Service Intelligence for the New Public IP Network

Introduction

The phenomenal success of the Internet and the universal adoption of the Internet Protocol (IP) are driving profound changes in the telecommunications industry. The infrastructure of today's Internet is being used as the foundation for a new Public IP Network, offering the universal reach and functionality needed to support a wide range of business-quality services, as well as today's "best effort" Internet applications. The revenue potential for network service providers is enormous if they successfully negotiate the transition from offering low-margin transport services to offering value-added IP services that will create new, high-margin, revenue streams. Innovative IP-based applications and expanding infrastructures promise to multiply demand for these services around the world.

To capitalize on this opportunity, service providers need more than just high-density edge devices, low-cost bandwidth, and fast routing switches in the core. Existing network devices, no matter how many bits per second they move, lack the intelligence required to switch users' sessions across a diverse array of applications and services. In the Public IP Network, for example, users will switch dynamically from service to service, accessing at any given time: the corporate intranet for telecommuting, a secure extranet for a business-to-business application, an IP telephony service for voice calls, or a multicast-based video conferencing service. Each session presents an unpredictable set of routing, addressing, security, performance, and protocol requirements.

The IP Service Switch (IPSS™) from Spring Tide Networks is a new class of product that delivers the network-based service intelligence required to enable the wide-spread deployment of IP Virtual Private Networks (VPNs) and other value-added IP services. The IPSS occupies the "service layer" in the Public IP Network, and provides the user-oriented, session-aware, processing and switching of traffic flows in support of a wide range of applications and services. Users benefit from network-based IP VPN services that are ubiquitous, affordable, and easy to use. Service providers are able to generate new, high margin, revenue streams that will continue to fund the explosive growth of the new Public IP Network infrastructure.

The Business-Quality Public IP Network

Despite its explosive growth and tremendous impact worldwide, today's Internet suffers from significant shortcomings; its one-size-fits-all, best effort service model, and its general lack of business-quality performance, security, and reliability negatively effect even the relatively few applications it does support (email, FTP, news groups, web browsing). Regardless, the immediate utility and commercial potential of the Internet has driven the networking world to adopt IP as a universal service interface. Equipment manufacturers are building a new generation of IP-based, carrier-class devices for data, voice, and video applications. Service providers are deploying these devices to offer new IP-based services in support of these applications. IDC projects that annual investment in public data network infrastructure will grow from $12 billion in 1999 to $22 billion in 2003.

This investment in IP network infrastructure is creating a new, business-quality, Public IP Network that will eventually encompass the Public Switched Telephone Network (PSTN), traditional leased line and shared facility data services such as Frame Relay, today's "best effort" Internet, and emerging broadband access services such as xDSL. This new Public IP Network is being constructed using IP-based equipment that operates over a high-bandwidth optical core, and provides the qualities we take for granted in today's PSTN:

- Services
- Connectivity
- Performance
- Reliability
- Security
- Simplicity
- Affordability
- Flexibility
- Scalability
- Ubiquity

The Public IP Network enables a new service model that supports:

- Class of Service (CoS), with multiple priority levels and delivery guarantees for different classes of network users.
- A rich set of IP services that provide the necessary Quality of Service (QoS) for a wide variety of data, voice, and video applications.
- The ability of a user to select a specific service for a given application, because the network can dynamically create the required connection based on security, addressing, protocol, CoS, and QoS requirements.
- Usage-based billing, enabling fine-grained control of service provisioning and billing that benefits both users and service providers.

The New Business-Quality Public IP Network

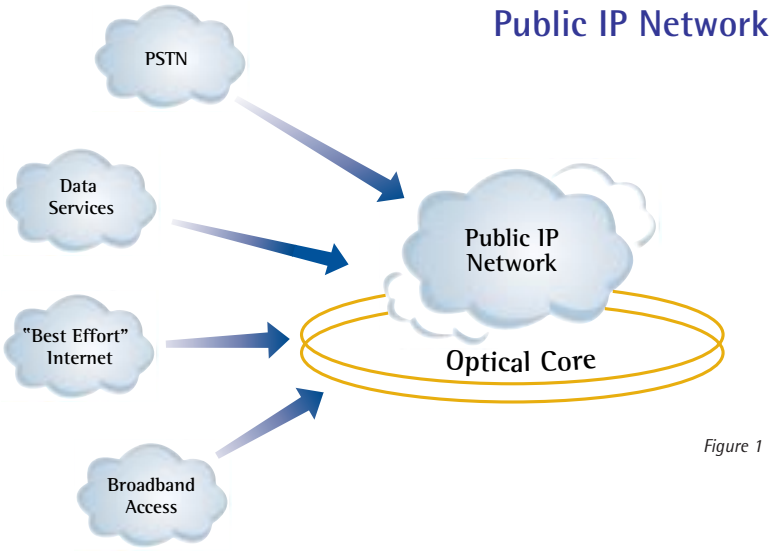

Figure 1

Enterprise IP VPN Applications

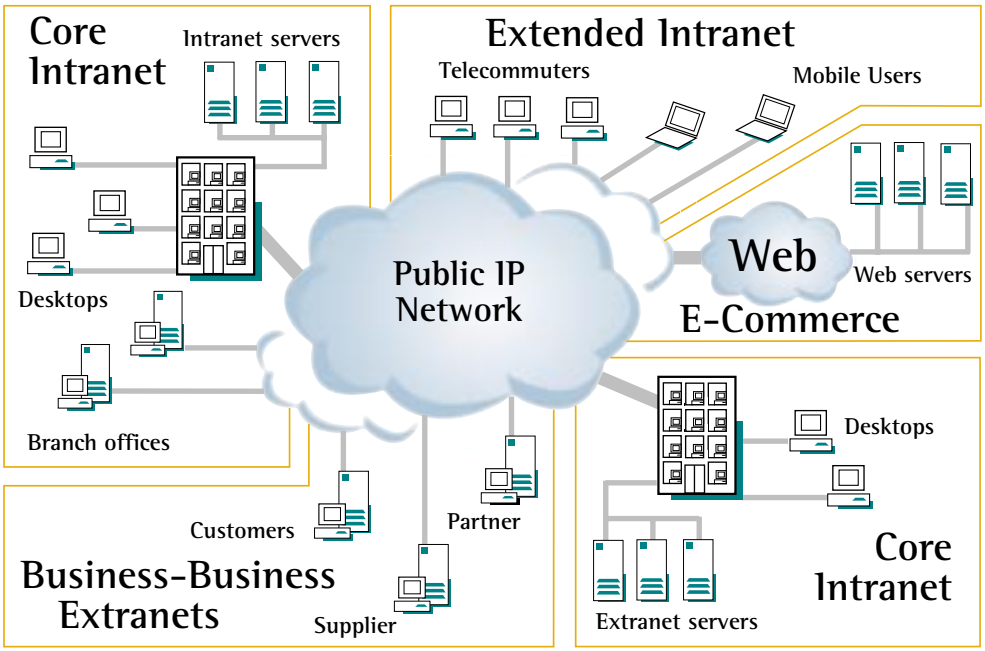

Figure 2

Enterprise IP VPN Applications

The business-quality Public IP Network enables a range of new enterprise IP VPN applications, including:

- Core intranet connectivity between campus sites and branch offices.
- Extended intranet connectivity for telecommuters, mobile workers, and Small Office/Home Office (SOHO) users.
- Dynamic network connectivity to customers, partners, and suppliers for secure business-to-business extranets.
- Web connectivity for a wide variety of e-commerce applications.

In contrast to today's enterprise networks, which often employ a variety of private and public network facilities for different applications, all these applications are supported over a common, IP-based, public network infrastructure. This benefits the enterprise by reducing the cost and complexity of managing the network, but also benefits the service provider by moving more traffic onto the shared facilities of the Public IP Network, generating additional revenue.

Projections for IP VPN Services

The revenue potential for IP VPNs is extraordinary. Table 1 lists revenue projections from several industry analysts. However, it is important to note that the bulk of these revenues will be derived from traditional enterprise networking customers, dominated by Fortune 500 companies. For instance, CIMI Corp. reports that today 80% of data service revenues are derived from only 28,000 sites. Although IP

Table 1

Projections for IP VPN Services

CIMI Corp.	$70.67B	Worldwide	2010
Frost & Sullivan	$13.59B	U.S.	2004
Infonetics	$8.85B	Worlwide	2001

80% of Data Service Revenues, Today, Are from Only 28,000 Sites

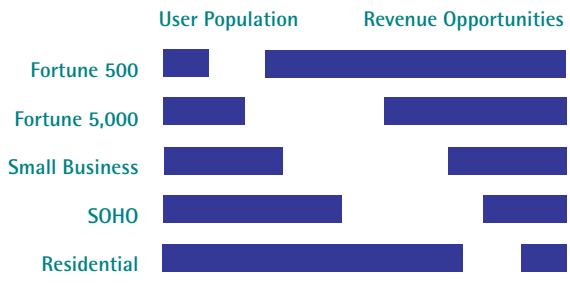

Figure 3

Dynamic, Diverse, and Unpredictable Traffic Flows

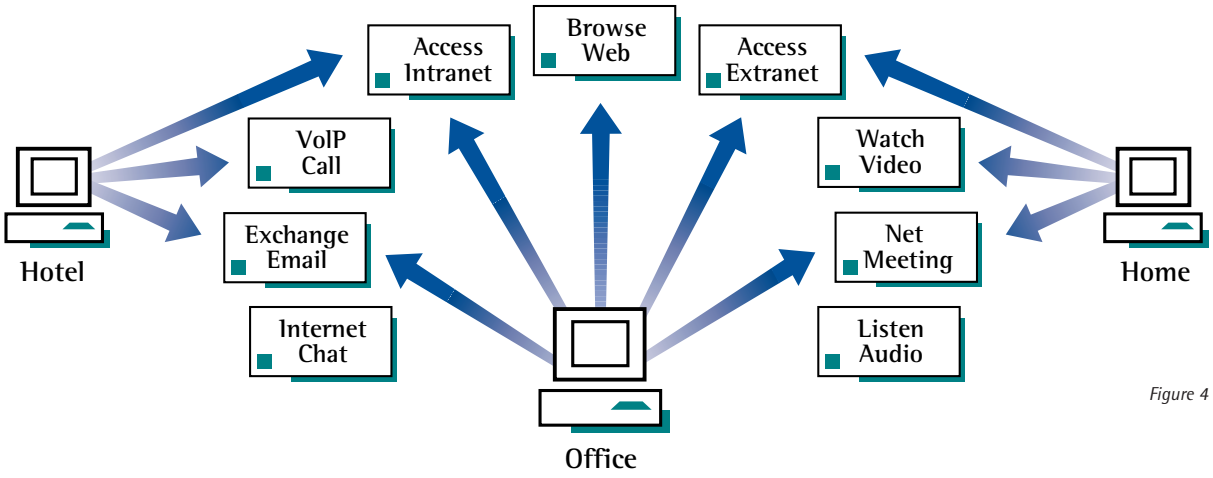

Figure 4

Any Service - Any Application - Any Place - Any Time

VPNs will extend the reach and scope of today's enterprise networks, enterprise customers will still account for a significant portion of service provider revenues, and more importantly, profits.

As IP VPN services are deployed to small business and residential customers, the number of users served grows, as do revenues. But since these customers have a fixed amount of money to spend on network services, these additional revenues will not generate nearly as much profit. There is a direct parallel to today's long distance calling services, where residential service is a low-margin, commodity market, and business services are the source of most profits.

Despite the weighting of IP VPN service revenues and profits toward enterprise users, the explosive growth of business and consumer e-commerce applications will continue to drive the deployment of IP VPN services to both small business and residential customers.

Dynamic, Diverse, and Unpredictable Traffic Flows

Since the initial commercialization of the Internet, service providers have struggled to keep pace with demand, and are constantly expanding connectivity and increasing capacity to accommodate new users and applications. At the same time, users have been treated as a homogenous source of undifferentiated traffic. This has simplified the network engineering challenges service providers have faced, but it has also worked against them, resulting in fierce competition for customers who pay low, flat-rate prices for unlimited Internet access, and who will readily switch service providers to get a better deal. Since the network has been engineered for undifferentiated, commodity bit transport, the service provider ends up in a low-margin, commodity business.

While the network engineering challenges are more complex, the new Public IP Network offers service providers the prospect of increased revenues and profits through the deployment of more highly differentiated IP services. In particular, the dynamic, diverse, and unpredictable nature of IP VPN traffic flows represents a significant network engineering challenge. At any given time, a given user can be accessing any service or application from any place in the network. Furthermore, each service or application has its own unique connectivity requirements, and users are entitled to different classes of service, depending on what they pay. Supporting IP VPNs, therefore, requires adding functionality to the network, which increases network engineering and operations expense. The return on this investment is satisfied, loyal customers, who are willing to spend more for value-added IP services.

Benefits of Enterprise IP VPNs

There are numerous benefits for customers deploying enterprise IP VPNs using the Public IP Network. First, multiple applications can be collapsed onto a single network, reducing cost and complexity. Second, leveraging a shared public network infrastructure results in further reductions in cost and complexity for all customers. Third, and most importantly, the new Public IP Network offers an unprecedented ability to support a diverse range of applications and services in a highly dynamic, flexible, and scalable fashion. To illustrate this, Table 2 summarizes the advantages of using an IP VPN vs. a private network solution.

Advantages of Enterprise IP VPNs vs. Private Networks

Private Network	IP VPN
Proprietary, closed	Standards-based, open
Layer 1 and 2 transport	Layer 3 transport (IP UNI)
Static topology	On-demand connections
Fixed bandwidth	Dynamic bandwidth
Complex, multi-device CPE	Simple, integrated CPE
Predominantly data	Data, voice, video
Uniform traffic flows	Diverse traffic flows
Site-oriented	User-oriented, session-aware
Strategic deployment	Tactical deployment
Slow to scale	Immediate scaling
Fixed reach	Global reach
Flat rate pricing	Usage-based pricing

Table 2

Service Intelligence Required in the Public IP Network

To fully realize the potential of IP VPNs, the new Public IP Network needs the service intelligence to ask each user: Who are you? Where do you want to go? Are you allowed to go there? What class of service do you require? And how much are you willing to pay for that service? It must then dynamically apply the necessary combination of security, performance, address management, and protocol functions to complete the connection.

This user-oriented, session-aware service model requires that the network support, at a minimum, the following service intelligence functions:

- User authentication and authorization – for secure access control
- Dynamic VPN selection – session level switching of traffic flows
- Encryption – to ensure privacy
- Compression – to conserve bandwidth, especially over access networks
- CoS and QoS – for tiered services and integration of data, voice, video
- Address management – network address translation, DHCP relay, etc.
- Accounting – fine-grained, for usage-based billing
- Customer network management – provisioning and management

But given that the network needs to support all these functions, we need to answer the questions: How are they implemented? Where do they reside?

The IP Service Layer

The new Public IP Network is comprised of two distinct parts: the access network, which aggregates user connections, and the backbone network, which transports these connections to their destination services. Neither of these, however, is a practical location for the network-based service intelligence required.

The access network is cost driven, and economic considerations are constantly driving down price per port for service provider edge devices. To remain cost-competitive, these devices

must specialize in concentrating and aggregating layer one and layer two connections, without being burdened with the additional processing power required to support a wide range of service intelligence functions.

The backbone network is performance driven, and is constructed using hardware-based core switches designed to transport millions of packets per second. The additional processing required to support service intelligence functions would slow them down.

Beyond cost and performance considerations, service providers need a common control point for provisioning, managing, and billing services to many thousands of users distributed across multiple access networks. To maximize service availability, service intelligence should not be tied to any one port, edge device, or physical link in the network, and it must be deployable in redundant configurations. Moreover, to handle mobile users, service intelligence can't be tied to any one port because a user must be able to connect to the network from anywhere.

The solution is to create an IP service layer between the access and backbone networks. At this strategic intersection, user traffic flows are already aggregated onto high-speed links. Service processing can be applied to many thousands of sessions without compromising price per port at the edge or the gigabit per second switching requirements of the core. A service layer between the access and backbone networks is also the ideal centralized location for large-scale provisioning, management, and billing of services.

Service Intelligence Powers the New Public IP Network

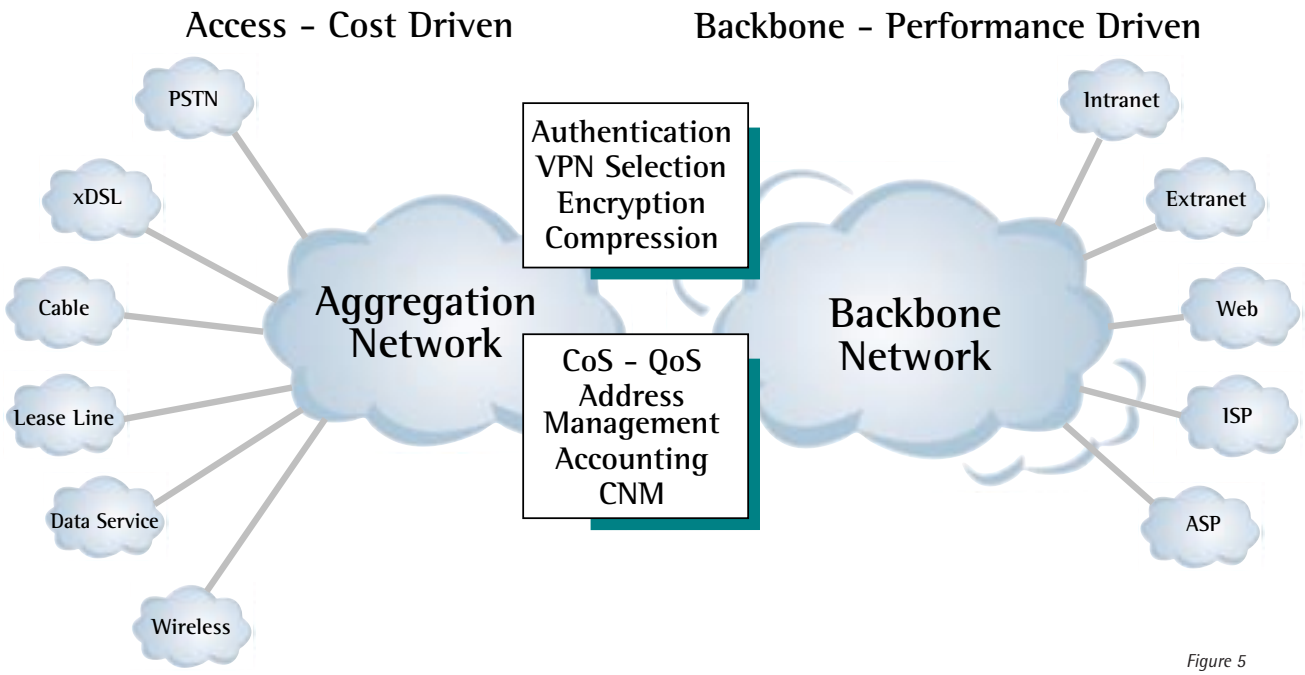

Figure 5

A New Class of Product: The IP Service Switch (IPSS™)

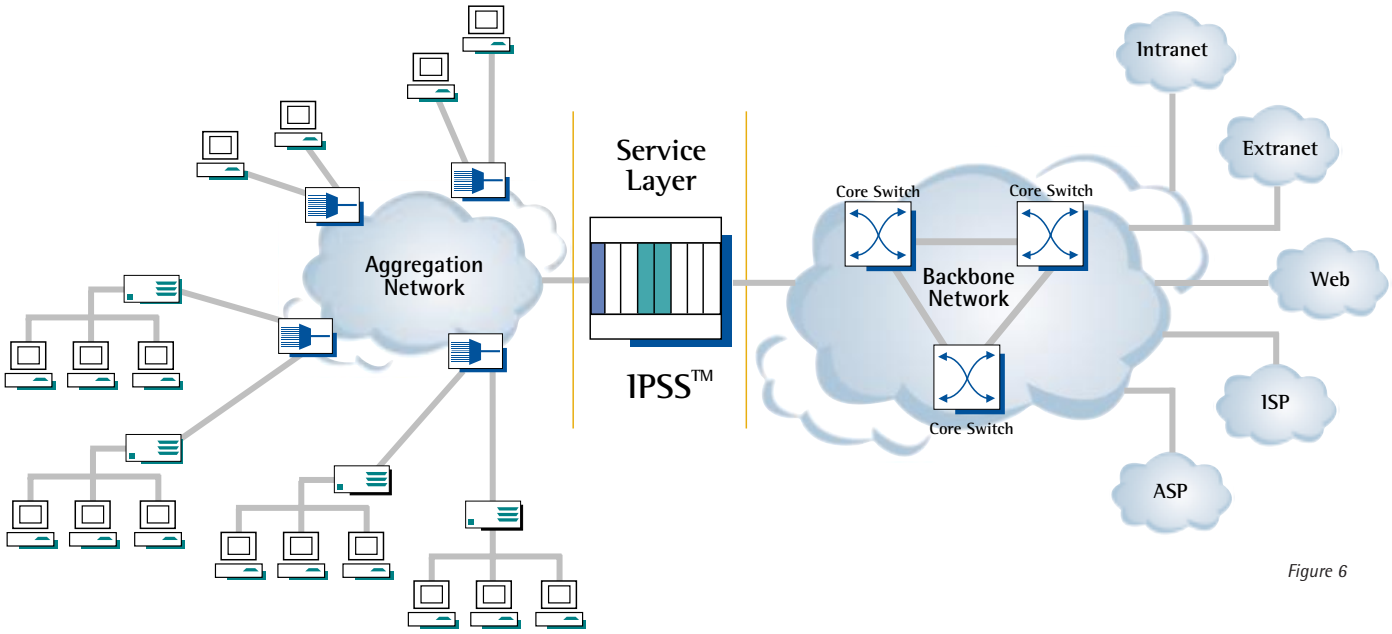

Figure 6

The IP Service Switch – First in a New Class of Product

The IP Service Switch from Spring Tide Networks is the first in a new class of product specifically designed for the IP service layer of the new Public IP Network. The IPSS integrates all user-oriented, session-aware, service processing functions in a single, highly scalable, carrier-class networking platform.

The IPSS is deployed in regional Points of Presence (POPs) where user traffic flows are aggregated from multiple local access networks. It supports a highly scalable, policy-based provisioning model, and fine-grained, usage-based billing of session-level connections.

The IPSS is a carrier-class platform that is NEBS-3 compliant, features optional redundancy of all common equipment, and hot-swap of all field-replaceable components. It's also based on a highly resilient software system to ensure maximum uptime.

Virtual Connection Termination and Switching

Figure 7 shows a conceptual model of how the IPSS terminates and switches virtual connections while individually applying the necessary service processing functions to each user traffic flow. Traffic enters and exits over high-speed interfaces such as OC-3 ATM, 100 Base-T Ethernet, and DS-3. Flows are conveyed in virtual connections that are aggregated over these interfaces. A virtual connection may be a virtual circuit, a PPP session, or a layer two or layer three tunnel of some type. A unique benefit of the IPSS is that it has been designed to handle a wide range of network types, and many different kinds of virtual connections across a variety of access and backbone networks.

Service Processing Functions in the IPSS

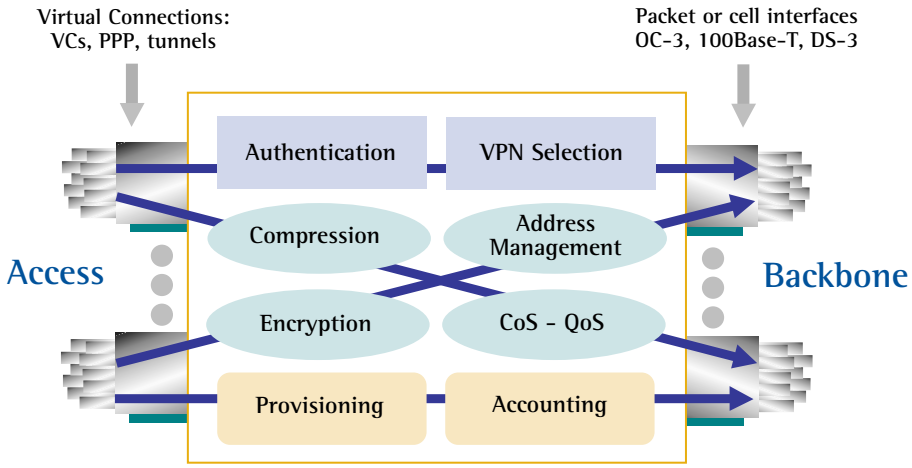

Service processing functions applied uniquely to each flow

Figure 7

The basic operation model of the IPSS is as follows:

1. The IPSS monitors virtual connections for new user traffic flows from the access network. When it detects a flow, it first dynamically constructs the input protocol stack required to process the flow (if necessary), and having done this, authenticates the user to access the network.

2. The IPSS then performs session-level filtering of the user's traffic flow to determine which VPN is being selected for a particular user session, and authorizes the user for access to the selected VPN.

3. Once a destination VPN has been selected, the IPSS then determines what route must be taken, and if necessary dynamically constructs the output protocol stack required for the virtual connection across the backbone network.

4. Having constructed the input and output protocol stacks for virtual connections across both the access and backbone networks, the IPSS applies the required service processing functions uniquely to each session-level traffic flow. These functions include: compression, encryption, dynamic address translation, and the classifying and queuing of flows to ensure the required CoS and QoS.

The IPSS is based on Spring Tide's breakthrough Pipelined PacketFlow™ architecture, which ensures that all flows are switched at full line speed with minimum latency. There is no performance penalty associated with putting the IPSS in the data path and having it perform multiple, processor-intensive service intelligence functions.

Policy-based Provisioning and Usage-based Billing

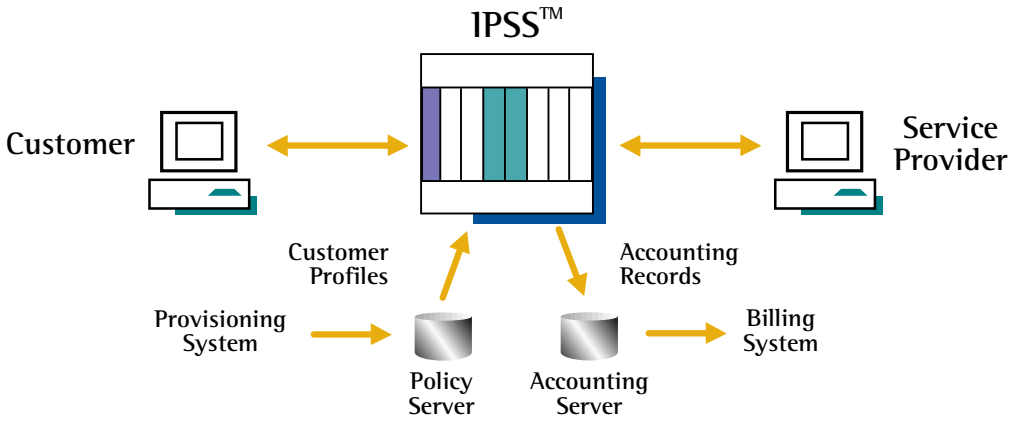

Provisioning and billing model scalable to millions of users

Figure 8

Policy-Based Provisioning

The IPSS enables a highly scalable, policy-based provisioning model because the dynamic processing of session-level flows is driven by customer profiles stored in an external policy server. Using standard protocols such as RADIUS or LDAP, the IPSS retrieves policy profiles that determine which service processing functions are uniquely applied to each user traffic flow. The dynamic behavior of the IPSS is controlled by customer profiles created in an external server using a service provisioning application.

This directory-enabled approach is the key to the scalability of the provisioning model. Device configuration of the IPSS is decoupled from service provisioning. Service providers can maintain hundreds of thousands or even millions of customer policy profiles in external servers without changing the configuration of any IPSS in the network. Furthermore, a common set of server-based policies can be shared by all policy-driven devices in the network.

Usage-Based Billing

Since the IPSS is the point where session-level connections are established between users and network services, it's also the natural place to perform the fine-grained metering of traffic flows required for usage-based billing. The IPSS can meter traffic flows using a wide range of billing criteria, and collects information on connections that is exported in detailed accounting records to an external server. This server then translates these records into the appropriate format, for instance, Bellcore AMA, required for the service provider's billing system.

Customer and Service Provider Management Views

Customer Network Management (CNM) is a key requirement in the new Public IP Network, and therefore the IPSS supports both customer and service provider management views. For example, customers need to be able to perform the following tasks without service provider intervention:

- Provisioning of user policy profiles
- Management of the customer's VPN
- Network diagnostics
- Service level agreement monitoring
- Online billing functions

The IPSS enables the service provider to provide secure access to these CNM functions, as well as control the degree to which customers can perform each management task.

Session-Level Tunnel Switching

Figure 9 illustrates a practical application of the IPSS session-level tunnel switching capabilities. In this example, the user connects to the IPSS over a single access tunnel, which provides the user with a "dial-tone" to access various services over the Public IP Network. The security and performance characteristics of an access tunnel are determined by the type of access network, and are independent of the services the user will select. Here let's assume the user is connecting via cable modem, in which case this could be a secure IP (IPSec) tunnel with compression enabled.

The user then runs various applications that require access to different network services. In this example, one application requires secure access to the user's corporate intranet (e.g., to download files), and the other a connection to the user's ISP (e.g., to check personal email). The IPSS has the ability to filter user traffic over the access tunnel and dynamically switch session-level flows to the destination network service over the type of tunnel determined by the user's policy profile. Here the tunnel to the corporate intranet could be an IPSec tunnel mapped to "premium" class of service, which the user's company is willing to pay for. The other tunnel to the user's ISP could be a non-secure IP over IP tunnel mapped to "best effort" class of service.

This examples demonstrates the basic operational model of the IPSS, which is to support session-level tunnel switching of authenticated user traffic flows to authorized network services while uniquely applying the service processing functions determined by the user's policy profile to each flow.

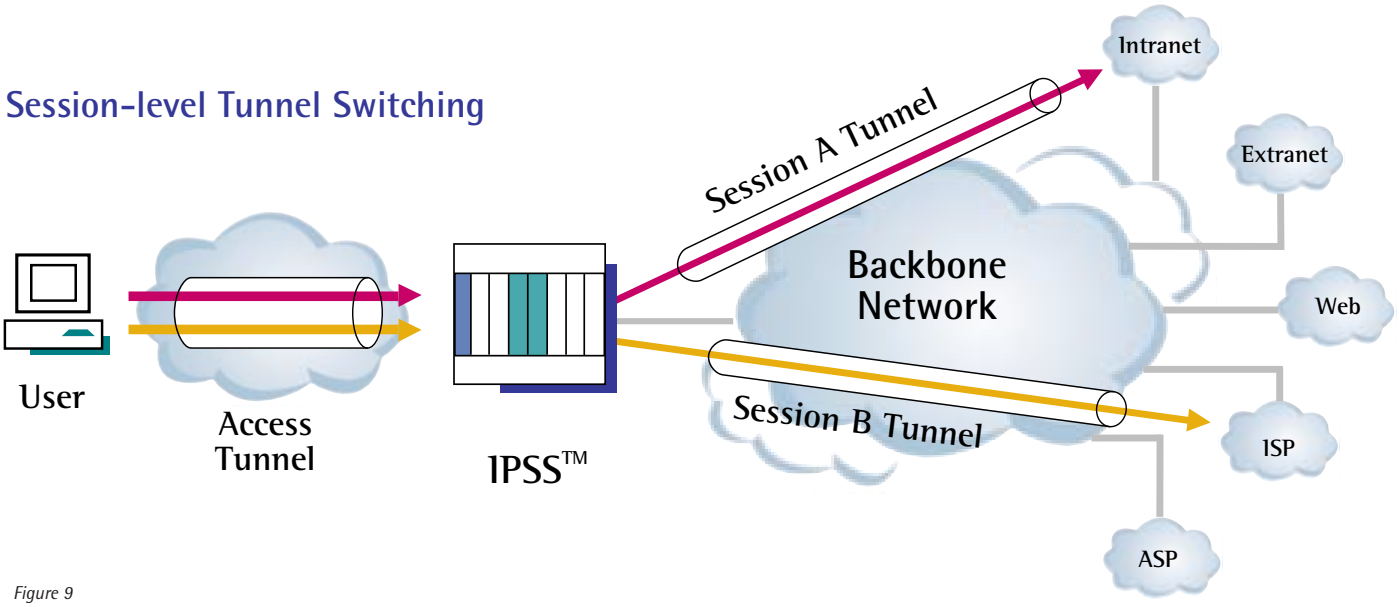

Figure 9

The IPSS Can Concentrate 1000s of Tunnels

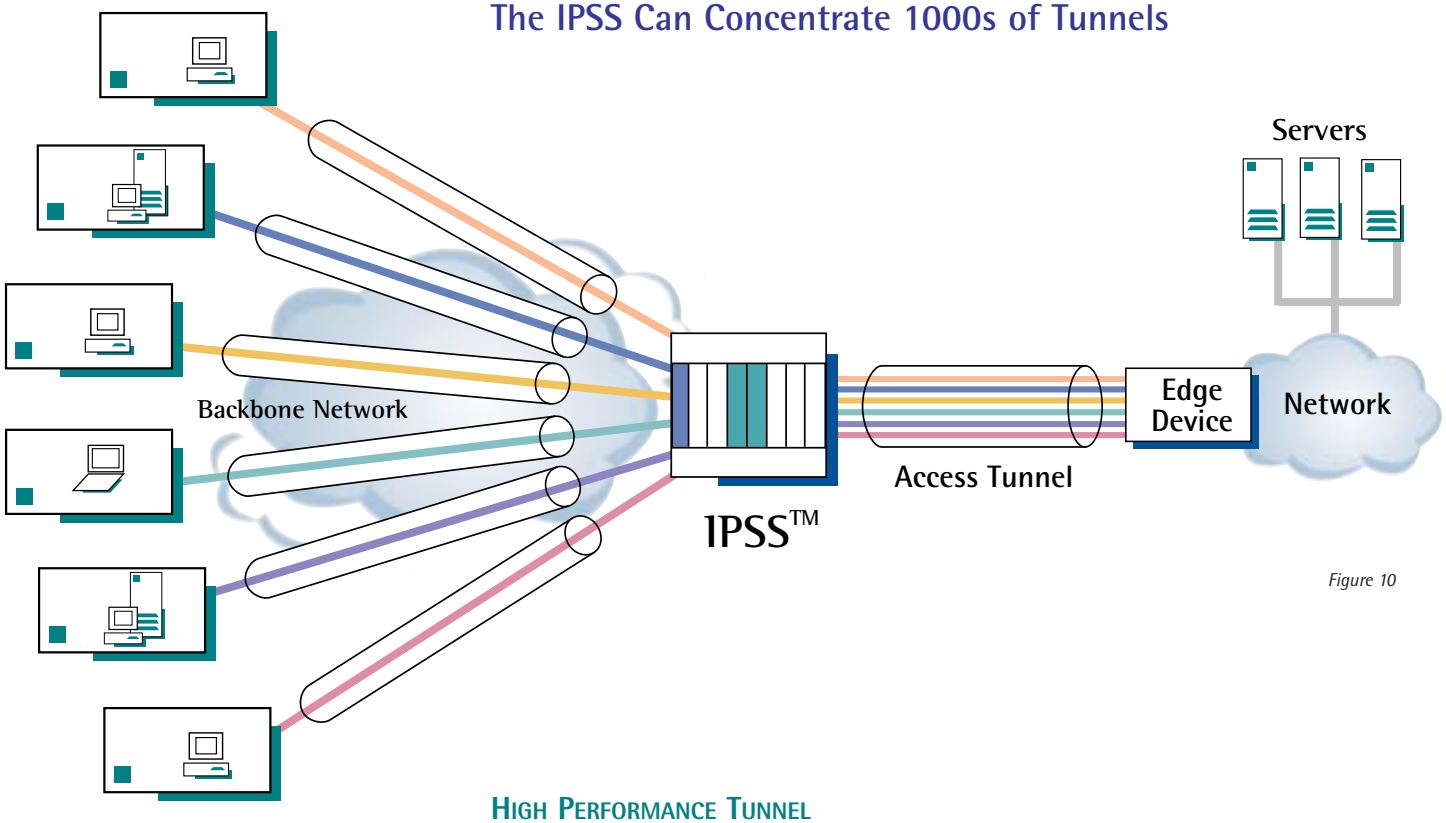

Figure 10

High Performance Tunnel Concentration

Figure 10 illustrates the high performance tunnel concentration capabilities of the IPSS, which can aggregate hundreds or thousands of flows, each conveyed in a separate tunnel, and switch these flows over a single access tunnel to a common destination network. The worst case scenario is when each flow is conveyed in an IPSec tunnel, since IPSec processing can be 25-50 times more computationally intensive than IP routing. If all the IPSec tunnels terminated at a traditional edge device, the processing load would be crippling, and performance would degrade badly. By terminating the tunnels at the IPSS, the processing burden on the edge device is greatly reduced, plus the administrative burden of managing hundreds or thousands of secure tunnels becomes the sole responsibility of the service provider.

Applications where IPSS tunnel concentration adds value include:

- Intranet access for remote users
- Business-to-business extranets
- Application hosting services

Each of these applications requires the large-scale termination of high performance, secure tunnels from hundreds or thousands of users accessing a common network service.

End-to-End vs. Edge-to-Edge Tunneling

Using the IPSS to deliver network-based IP VPN services enables an edge-to-edge tunneling model, as opposed to the end-to-end tunneling model that purely CPE-based IP VPN solutions employ. The benefit of this is that the edge-to-edge tunneling model is significantly less complex to deploy and manage.

In the end-to-end tunneling model, each CPE device must have complete information on how to establish and maintain tunnels to all other CPE devices. As the number of CPE devices grows, the management complexity of this model grows on the order of N-squared, where N is the number of CPE devices.

In the edge-to-edge tunneling model, each CPE device only needs enough information to establish and maintain a single tunnel to an IPSS in the service provider network. The network maintains all the information needed to establish connections to any other CPE devices. As the number of CPE devices grows, the complexity of this model remains order of N, where N is the number of CPE devices.

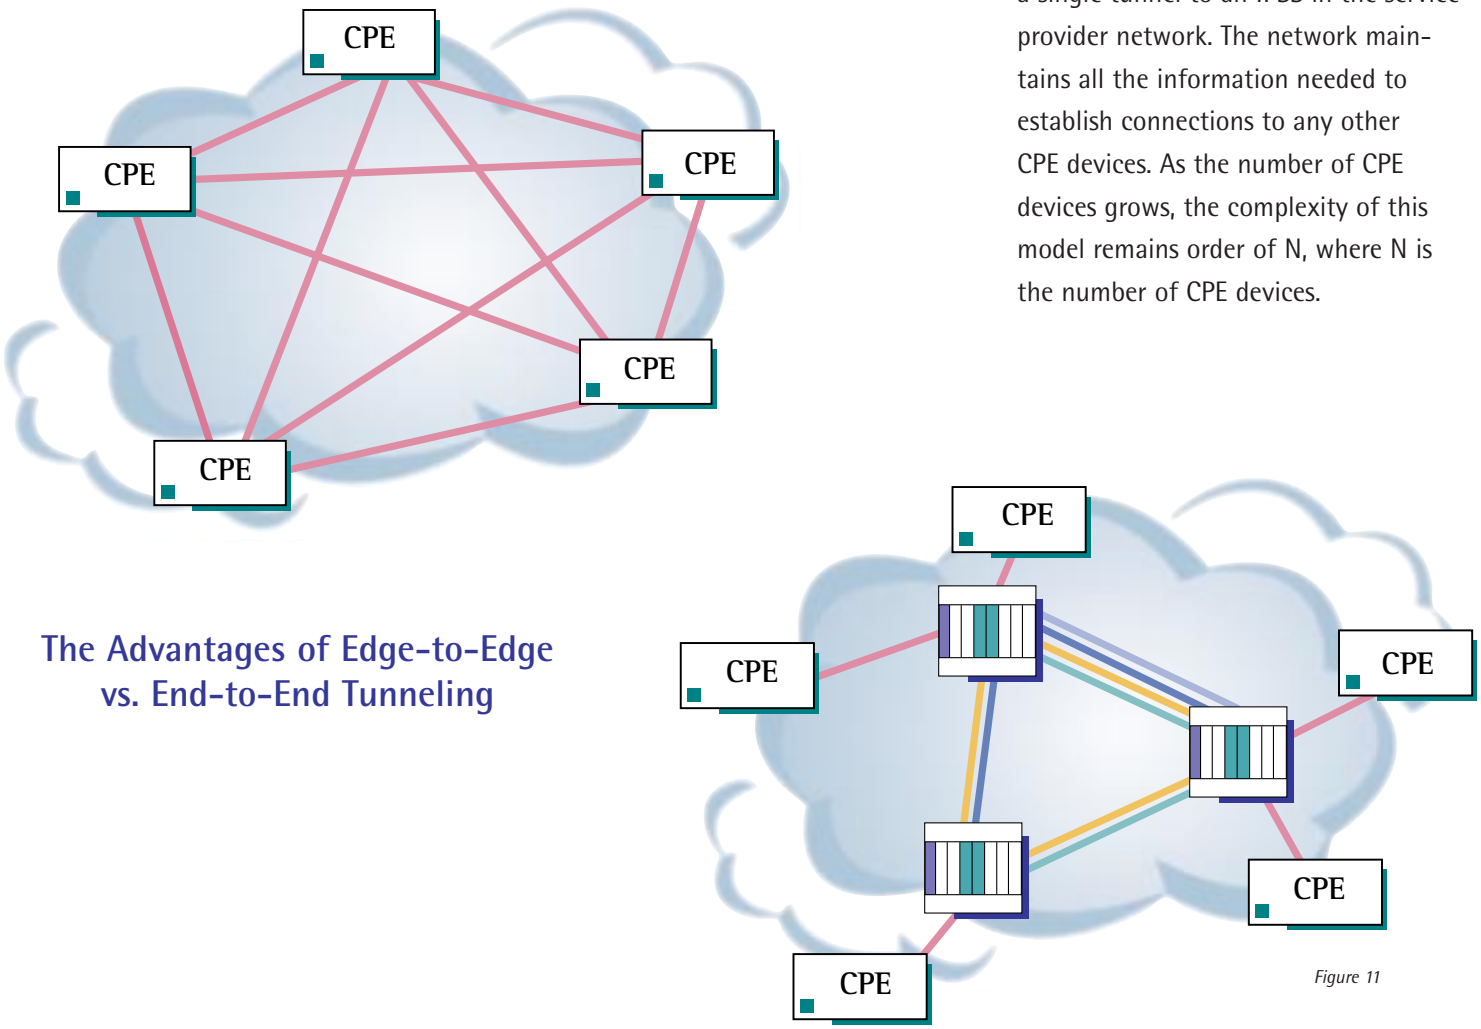

Figure 11

The Advantages of Edge-to-Edge vs. End-to-End Tunneling

Pipelined PacketFlow™ Architecture

The IPSS is based on Spring Tide's breakthrough Pipelined PacketFlow architecture, which combines the capabilities of three devices in a new type of architectural hybrid, incorporating: the user-oriented, session-aware capabilities of a remote access server, the IP intelligence of an IP router, and the performance and QoS of an ATM switch.

The IPSS sets a new benchmark for IP VPN equipment performance. It is capable of terminating and switching over 100,000 simultaneous virtual connections. The baseline for IPSec tunnel switching is full line-speed switching of eight OC-3s of IPSec tunnel traffic, with less than 50 microsecond latency.

The **Pipelined PacketFlow™** hardware architecture is highly modular and processor-intensive, featuring multiple general and special purpose processors. The IPSS is able to apply multiple service processing functions to all flows at full line-speed because it dedicates processing power and memory space for discrete functions, including: flow classification and queuing, routing and session management, packet forwarding, encryption and compression, and key generation.
At the core of the system, all processor and interface modules are interconnected by a multi-gigabit switch fabric. The IPSS is based on a highly versatile software architecture designed from the ground up to meet the unique requirements of user-oriented, session-aware processing and switching of traffic flows. It can terminate tens of thousands of virtual connections, and supports hundreds of fully partitioned virtual routers in a single platform. The IPSS supports policy-driven flow processing, and has an unprecedented ability to handle a wide variety of virtual connection types and protocol encapsulations. It does this by dynamically creating a unique protocol stack for each user traffic flow based on the policy rules from the user's policy profile.

Service Provider Benefits

The IPSS is a versatile, service-enabling platform for the widespread deployment of network-based IP VPNs and other value-added IP services. Occupying the service layer in the Public IP Network, it serves as a common control point for policy-based provisioning, usage-based billing, and customer network management.

The IPSS can be deployed in a variety of service provider environments. A complementary addition to existing access and backbone networks, it adds value while protecting the service provider's investment in installed equipment.

By integrating multiple service intelligence functions in a single device, the IPSS enables service providers to reduce equipment and operations costs, which are then amortized over many thousands of users.

But most importantly, the IPSS is the critical network element that will enable service providers to fully realize the potential of the new Public IP Network, and deploy new IP VPN services that will generate new, high margin, revenue streams.