THE POSSIBILITIES ARE INFINITE FUJITSU

Transport for Enterprise VoIP Services

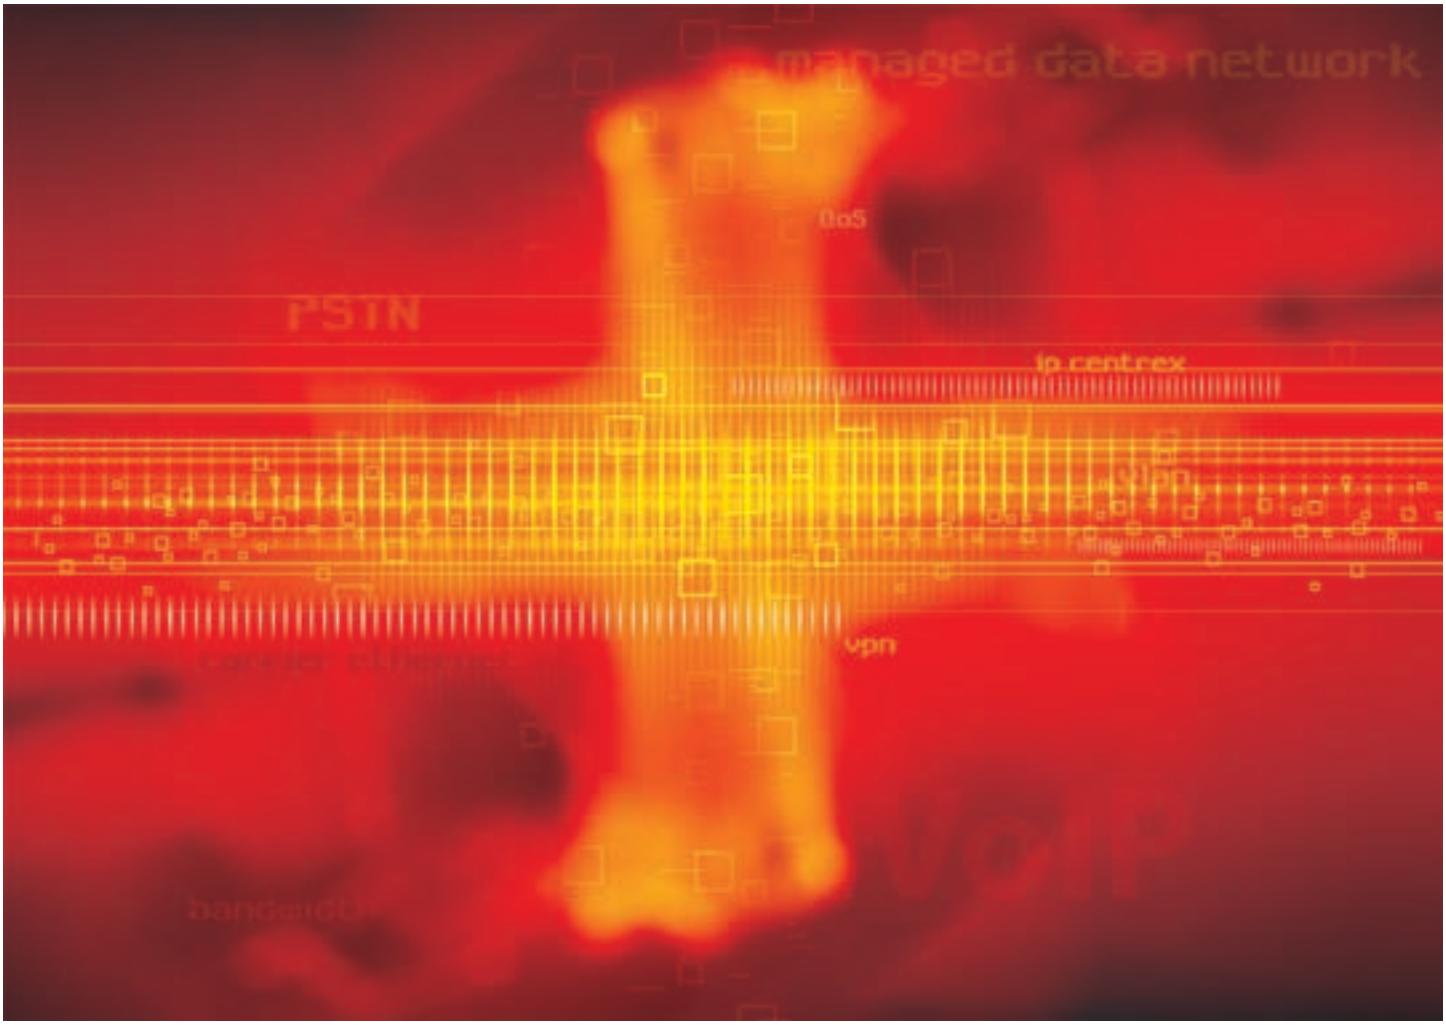

Introduction

Many carriers are looking to advanced packet services as an opportunity to generate new revenue or lower costs. These services, which include VoIP, IP Centrex and video services, typically require very tight QoS controls. Preparing for these new services requires understanding current and future network traffic demands, determining network policies for traffic and applying QoS controls over the WAN. Simply over-provisioning the network and hoping to obtain adequate QoS is insufficient. New technologies allow the network to efficiently provide the characteristics required by these demanding services.

The General Reference Architecture

There are a variety of architectures for implementing advanced services, many of which involve hanging various equipment types off of a managed data network.

A typical IP Centrex architecture is shown in Figure 1.

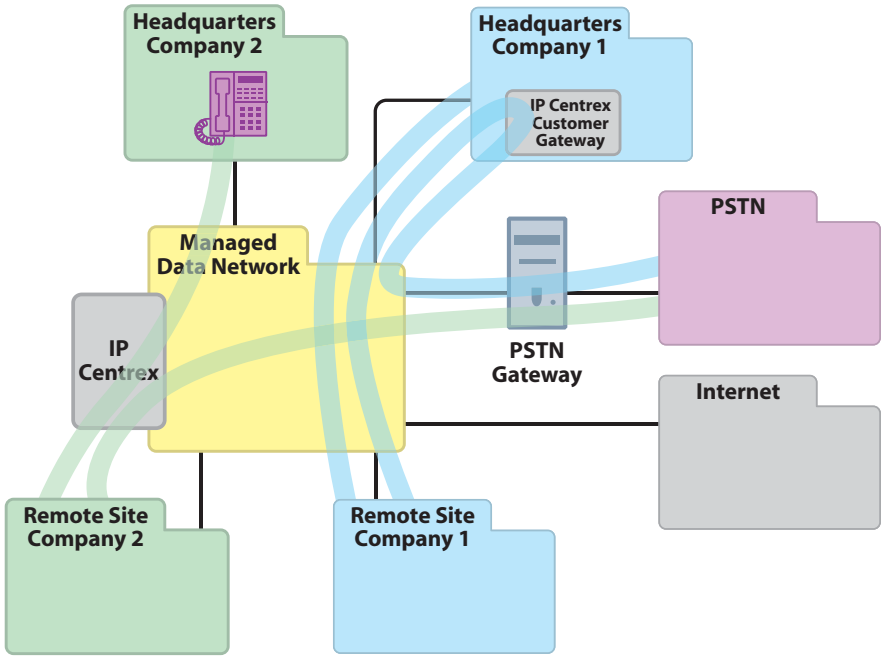

Figure 1: IP Centrex Service

As seen in Figure 1, a managed data network is used to provide both standard data services such as Internet access and Layer 2 VPNs, as well as advanced data services such as VoIP. With such a network, a service provider can provide data services, voice services and video services using the same network. Historically, service providers provide advanced services via dedicated overlay networks such as an overlay TDM network for voice services. However, a converged network is much more economical since there are gains from statistical multiplexing of the available bandwidth. Not only is this practice more economical, but it is also simpler in some ways, specifically since there is only one network connection at each customer site. Previous approaches, using overlays, would require multiple connections to each site.

With this architecture, an IP Centrex server and a PSTN gateway are added to the managed data network. There are two different ways to implement IP Centrex services. In Figure 1, Company 1 has chosen to operate its own IP Centrex server. Remote traffic for Company 1 is transported via the managed data network to the HQ location for service by the IP Centrex server. Internal traffic is serviced by the IP Centrex server and delivered to the destination phone. Should the call be destined for some external party, the call is transported via the managed data network to a PSTN gateway for hand-off to the PSTN. Company 2, on the other hand, has chosen to utilize an IP Centrex service provided by the service provider. In this scenario, the IP Centrex service is part of the managed data network. Calls are handled in the same way as the first scenario, but the location of the IP Centrex server is different.

This architecture allows a service provider to bring advanced next-generation telephony to customers who wish to limit their investment in legacy IP Centrex equipment. Solutions exist so that IP Centrex feature sets work with analog phones and include even more features when deployed with IP telephones. This approach can deliver a shared feature set across multiple business locations and home-based employees. Other advanced services, including video services, are provided in a similar manner where video sources and servers are included within the managed data network.

The Managed Data Network

Historically, service providers simply add bandwidth to the managed data network to provide more capacity for demanding services, hoping to improve performance. This may help in the short term, but typically bandwidth-intensive applications such as e-mail and Web browsing soon gobble up the added bandwidth. After uncontrolled applications consume an inordinate portion of network bandwidth, performance-sensitive applications such as VoIP and video begin to suffer. In essence, without proper network mechanisms in place, the least important (i.e. lowest priority) applications control the network.

The Network Requirements

Since adding bandwidth has been shown to be an ineffective technique to provide for advanced services, other approaches must be considered. The biggest challenge in providing for demanding voice and video systems in new, more economical packet-switched technologies is obtaining adequate QoS across the network. QoS is the capability built into the network to guarantee that information traverses the network within certain parameters.

According to data collected from www.TestYourVoIP.com, Brix Networks' free voice quality testing portal, nearly 20 percent of Internet telephone test calls experienced unacceptable call quality over the last 18 months. From late-2004 through mid-2006, their test results have shown a consistent decrease in overall voice quality as calculated via an MOS, a common objective measure of conversational voice quality that rates calls on a scale from one (bad) to five (excellent). Test calls with an MOS of 3.6 or better are typically regarded as having satisfactory quality. The number of test calls throughout this time that achieved an MOS of 3.6 or higher was only 81 percent.

For packet networks supporting voice, video and data applications, it's clear that adequate QoS is critical to preserve both mission-critical data in the presence of voice and video, as well as voice and video quality in the presence of bursty data traffic. Network quality of service is evaluated by measuring four key parameters:

- **Bandwidth:** The average number of bits per second that can travel successfully through the network.
- **End-to-End Delay:** The average time it takes for a packet to traverse the network from a sending device to a receiving device.
- **Jitter:** The variation in end-to-end delay of sequentially transmitted packets.
- **Packet Loss:** The percent of transmitted packets that never reach the intended destination.

In order to provide these objectives, a network must provide certain capabilities:

- **Connection Orientation:** In order to provide repeatable data delivery performance, a connection-oriented network is required. Connectionless networks do not provide the performance characteristics needed for advanced services.
- **Intelligent Traffic Mapping:** An interface from a customer will often deliver several different kinds of traffic to the network, and each of these traffic types correspond to different services. The network must be capable of accepting multiple traffic types from an interface and assigning them to different connections.
- **SLAs:** Data from differing applications require different performance characteristics. These performance characteristics are specified in an SLA. SLAs typically specify the jitter, latency, guaranteed bandwidth and protection mechanism.
- **Guaranteed Bandwidth:** Many services require guaranteed bandwidth and some services sell better if some amount of bandwidth is guaranteed. Consequently, the ability to provide guaranteed bandwidth is a necessary feature of a network designed for advanced services.
- **Best-Effort and Over-Subscription:** Providing connections for un-guaranteed or best-effort traffic is appropriate for some services such as e-mail and web browsing and allows the network to be over provisioned, thus more effectively using the network while still meeting varied performance characteristics.
- **Jitter Performance:** Some services are sensitive to jitter so control of jitter is paramount for those services.
- **Latency Performance:** Latency is crucial to some services and is not critical for other services. Consequently, being able to provide better latency for some connections is necessary.
- **Network Resiliency:** High availability, including sub-50 ms protection, is a necessary capability for networks that support advanced services.
- **Traffic Engineering:** Precise control of traffic in the network enables the control needed to effectively provide advanced services.
- **Traffic Policing:** The use of best effort traffic allows for oversubscription, but this traffic is not guaranteed and needs to be preempted if traffic for a service that has more stringent requirements needs the bandwidth. As such, the network needs mechanisms to police traffic and give priority where and when appropriate.

The Fujitsu Carrier Ethernet Solution

The Fujitsu Carrier Ethernet solution is a high-bandwidth MAN architecture that provides carriers with the scalability, flexibility, and efficiency required to satisfy growing bandwidth demands, decrease equipment and operational costs, and provide new revenue generating services. The solution, built on Carrier Ethernet platforms from our partner, Atrica, combines the latest in packet switching technologies, traffic engineering enhancements, and optical technologies to produce a new breed of low-cost, carrier-optimized platforms suited for delivery of differentiated services, including data, voice (including IP Centrex) and video.

Carrier Ethernet Connections

In order to minimize the management and service provisioning work and achieve the highest scalability, the Carrier Ethernet network aggregates traffic on a per-service basis using TLS. Point-to-Point connections can also be applied to services that are more appropriately served with a single connection.

Traffic Mapping

When establishing a connection (TLS or Point-to-Point) in a Carrier Ethernet network, operators need to choose a mapping scheme to assign traffic from customer interfaces to the connection. The schemes supported by the Carrier Ethernet nodes include:

- By physical port (i.e. port mapping)
- By the outer VLAN tag in Ethernet frames (i.e. VLAN mapping)
- By source/destination IP address
- By 802.1p bits in the outer VLAN tag (i.e. 1p mapping, which is usually used along with VLAN mapping)
- By DSCP bits in the Ethernet frame carrying IP packets (i.e. DSCP mapping, which can be used together with VLAN mapping or used independently)
- By EtherType field in the frame to map non-IP packet (e.g. ARP) to a connection

When voice services are implemented with VoIP, DSCP bits are used to indicate that voice traffic is contained within the packets. By using a combination of VLAN and DSCP mapping into connections, voice traffic can be differentiated.

Operators can use these mapping schemes to classify, segregate and aggregate different types of traffic from end points onto different connections based upon a flexible set of criteria such as the type of traffic or the group of subscribers. These mapping schemes allow maximum freedom for service providers in supporting advanced services. Additionally, multiple logical connections can be created in a single customer port. In other words, if a customer is being served with a single 10/100Base-T Ethernet connection, a customer can mix voice, video and data traffic on that link and each traffic type can have its own unique connection (where each of these connections can have different destinations and SLAs). Consequently, each traffic type from the customer can have an SLA appropriate for that service, even though they all share the same access port.

The Importance of QoS
The delivery of customized, tiered services depends on the ability to distinguish between traffic types and the capacity to regulate the sending and receiving conditions of individual data flows. QoS is required to support latency-sensitive applications or the wide range of today's networked business applications. Without it, providers cannot assure that all communications, regardless of the resources and bandwidth they consume, will be delivered according to every customer's satisfaction. A comprehensive system must be deployed, enabling network operators to configure services and establish customer-specific SLA attributes. In addition, traffic must be monitored on an ongoing basis to measure characteristics such as packet loss, throughput, delay and jitter to ensure SLAs are being met.

Establishing Connections
The Carrier Ethernet solution for QoS management is based on a connection-oriented approach in which end-to-end connections for each customer's traffic are established across the network. Definition of connections is accomplished using MPLS-based LSPs with the aid of the OSPF-TE protocol, which discovers network topology and updates the registry of network elements and resources. The RSVP-TE protocol is also used to reserve bandwidth for each customer's connection and signal that connection along the shortest or explicit path.

These mechanisms greatly streamline and speed-up the provisioning process. The alternative methods involve the use of multiple management systems to configure and provision services, plus require multiple manual handoffs between different teams. By contrast, a single network operator using the Carrier Ethernet solution can provision services in seconds with a few clicks of a mouse button from a single workstation.

Separating traffic into individual paths or tunnels enables a specific customer's traffic and services to be segregated from those of other subscribers as they traverse the network core. This technique accomplishes two things: It provides more efficient and faster traffic delivery than alternative solutions. Plus, it enables bandwidth and other resources to be reserved, eliminating the possibility that they will be degraded by other services. This method differs greatly to those used in many of today's Ethernet and IP networks in which services can only be guaranteed on a best-effort basis.

Defining SLAs
Once traffic has been separated into individual paths, it becomes easy to assign and apply unique attributes for quality and protection for each one. These attributes form a contract, or SLA, defining the service level according to predefined or customized values.

Parameters corresponding to the following classifications are entered into the management interface:

- CIR corresponds to the guaranteed transmission rate. This bandwidth is always available to the customer from end-to-end.
- EIR sets the maximum best-effort traffic rate. The actual traffic rate depends on the available bandwidth on the selected end-to-end path and is thus not guaranteed (i.e. when congestion occurs, EIR traffic may be dropped). EIR corresponds to the best-effort service within the allowed rate.
- Traffic priority assigns priorities to different types of traffic that have different sensitivities to latency and jitter. Priority queues are used to manage delay and jitter for four types of user traffic: normal, business-critical, delay-sensitive, and TDM CES traffic. Priority for network control and management is set between that of business-critical and delay-sensitive traffic. Delay priority and jitter values can be assigned per-connection or per-packet. The Carrier Ethernet solution utilizes a WRR scheduler over the queues for normal, business-critical and network control traffic so that each of these types of traffic may have fair opportunity to be sent. A strict priority queuing mechanism is used with the queues of delay-sensitive and TDM CES traffic to guarantee the minimum delay and jitter for these two types of traffic.
- Protection level ensures traffic will reach its destination with little or no disruption in service in the event of a span cut or other failures such as a configuration error or partial network element failure. Network operators can select the restoration level that fits customer needs: sub-10 s protection implemented by the End-to-end protection mechanism, or the SONET/SDH-like sub 50 ms restoration using fast link protection.

Traffic Policing and Congestion Management

As described above, the network is managed in relation to each customer's SLA. As live traffic traverses the network, the Carrier Ethernet system checks SLA values (CIR and EIR) using patented traffic policing mechanisms for each connection. During this process, traffic is allowed or marked for DE, or dropped depending on established parameters. If traffic rate is below the CIR, all the packets will always be forwarded and will not be impacted during periods of peak network congestion. If the CIR is assigned to zero, the customer's traffic will always be in excess and therefore be handled on a best-effort basis.

If the customer's traffic is above the CIR but below EIR, the excess traffic will be marked for DE, signifying that it will be dropped during times of congestion. In many cases, EIR traffic will be allowed to continue as long as there is enough bandwidth available on the path of the connection. Otherwise, the switches in the Carrier Ethernet solution will only drop those packets with DE marked so that CIR is always maintained. If the amount of traffic exceeds the EIR, then the traffic that is beyond EIR will not be permitted to go through the ingress interface of the connection.

All policing is accomplished within the Carrier Ethernet platforms, which support more than a thousand traffic policies per GigE interface. This results in much more efficient traffic management than in other vendors' networks and more scalability for a large number of QoS-guaranteed traffic flows.

Network Resiliency
The Carrier Ethernet solution provides not only an affordable high-bandwidth metropolitan transport network, but also a carrier-class infrastructure that includes reliable sub-50 ms SONET/SDH-like resiliency. Sub-50 ms protection is required to support IP Centrex voice applications.

The Carrier Ethernet solution utilizes a standard MPLS and VLAN tagging transport infrastructure, which are used to create backup paths and tunnels in order to guarantee data delivery in the event of a main path failure. This ensures a predictable convergence time.

When provisioning point-to-point or VPLS services, an SLA is configured which includes not only bandwidth and delay characteristics, but protection level as well. Each time a connection is provisioned via the simple point-and-click interface of the Atrica ASPEN EMS, a selection between different types of protection modes is specified. The Carrier Ethernet solution can provide not only a redundancy level comparable to SONET/SDH-based solutions, but also support multiple grades of restoration to be provided as a part of the connection SLA parameters.

During the connection provisioning process, a back-up path is established. In order to ensure rapid switchover to the backup path without service degradation in the event of failure, bandwidth is reserved for the back-up path as well. One of the key advantages of Carrier Ethernet-based protection over other technologies is its ability to utilize bandwidth for EIR traffic during stable network conditions.

The protection scheme of the Carrier Ethernet solution, which provides different restoration times ranging from 10 seconds to sub-50 ms, uses two different protection mechanisms: End-to-end connection Protection (compliant to EEPP as specified by the MEF) and fast facility protection (based on MPLS fast reroute).

Summary

The Fujitsu Carrier Ethernet-based network can support advanced data services, including VoIP and triple-play services for residential and business subscribers, while it also opens the opportunity of supporting other types of services (i.e. IP Centrex) over the same network. It is a connection-oriented Layer 2 solution that has the following advantages and value propositions to service providers:

- Hard QoS for multiple service types
- Fast protection against network and service failures
- Scalability in supporting a large number of subscribers
- Flexibility that gives service providers the freedom to provide innovative services
- Simplicity that eases the work of configuration and service provisioning

Acronym	Descriptor
ASPEN	Atrica Service Platform for Ethernet Networks
CES	Circuit Emulation Service
CIR	Committed Information Rate
DE	Discard Eligibility
DSCP	DiffServ Code Point
EEPP	End-to-End Path Protection
EIR	Excess Information Rate
EMS	Element Management System
GigE	Gigabit Ethernet
HQ	Headquarters
IP	Internet Protocol
LSP	Label Switched Path
MAN	Metropolitan Area Network
MEF	Metro Ethernet Forum
MOS	Mean Opinion Score
MPLS	Multi-Protocol Label Switching
OSPF-TE	Open Shortest Path First with Traffic Engineering
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RSVP-TE	Resource Reservation Protocol with Traffic Engineering
SLA	Service Level Agreement
TDM	Time Division Multiplexing
TLS	Transparent LAN Service
VLAN	Virtual Local Area Network
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
WRR	Weighted Round Robin

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom