence, subscriber services, and aggregation capabilities for lower-speed connections. In the core, the routers are built for high-speed packet processing and forwarding, fast convergence, efficient transmission of millions of traffic flows, and interoperability with optical networks.

Core IP Network Challenges

Traditional core router implementations have created challenges for the engineering, network planning, operations, and marketing organizations of service providers. Some of the biggest core IP network challenges for service providers include the following:

- Scaling the Network
- Improving Platform Longevity
- Improving Network Reliability and Availability
- Improving Performance
- Simplifying the Network and its Management
- Reducing Capital and Operational Expenses
- Generating Revenues from New Services
- Integrating with Front- and Back-Office Systems

Scaling the Network

Escalating traffic demands require highly scalable backbone networks. Dense Wave Division Multiplexing (DWDM) was developed to help solve the scalability issues in the optical core. DWDM increases the number of wavelengths per fiber, which translates into increases in port requirements within the IP core; however, one problem remains. Many of today’s IP routers have difficulty scaling to keep pace with the optical network. These routers may not have higher-speed interfaces, high port densities, or high-capacity switch fabrics.

Traditional core IP routers have a single switch fabric that has a capacity limitation, typically designed to support a set of high-speed line cards. However, as the industry develops faster interface speeds, routers with fixed switch fabric cannot accommodate the higher-speed line cards. Therefore, router vendors might have a family of routers, perhaps small, medium, and large, with different switch-fabric capacities and line-card speeds supported. For example, the small version might support OC-12 and OC-48 interfaces while the bigger version might support OC-48 and OC-192 interfaces. If service providers start with the small version, as traffic grows and as they need to aggregate more traffic onto a larger pipe, they would need to purchase the bigger version to augment the smaller version and to get the higher-speed OC-192 interfaces. Traditional core routing architectures force service providers to deploy a mesh of discrete routers at a particular node or POP to achieve the scalability needed. Figure 2 shows a mesh of discrete routers supporting OC-12, OC-48, and OC-192 interfaces.

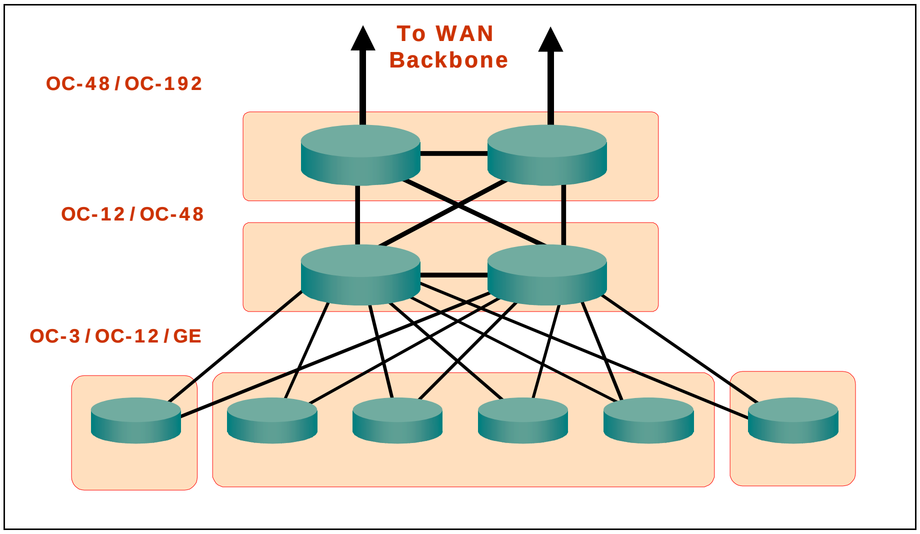

Figure 2. A Mesh of Discrete Routers

Deploying a mesh of routers consumes line cards needed to link the routers within the mesh, creating a more complex configuration more likely to cause instability in the network. When primary and backup routers are configured for high availability or redundancy, the problem is exacerbated.

A mesh of discrete routers makes it extremely difficult for service providers to scale their networks quickly and cost effectively. Service providers typically have to make the tough decision of choosing a larger platform with a higher-speed switch fabric that's expensive and is overkill for today's traffic but that allows them to scale rapidly versus a smaller, more cost-effective router that supports today's traffic adequately but has limited growth potential.

Improving Platform Longevity

Related to scalability, replacing or augmenting smaller routers with a bigger router every 12 to 18 months to keep pace with traffic growth severely disrupts the service provider's operating environment. Constant evaluation, integration, and deployment of new products is time-consuming and expensive.

Service providers need to have redeployment plans and migration strategies for the older routers, to extend the life of their current investment. The smaller routers may become part of a router mesh or be redeployed at the edge or access layers of the networks. Frequent reshuffling of equipment and redesigning of the network is a painful and inefficient undertaking.

Improving Network Reliability and Availability

Traditional routers were not built to carrier-class reliability and availability standards. As a consequence, service providers now have to build reliability into the network to overcome the reliability inadequacies of current routers. Service providers might duplicate and separate the routing functions across discrete devices. They also may design multi-homed configurations. These work-around solutions significantly complicate network design and increase implementation costs.

Most routers typically offer redundancy options for both line and control cards. However, redundancy for the control cards typically means warm-standby. When the primary control card that contains the routing tables, policies, and adjacencies fails, the backup control card then assumes the primary card's responsibilities. However, this backup card generally needs to recalculate routes and re-establish adjacencies before it will receive or send traffic again. Typically, such network reconvergences take two to three minutes, and the network sees that router as offline throughout that process. This architectural limitation also makes it impossible to do in-service upgrades, so service providers need to take routers offline to perform software upgrades.

Improving Performance

Packet processing rates can vary greatly among different router vendors. Some cannot support line-rate processing. Some can achieve line-rate packet processing but not for all packet sizes. Others can process packets at line rate but cannot sustain line-rate throughput when other features such as packet prioritization, filtering, and quality of service are enabled.

Variations in packet processing rates pose a major hurdle for service providers as they develop and launch more intelligent and more comprehensive IP services. Network and service performance should not suffer when service providers package value-added features with their transport services.

Simplifying the Network and its Management

As mentioned previously, a mesh of discrete routers and duplicate routers to improve scalability and reliability is complex and hard to manage. The service provider needs to manage each router individually within the mesh. Service providers' operations groups can manage a small number of router meshes; however, management quickly becomes a nightmare as the number of meshed routers and groups of meshed routers grow.

Reducing Capital and Operational Expenses

Service providers need to balance the competing demands of supporting traffic growth while managing capital and operational expenses. A fine line exists between purchasing enough equipment to support traffic growth and reducing capital and existing operational expenses. Previously, many service providers spent two to three years in a build-out frenzy, constructing and building "next-generation networks." Figure 3 illustrates the fundamental goal of these build-outs. The X axis and Y axis depict time and unit cost, respectively. As a specific technology evolves over time and as network utilization improves, service providers can drive their per-unit cost much lower. The parallel technology "hockey sticks" show different technologies' ability to drive down costs. Some technologies are able to drive down costs more than others; therefore, service providers are constantly searching for technologies that allow them to achieve this goal.

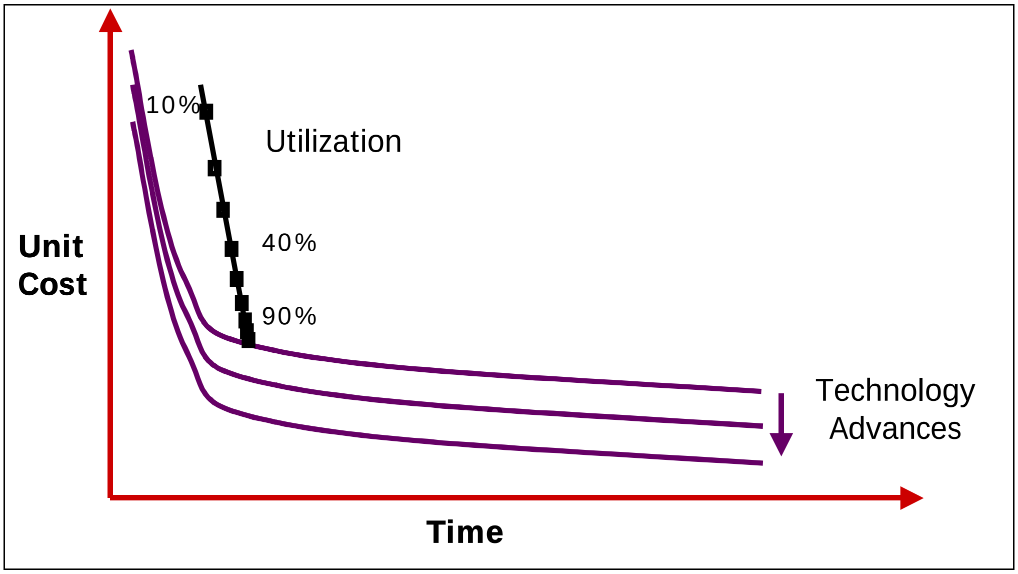

Figure 3. Fundamental Goal of Network Build-outs

Recent stock market declines have resulted in capital "belt tightening," making it more difficult for service providers to keep pace with growing traffic volumes while attempting to reduce cost. Recurring network builds and overbuilds intensify the challenges associated with scaling the network. Deploying meshed routers and duplicating routers for improved scalability and reliability is expensive, consumes significant space and power, uses up line cards for inter-router connections within a mesh, and complicates the network management and troubleshooting process.

Limitations in technology and the work-arounds needed to overcome them prevent service providers from reducing capital and operational expenses. Their remaining option is to offset additional capital and operational expenses by lowering costs in other areas or to generate more revenue. Service providers need to introduce new and differentiated services to improve their revenue streams.

Generating Revenues from New Services

Many service providers face significant challenges developing and launching new services rapidly. Among the most significant are training and building an organization to support new services, testing the technology within the network, and integrating the solution into existing billing and Operations Support Systems (OSS).

Given these challenges, service providers may need more help from their vendors to develop these new IP-based services, particularly as vendors introduce enhanced features and capabilities. For newer providers that may have layer 3 expertise, often the issue is a lack of resources. Most core platform vendors do a great job of communicating the features and capabilities of the equipment. However, very few have programs in place to help service providers develop and launch consumable services based on these features, packaged and positioned appropriately for the service providers' target markets. As a result, many service providers end up offering commodity IP services.

Without assistance, service providers take much longer to launch services, hindering their ability to achieve the higher returns on investment from offering multiple, differentiated services leveraging the same network. Furthermore, there's a higher probability of failure when competing in the commodity business.

Integrating with Front- and Back-Office Systems

Some traditional router platforms do not provide seamless and comprehensive integration with the front- and back-office systems. Most routers can provide information to network management systems, Service Level Agreement (SLA) reporting, and trouble ticket systems. Fewer router vendors integrate and communicate with order entry, provisioning, and billing systems; therefore, service providers need to proactively monitor data integrity to eliminate errors and the risks associated with manual entries.

Guidelines for Evaluating Core Routing Platforms

With new innovations in technology, service providers can now expect features and capabilities not readily available in the past. The following list includes some key evaluation criteria for choosing an IP core routing platform:

- **Higher-Capacity and Higher Port Density Routers.** Single-chassis-based platforms provide an effective solution to support rapid growth, offering high port densities and high-speed capacity while saving on space and power requirements and simplifying the network. However, service providers currently or will eventually have the need to deploy multiple chassis within a single node or POP to support increasing IP traffic. Service providers need to evaluate the platform's inherent ability to interconnect multiple chassis at high speeds (i.e., the chassis are linked and function as a single router, not as individual entities), therefore eliminating the consumption of line cards for inter-chassis connections. Having this capability allows service providers to scale their routers well beyond the scalability limitations of traditional routers. A scalable architecture eliminates the need to mesh discrete routers and enables service providers to collapse or combine network layers, creating a more simplified network.

- **Management as a Single Unit.** To simplify the network and the management of the network, the system needs the capability to manage and control each multi-chassis router. With these designs, service providers do not increase management complexity when they add more chassis to the system.
- **Wide Range of Optical Interface Speeds.** Supporting a wide range of optical interface speeds eliminates the need for service providers to deploy multiple boxes to support an array of line rates, with each box supporting one or just a couple optical speeds. This solution flattens and simplifies the network by also eliminating unnecessary speed-based layers within the core. It makes sense to build function-specific tiers, but it doesn't always make sense to add layers just to work around current box speed limitations. Equipment that supports a wide range of optical speeds reduces capital expenses, saves space and power, and eases management and operations. A distributed architecture with the ability to expand the switching capacity on an as-needed basis can support a wide range of interface speeds because the platform is not constrained by a fixed-capacity switch fabric.
- **Automatic and Transparent Failover of Control Cards.** In addition to having a backup control card to improve reliability, this backup card should operate in hot-standby mode, meaning the backup control card is aware of the current routes and adjacencies and does not trigger a network reconvergence if the primary card fails. Primary control card failures become transparent and do not disrupt packet forwarding functions.
- **In-Service Scaling and Upgrading.** Multi-chassis routers that have distributed architectures can allow service providers to scale their networks with little to no downtime. Adding a new chassis (populated with line cards) to an existing multi-chassis router becomes just like adding a line card to a chassis, providing near-instantaneous scaling of networks as traffic loads increase. Similarly, if a platform supports flexible expansion of switch-fabric capacity, adding more fabric to the router is transparent and non-service disrupting. Hot-swappable modules and in-service software upgrades are also important. In-service upgrades can be achieved with distributed architectures that feature hot-standby control plane implementations. Minimizing downtime improves network availability and allows service providers to more easily meet their SLA guarantees.
- **Independent Routing Protocols.** Maintaining independence between the different routing protocols by having dedicated control processors for each protocol, for example, can improve the network's reliability. A failure or problems in one of the routing protocols does not impact other protocols. Coupling segregated control modules for each routing protocol with hot-standby control modules provides the highest level of availability for protocols.
- **Simple Port Speed Upgrades.** As traffic volumes increase, the platforms' ability to quickly and easily upgrade port connections plays an increasingly important role in core networks. Through port aggregation, service providers can combine multiple interfaces of different speeds into one logical port, simplifying interface configuration as the logical port uses only one IP address. It also allows service providers to load-balance traffic between the interfaces that comprise the logical port. This design improves reliability as the connection continues to carry traffic over the remaining active interface(s) when other interface(s) in the logical port grouping fail.

- **Line-Rate Packet Processing with Enhanced Features.** Ideally, the platform should have the capability to process packets at line rate, regardless of packet size and features enabled. As core networks become intelligent, support a wider variety of traffic and applications, and provide more enhanced features and functions, line-rate packet processing while features are enabled becomes more important.
- **Tunable Quality of Service (QoS).** Core platforms need to support tunable QoS so service providers can truly offer intelligent services that provide the appropriate level of service and QoS based on application, traffic type, source and destination IP address, and other parameters. Tunable QoS gives service providers the ability to control and tune different QoS parameters such as delay, jitter, restoration, packet discard, and queue priority. Tunable QoS provides a high degree of service-level customization that the service provider can offer, beyond standard gold, silver, and bronze service levels. Depending on the sophistication of the service provider's customers, the service provider may allow specific customers to specify and configure the QoS parameters.
- **Intelligent Provisioning and Planning.** Ideally, core platforms should offer some level of intelligent provisioning to ease network planning and operations functions. For example, the platform might monitor load levels on different links in the network and report back when a link exceeds a specific load threshold. Service providers can then react to changing networking requirements in near real time.
- **Integration with Front- and Back-Office Systems.** Integration with front- and back-office systems becomes critically important as networks expand, become more intelligent, and increase in complexity. CORBA interfaces, Application Programming Interfaces (APIs), and partnerships with third-party Operations Support Systems/Business Support Systems (OSS/BSS) vendors better position router vendors to more seamlessly connect with OSS and BSS. These features will also allow the vendors to offer self-ordering and self-provisioning capabilities more effectively.
- **Professional Services.** Service providers should evaluate the vendors' professional services programs designed to help service providers with engineering, planning, operations, and marketing. These services offer network operators additional resources and expert partners to help them design, implement, maintain, and manage the network, as well as develop and launch services quickly. The programs also help providers generate revenues and realize a return on their investments more rapidly.

Summary

Traditional core routing platforms leave service providers with sub-optimal solutions for supporting IP-based services. These solutions typically involve forklift upgrades, a mesh of discrete routers, duplicate routers, multi-homed configuration, addition of unnecessary network layers, reliance on other technologies for quality of service and restoration, significant investments in space and power, and time-consuming network management.

Fortunately, innovations in technology help service providers minimize or eliminate some of the most painful networking challenges, including scaling the network, improving performance, improving reliability, simplifying the network and its management, reducing expenses, generating new revenues, and integrating with front- and back-office systems.

A new breed of multi-chassis routers with fully distributed architectures is positioned to address these challenges. These routers can scale well beyond the limitations of traditional single-chassis routers and offer a wide range of high-speed optical interfaces. These purpose-built architectures also provide carrier-class reliability in hardware, software, and system redundancy. In addition, hot-standby capabilities enable in-service upgrades and scalability. Finally, these routers can provide line-rate throughput while supporting a variety of enhanced features such as quality of service, packet filtering, prioritization, and security.

With the market focusing on sustainable profitability, these new platforms better position service providers to achieve success, allowing them to reduce their capital and operational costs and to develop and launch services more rapidly.

TeleChoice

TeleChoice, the leading strategic catalyst™ for the telecom industry, is recognized worldwide as the expert in launching telecom innovations. TeleChoice focuses at the intersection of strategy development and technology for telecom service providers and the vendors who serve them. TeleChoice helps companies crystallize business and market strategy, ignite new market categories, distill the value of innovation, and accelerate adoption of new technologies. Strategic catalysts™ like TeleChoice use rapid development methodologies to accelerate project completion. TeleChoice clients can achieve in days or weeks the strategic milestones that usually take months or years.

Pluris

Pluris develops IP core routers with the scalability, reliability, and performance to anchor the next-generation Internet. The TeraPlex family of routers features the industry's only optical backplane and chassis interconnect, supporting flexible multi-chassis deployments managed as a single router. Service providers can add capacity incrementally, with OC-3, OC-12, OC-48, OC-192, and Gigabit Ethernet interfaces, all providing line-rate throughput with extensive features enabled. For more information, contact the company at 408-863-9920 or visit the Pluris website at www.pluris.com.