WHITE PAPER

SERVICE PROVIDER VOICE OVER IP (VOIP)

Summary

The trend of offering voice services over a packet-based network has provided many opportunities to new and existing telecommunications providers. This market is expected to grow as providers offer additional multimedia services over the same network. This makes it imperative that the underlying network is scalable and designed with diligence. Voice services demand a highly reliable network that can maintain voice quality. This is achieved through networking products that offer high reliability, low latency, and can prioritize traffic based upon packet priorities. Routers and switches should also be able to support traffic engineering capabilities that can guarantee consistent and predictable behavior through the network. Foundry Networks routers and switches enable a wide variety of solutions that make it easier for service providers to deliver these advanced multimedia services.

Contents

Figures

1 Introduction

The ability to packetize and carry voice over an Internet Protocol (IP) network is called Voice over IP (VoIP). The technology provides IP Telephony service that can be combined with data and various multimedia services over a converged IP network. While IP Telephony can be applied within enterprise networks as well, this paper concentrates on architectures and requirements for service providers offering services to residential and small-medium business (SMB) customers.

Voice has been traditionally carried on circuit switched networks with resources dedicated for each call. These networks are highly reliable and have set the standards for voice quality. Data networks have co-existed with circuit switched networks, and in the last decade their growth has outpaced the growth in voice networks. At the same time there have been vast improvements in the data networking standards and in the switching/router technology. Today data networks are capable of delivering toll-quality voice over a network that is much more scalable. Service providers can now provide Public Switched Telephone Network (PSTN) scale solutions with additional services like email, unified messaging, and cost effective business services with the flexibility to add more multimedia services in future.

2 Motivation

VoIP technology has many proven and potential benefits for established voice carriers currently offering PSTN services, as well as for Internet Service Providers (ISPs) and cable operators newly entering the voice market. The established telecommunication companies have seen considerable customer churn and declining voice revenues due to competition from cellular and Internet-based voice providers. They need the ability to create new “sticky” services to help reduce customer churn and increase Average Revenue per User (ARPU). IP Telephony provides additional revenue opportunities through rapid creation and delivery of new multimedia services. Further, a move towards converged services with voice, video, and data over a single IP-based network promises to reduce CapEx and OpEx.

The technological innovations in IP telephony and wide availability of cost effective IP infrastructure network have given rise to new business models. Many new service providers have emerged and some ISPs are expanding into enhanced voice services. Customers demand services similar to PSTN at a lower cost point, and additional value-added services including video, unified messaging and conferencing.

An example of emerging value-added services is Hosted VoIP, which provides voice service with full PBX like functionality with no on-site hardware requirements. In-Stat market research firm projects strong growth for Hosted VoIP over the next few years, driven mainly by cost savings as well as additional benefits for companies with a distributed workforce. The research firm has also projected strong revenue and subscriber growth for residential and SMB VoIP service offerings. It forecasts residential IP telephony revenue to be over $3.5 billion and Hosted VoIP to be over $2B by 2010[1]. The growth in subscribers is coming from the Telcos, cable Multiple System Operators (MSOs), and new service providers.

[1] In-Stat research: In-Dustry Update: Hosted VoIP: Steady Growth, But Will the Boom Come?
In-Stat research: Consumer VoIP Subscribers Skyrocket, New Voice Contenders Entering Market

3 Key Considerations for Designing Packet Voice Infrastructure

3.1 Services offered

VoIP has enabled many new services, the most basic being Internet Telephony service to users over an IP network. This may involve PC-to-PC connection or PC-to-phone connectivity using PSTN gateways. This could either be a best-effort service over the Internet or be offered as a service by a service provider with required guarantees. The service has already evolved to video telephony and many service providers are offering video telephone service for users that own a video phone. Further, business could benefit from value-added services such as videoconferencing and application sharing. VoIP networks eventually will evolve to offer a variety of multimedia services over a converged IP network.

The services can be categorized as follows:

- Residential VoIP – Voice over broadband services sold to residential and home office customers.
- Hosted VoIP for businesses – Voice service with full PBX like functionality for SMBs.
- Long distance bypass – IP trunking service offered by carriers to enable long haul voice providers to bypass long distance toll networks.
- IP trunking services – Connect islands of PSTN networks using private IP networks.
- VoIP peering – Allows direct peering of VoIP networks to completely bypass PSTN networks wherever possible, providing relief from PSTN regulations and tariffs.

3.2 Quality of Service

Voice quality has been one of the major considerations in deploying these networks. Noise, voice delays, and echo interfere with regular voice conversations and must be addressed to deliver a toll quality voice service. To achieve the high Quality of Service (QoS) requirements for toll quality voice, various techniques have been deployed in the network. Echo is a serious issue in voice networks and can be due to impedance mismatches in a classical circuit switched network or simple acoustic coupling between microphone and speaker of a phone, which can occur in either classical or VoIP networks. Echo becomes more pronounced as round trip delay goes above 50ms. In a VoIP network, every device and link introduces delay causing mouth-to-ear delay to be the sum of encoding delay, packet serialization delay, propagation delay, network equipment delay, and play-out buffer delay. Since such mouth-to-ear delay can be as high as 100-150ms, echo cancellers are often used. Today echo cancellers have become very effective even in long delay networks.

Voice networks are sensitive to packet loss as well, and steps must to be taken to avoid end-to-end packet loss. In IP networks, packet loss is usually a result of congestion or the bursty nature of traffic. Hence voice packets are assigned a higher priority and the router queues deliver preferential treatment based upon the priority of packets. A few of the approaches used are described below:

1. Differentiated Services (DiffServ): Differentiated Services (RFC 2475) can be used to provide class based traffic management capabilities in the network without the need for per-flow state and signaling at every hop. The method is based on classification and marking of packets at the edge into a limited number of classes. The Per Hop Behavior (PHB) of each packet will be determined by the class of the packet. The RFC has reused the IP precedence bits as 8 bit DS field. Six bits of DS field are specified as Differentiated Services Code Point

(DSCP), which determines the PHB of a packet. Although there are 64 different DSCPs, most networks map these to the following standardized per hop behaviors (PHBs):

- Expedited Forward – Loss sensitive traffic requiring minimum delay and delay variation (jitter)
- Assured Forward – This has multiple classes and applies to loss sensitive traffic. Traffic forwarding is assured as long as the traffic is within the service specifications. In other words, each class is guaranteed a certain amount of bandwidth.
- Best Effort

2. Layer 2 CoS (IEEE 802.1p): The approach is to use 802.1p priority bits to prioritize traffic. The packets are remarked as they traverse to a Layer 3 network.
3. MPLS Traffic Engineering (TE): MPLS-TE provides precise control over the path packets traverse, which can be set to avoid network congestion points. RSVP-TE is used to set up a traffic engineered tunnel from ingress point to egress point by reserving bandwidth across the path. Once established, many individual flows can be aggregated on this tunnel without the need to allocate bandwidth for each flow. Further, priority bits from Layer 2 CoS or Layer 3 DiffServ can be copied to EXP bits in MPLS header, and the packets can be prioritized across an MPLS transport network.
   MPLS protocols are designed with resilience and recovery mechanisms such as Fast-Reroute and path protection to reroute traffic around failures in the network in as little as 50ms. Further, they work with the robust IP control plane and take advantage of the resiliency and fast convergence mechanisms of IP.

If the networks are not designed with sufficient bandwidth based on peak load requirements, packet loss might still occur but it should be kept to a minimum. A good number of VoIP gateways and IP phones also implement sophisticated concealment methods that can mask the effect of minor packet loss. Some of the simple concealment methods include rerunning the last correct packet. This makes the packet loss less perceptible, although it might add additional delay in the network. In VoIP, networks routers with hardware based QoS are highly desirable.

Jitter is the variation of end-to-end delay from one packet to the next packet. It is a major problem for voice networks. In order to remove the effects of jitter, many IP phones employ buffers to collect the packets and smooth out the variation before play out – hence, called play-out buffers. This process introduces some additional delay and potential for packet loss – if not tuned correctly. Network designers must carefully tune the appropriate buffer size. Too small a buffer will yield low delay but may result in an unacceptable packet discard rate. Too large a buffer will yield the opposite. To address this, and simplify design and operations, modern IP phones sometimes employ adaptive buffers that can dynamically adjust the buffer size to optimize the delay to discard trade-off.

3.3 Signaling Protocols

Signaling protocols are required for creating, modifying, and terminating sessions between endpoints in a VoIP network. The signaling protocols can be divided into two broad categories. Peer-to-peer call signaling protocols that include Session Initiation Protocol (SIP) and ITU H.323 were initially designed for two intelligent endpoints to communicate with each other. They allow users to find the remote device, and provide call setup, call teardown, capabilities exchange and call control functions to establish multimedia communication (voice, video, etc.).
SIP is defined in RFC 3261 as “An application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences.” ITU H.323 is an umbrella standard that came from the telephony world and defines protocols for call establishment and teardown in a packet based network. The standard incorporates specifications

for audio codecs, call control and signaling and data/fax. Although H.323 and SIP protocols have some differences, they offer very similar capabilities and are both widely deployed.

The second category is comprised of client-server protocols such as H.248 (MEdia GAteway COntrol/MEGACO) and Media Gateway Control Protocol (MGCP). This model assumes very little intelligence at the end terminal and provides the intelligence in core. It uses low cost phones, Media Gateway Controllers (MGCs) for call control and signaling, and Media Gateways (MGs) to interface to PSTN. MGCP and MEGACO provide a method of communication between MGC (also referred to as a call agent or softswitch), and the MG. Hence, MGCP and MEGACO protocols can be complementary to SIP or H.323.

3.4 High Availability

Telephone users are accustomed to extremely high availability on current circuit-switched telephone networks. The networks are also used for many mission-critical applications for businesses that demand networks to always be available. Such consistent service is expected of VoIP networks as well. High availability is achieved through reliable systems and through network-level high availability features.

3.5 Regulatory Issues

Regulatory compliance might necessitate lawful intercept and emergency services on publicly available VoIP networks. The US FCC has mandated that VoIP networks to comply with the Communications Assistance for Law Enforcement Act (CALEA). Currently not all countries have such regulations, but it is recommended that networks should be designed with this in mind. For additional information on CALEA please refer to the Foundry Networks whitepaper titled "Lawful Intercept and Enablement in High-Performance Networks" [3].

3.6 Security

Threats to the VoIP network can involve service disruption, capturing voice conversations, theft of service, and attacks to other devices through VoIP servers. Each of these potential risks to the network should be addressed while designing the network.

3.7 Future-proof Infrastructure

It is important to build a scalable infrastructure that can adapt to changing standards and grow to accommodate future VoIP capacity needs as well as additional higher bandwidth services.

4 Reference Architecture for IP Telephony

Figure 1 - Reference Architecture for an IP Telephony Network

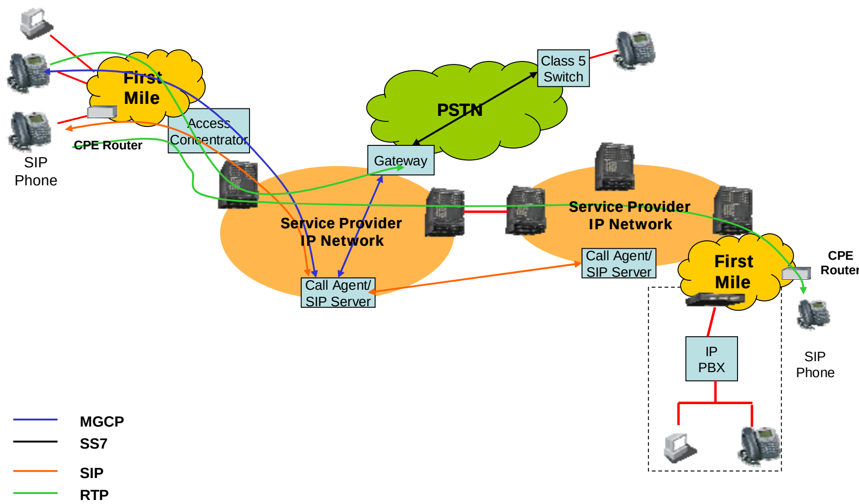

5 Components of a VoIP Network

5.1 End Systems

These voice service endpoints include IP Phones and PBX systems. They interact with the SIP Server or call agent using a call signaling protocol like SIP or H.323 as described below or use a device control protocols such as H.248 or MGCP. Depending upon the network solution a variety of end user terminals can be supported including Plain Old Telephone System (POTS) phones, IP phones, Private Branch Exchanges (PBXs) as well as Personal Computer (PC) soft clients like Skype and other Web based applications.

5.2 Voice Codecs

At the user edge the codecs (Coder-decoders) encode and decode the voice using one of the ITU-T standards. The most commonly used audio codecs are G.711, G.723 and G.729. The G.711 standard operates at 64Kbps and is frequently used in Local Area Network (LAN) environments. The G.729 standard operates at 8Kbps and is often used over the Internet and lower speed links where better compression is desired. The typical frame based voice coders wait for the audio data stream to reach a certain size before they process it. The audio codec frames may be compressed and placed in a Real-time Transport Protocol (RTP) packet. RTP is a User Datagram Protocol (UDP) based protocol carried over IP networks. Since there are many layers involved, each layer introduces a fixed overhead on the voice packets (Figure 2). Various schemes can be used such as better compression and multiplexing to reduce overhead, but they add to the delay introduced by the packetization process. The packet is then carried over an IP network that introduces additional delay and jitter.

Figure 2 - Audio stream packetization

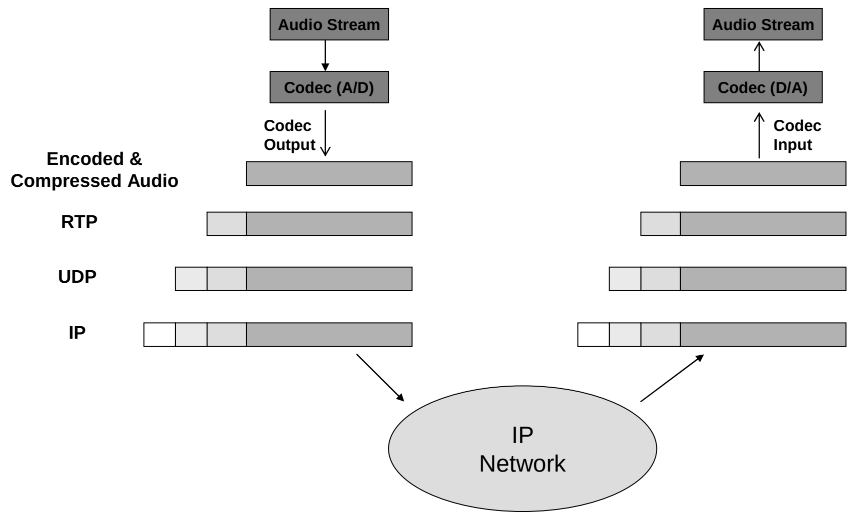

5.3 *Call Agent*

Typical functions of a call agent include methods to connect the endpoints, maintain call state, event reporting, add or subtract media streams in a session and report call statistics. Besides call control they may provide additional services and/or an interface to the application server for other services. In a SIP signaling network, SIP Servers and SIP Clients provide a role that is similar to call agent. The call agents are also known as soft switches, Media Gateway Controllers or Call Servers.

5.4 *Access Concentrator*

These service provider devices terminate the first mile WAN links. Examples of these first mile concentrators include Digital Subscriber Access Line Multiplexers (DSLAMs) for Digital Subscriber Line (DSL) networks and Cable Modem Termination Systems (CMTSs) for cable networks.

5.5 *Router*

The router, the essential part of the carrier IP network, routes IP traffic into the backbone. The responsibility of the edge and core routers to maintain a reliable network is the key to providing high quality voice networks. The routers are also expected to provide a low packet loss in the network, with low latency and jitter. Advanced traffic engineering is frequently used to provide good QoS in such networks.

5.6 *Gateway*

The media gateway acts as a gateway between the Service Provider's IP network and TDM-based PSTN network. It is part of the service provider network and provides transcoding and trans-signaling between VoIP and PSTN networks. The media gateway is controlled by the call agent using MGCP or MEGACO.

6 Foundry's Solutions for VoIP

Foundry Networks provides a wide variety of network infrastructure options for delivering services that require high availability, ultra-low latency and jitter, and strict Service Level Agreements (SLAs). The products offer network level resiliency, hardware based QoS, easy manageability, and flexible network processor based architecture. Foundry products include:

- The NetIron XMR Series of high-end, carrier-class, MPLS backbone routers. Available in four form factors, the NetIron XMR 4000, XMR 8000, XMR 16000 and XMR 32000 routers are designed from the ground up for high performance and scalability to address the needs of the most demanding converged networks. They offer full Layer 2, Layer 3 and advanced MPLS capabilities. The XMR routers offer wire-speed performance on all ports, full IPv6 routing today, and 100-Gbps of user bandwidth per full slot.
- The NetIron MLX Series of metro switching routers, with full Layer 2, Layer 3 and advanced MPLS capabilities. Available in four form factors, the NetIron MLX Series is the only MPLS-enabled metro switching router that offers wire-speed performance on all ports, full IPv6 unicast and multicast routing today, and 100-Gbps of user bandwidth per full slot. The NetIron MLX Series offers unparalleled flexibility to providers in choosing the right network design, and is specifically designed for advanced capabilities desired for voice, video and data edge networks.
- The NetIron CES 2000 Series of compact edge/aggregation switches are purpose-built for Carrier Ethernet service delivery. The switches offer 24-48 ports of 1 GbE and optionally, 2x10-GbE ports in a 1RU compact form factor. Built for compliance with MEF9 and MEF14 specifications, the NetIron CES 2000 Series of switches support both Provider Backbone Bridging (PBB), Provider Bridging (PB) technologies along with comprehensive operation, administration, and maintenance (OAM) capabilities for Ethernet. When used in combination with the NetIron XMR/MLX product families, these platforms can be used to deliver scalable voice, video and data services from the edge of the network to the core.
- The NetIron M2404 metro access switches are high-density, versatile, compact MPLS metro access switch family occupying only a single rack unit (RU). The device features 24 10/100 ports, two 1-GbE fiber/copper combo ports, and two enhanced services 1-GbE ports. It is available in both fiber and copper (RJ-45) models. In addition to the rich Layer 2 Ethernet features, it provides native support for MPLS-based Layer 2 VPN services such as Virtual Leased Line (VLL) and multipoint Hierarchical VPLS (H-VPLS). It also supports comprehensive OAM capabilities for Ethernet to facilitate fast detection and troubleshooting of network faults. These capabilities make the NetIron M2404 switches the most versatile MPLS MTU devices available in the industry today.
- The BigIron RX Series of advanced Layer 3 switches provide full Layer 2 and Layer 3 functionality. Available in four form factors, the BigIron RX Series is an ideal fit for data center applications where high-performance server aggregation is needed. The platform offers wire-speed performance on all ports and massive bandwidth scalability up to 100-Gbps of user bandwidth per full slot. In addition, the BigIron RX Series Layer 3 switches provide hardware-based IPv6 routing to facilitate the inevitable transition to IPv6 in the future.
- The FastIron Edge X (FESX) Series of compact switches offers 24-48 ports of 1 GbE and optionally, 2x10-GbE ports in a 1.5RU compact form factor. Availability of both copper and fiber interface options for the 1 GbE ports makes FESX an appealing choice for aggregating traffic from several access nodes that are either co-located or geographically distributed.
- The FastIron Edge Switch (FES) Series of compact switches offers a highly adaptable feature set combined with the highest 10/100Base-TX and Gigabit Ethernet port densities in its class. The FES switches provide feature-rich switching and Layer 3 multi-protocol

routing capabilities, comprehensive hardware and software redundancy, complete QoS controls including prioritization and rate limiting, and integrated Copper Gigabit Ethernet ports. This makes FES a great fit for a range of applications from aggregation to converged voice, video and data applications.

Figure 3 - Example Architecture for VoIP network using Foundry Routers and Switches

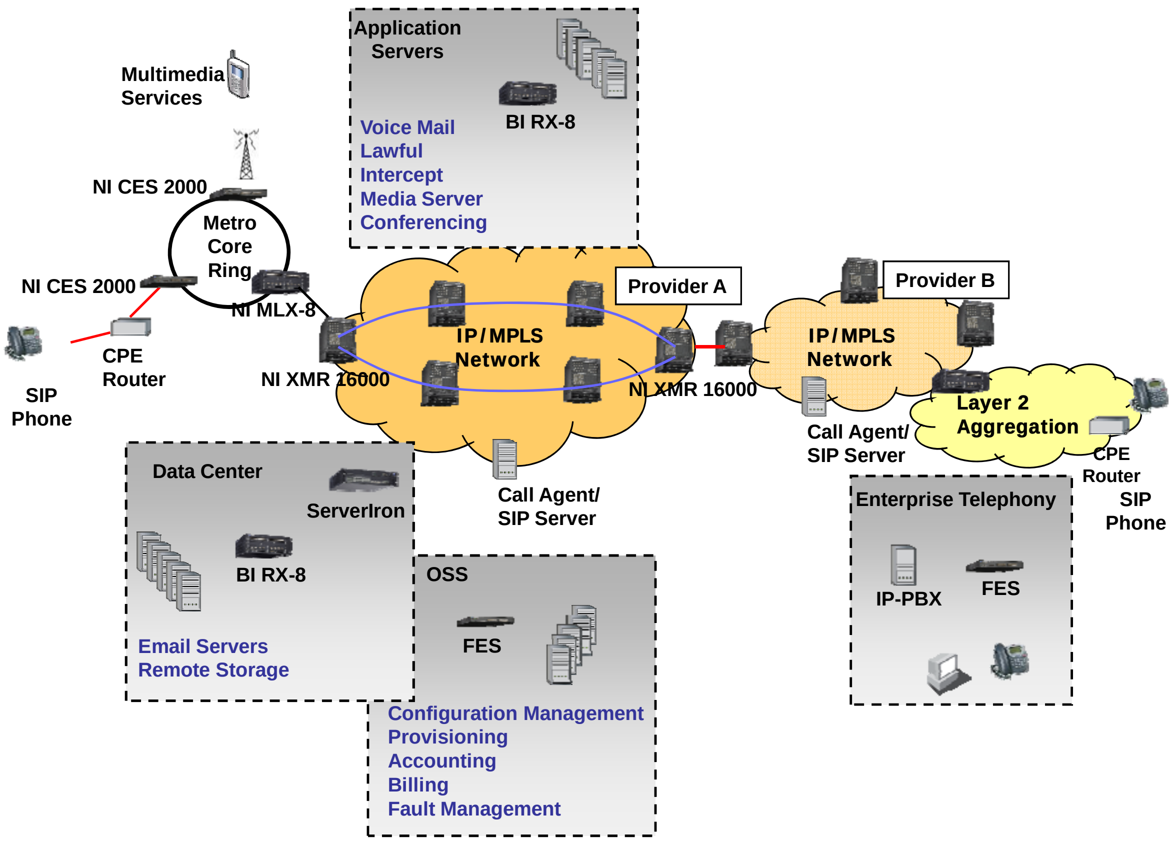

7 Advanced Capabilities for VoIP Service Providers

Foundry's broad range of high performance routing and switching products offer advanced capabilities for VoIP networks. The products provide multiple form factors to suit the needs of most network infrastructures. They also offer full featured Layer 2, Layer 3 and advanced MPLS capabilities that give service providers the flexibility to choose a network design that offers converged voice, data, and other multimedia services.

7.1 High Availability

High availability is achieved through a combination of hardware and software architecture. Foundry's NetIron XMR and MLX, and BigIron RX architectures feature a fully redundant design with no single point of failure. All the system modules are hot pluggable and removal of any system module does not impact the performance of the rest of the platform. The modular architecture of the Multi-Service IronWare® operating system has several high availability features that distinguish it from legacy operating systems that run on other routers:

- Industry-leading cold restart time of less than a minute
- Support for hitless software upgrade
- Hitless Layer 2 and Layer 3 failovers
- Sub-second switchover to the standby management module if a communication failure occurs between active and standby management modules.

These capabilities increase network-level reliability by minimizing the impact of a node failure. The MPLS feature set on NetIron XMR/MLX includes path protection and Fast Reroute (RFC 4090) to ensure 50ms failover around failed links or nodes. The combination of hardware and software architecture, as well as network level reliability features makes NetIron and BigIron platforms perfectly suited for VoIP networks.

7.2 Quality of Service

Foundry's NetIron XMR/MLX platforms offer hardware based QoS mechanisms to prioritize the use of available bandwidth and to manage congestion in the network. Weighted Random Early Discard (WRED) can be used for congestion avoidance. WRED enables the system to detect signs of congestion and take corrective action. This way the system can selectively discard lower priority traffic when the system gets congested.
When the respective QoS features are enabled, arriving traffic is classified and processed based upon the packet priorities. Packet prioritization can be based upon the following:

- Layer 2 CoS as defined in IEEE 802.1p
- Layer 3 IP precedence as defined in RFC 791
- Layer 3 DSCP
- MPLS EXP

The platform has extensive packet marking capabilities to allow changing of QoS information for treatment in next hop. This is useful when the packet is traversing from one network to another. The platform also supports remarking of packet priority based on the result of 2-rate 3-color policer. Additionally, tiered QoS guarantees can be obtained by using scheduling mechanisms such as strict priority, weighted fair queuing or a combination at the output port.

A unique characteristic of the NetIron XMR/MLX routers and BigIron RX switches is the use of a distributed buffering scheme that maximizes the utilization of buffers across the whole system during congestion. This scheme marries the benefits of input-side buffering (Virtual Output Queuing) with those of an output port-driven scheduling mechanism. Input buffering using virtual output queues ensures that bursty traffic from one port does not hog too many buffers on an output port. An output-port driven scheduling scheme ensures that packets are sent to the output port only when the port is ready to transmit a packet. Additional details can be found in the respective product architecture brief [4] [5].

For the network edge, Foundry's NetIron CES 2000, M2404, and FES/FESX platforms provide a rich set of QoS controls. The FES/FESX platforms support packet prioritization based on 802.1p, Type of Service (TOS), DSCP, and Access Control Lists (ACL). They offer flexibility in queuing methods supporting Weighted Round Robin (WRR), strict priority queuing, or a combination. In addition, the NetIron CES 2000 supports up to 8 queues per port, each with a distinct priority level. Advanced QoS capabilities such as the use of 2-rate, 3-color traffic policers, Egress shaping and priority remarking may also be applied to offer deterministic "hard QoS" capability to customers of the service.

Foundry's NetIron XMR/MLX routers, BigIron RX, NetIron CES 2000, NetIron M2404, and FES/FESX switches provide full line-rate forwarding throughput for all ports. These platforms offer hardware based QoS, low latency, and efficient bandwidth management, which can be used to build networks that provide very low latency, jitter, and packet loss. On the NetIron XMR and MLX platforms, MPLS Traffic Engineering can be used to establish dedicated RSVP-TE tunnels and allocate bandwidth in the network. VoIP calls can then be routed over these tunnels to guarantee resource availability through the network.

7.3 Flexible Interfaces

Foundry's routers and switches support a wide variety of interfaces featuring one of the highest densities in the industry. Supported interfaces include 10/100/1000Base-T, 100Base-FX, Gigabit Ethernet, 10Gigabit Ethernet and SONET/SDH (OC-12c/STM-4, OC-48c/STM-16, OC-192c/STM-48).

7.4 Ultra-High Capacity

Foundry's NetIron XMR/MLX routers and BigIron RX switches provide full line-rate forwarding throughput for all modules. The platforms use Foundry Direct Routing that provides consistent wire speed routing without any CPU lookups. This is achieved by maintaining large routing/forwarding tables in hardware at the interface module level. These platforms can scale up to 3.2 Tbps data capacity and up to 2 Billion packets per second (Bpps) per system processing capacity.

7.5 Scalability

NetIron XMR/MLX and BigIron RX platforms are highly scalable and available in four different form factors 4, 8, 16, and 32 slot chasses, offering service providers flexibility to choose a platform that fits their current and projected future needs. The XMR supports 4K VLANs and up to 2 million MAC addresses, 1 million IPv4 routes in hardware, 240K IPv6 routes, 10 million BGP routes, 16K VPLS instances and 2K BGP/MPLS VPNs. The MLX platform supports 4K VLANs and up to 1 million MAC addresses, 512K IPv4 routes in hardware, 112K IPv6 routes, 2 million BGP routes, 4K VPLS instances and 400 BGP/MPLS VPNs . The BigIron RX platform supports 4K VLANs and 64K MAC addresses, 512K IPv4 routes in hardware, 64K IPv6 routes and up to 4 million BGP routes. These platforms lead the industry in port density and line rate throughput.

7.6 Purpose-Built for VoIP and Converged Networks

Foundry Networks has been powering major voice, video and data networks, and its products have enabled service providers to deliver the strict SLAs required for these services. Foundry's product portfolio includes customer premises devices, Layer 2 and Layer 3 aggregation devices, edge routers, and core routers.

Requirements for VoIP Networks	Foundry's Solution
Quality of Service	NetIron, BigIron, and FastIron products offer ultra-low latency, packet prioritization based on CoS, Layer 3 IP precedence, and DSCP. NetIron XMR/MLX platforms, in addition, provide prioritization based on MPLS EXP, and offer MPLS traffic engineered tunnels to select a path around congested links. They also offer path protection and fast-reroute for protection against node or link failures.
Latency	Ultra-low latency on all the products including NetIron XMR/MLX, BigIron RX, NetIron CES 2000 and FastIron Edge Switches.
High Availability	Fully redundant design, hot pluggable modules, hitless software upgrade, hitless Layer 2 and 3 failovers, MPLS FRR
Flexibility in Service Delivery	Support for full Layer 2 and Layer 3, and advanced MPLS functionality provides flexibility in choosing most suitable service delivery model.
Throughput	Full line-rate forwarding using Foundry Direct Routing
Interfaces	Variety of low- and high-speed interfaces available. Industry-leading port density. 100-Gbps of user bandwidth per full slot.
Future Proof Infrastructure	Scalable systems available in various form factors. Support large MAC tables, Layer 3 forwarding tables, Layer 2 and Layer 3 VPNs.

8 Summary

VoIP has provided service providers the ability to offer voice, video, and other multimedia services over a packet-based network. As competition grows, service providers need to differentiate their offerings to reduce churn and increase ARPU. Designing an underlying IP infrastructure that can handle the demands of new services is crucial to their success. Foundry Networks routers and switches enable a wide variety of solutions that make it easier for service providers to deliver advanced VoIP and multimedia services. The products offer various degrees of resiliency through design features that include redundant and hot-pluggable hardware, hitless software upgrades, and graceful restart for routing protocols[2]. The products offer a wide range of capabilities including robust Layer 2 and 3 protocols, IPv6 routing, MPLS, pseudowires, and Layer 2 and 3 Virtual Private Networks (VPNs). These features, along with advanced Quality of Service (QoS), high platform capacity, low latency, and jitter make Foundry products an ideal fit for scalable, cost effective VoIP deployments.

[2] See datasheet of the respective products.

FOUNDRY NETWORKS

9 Acronyms

ARPU	Average Revenue per User
CMTS	Cable Modem Termination System
CoS	Class of Service
Codecs	Coder-decoders
CALEA	Communications Assistance for Law Enforcement Act
DiffServ	Differentiated Services
DSCP	Differentiated Services Code Point
DSLAM	Digital Subscriber Access Line Multiplexer
DSL	Digital Subscriber Line
IP	Internet Protocol
ISP	Internet Service Providers
LAN	Local Area Network
MEGACO	MEdia GAteway Control
MGCP	Media Gateway Control Protocol
MGC	Media Gateway Controller
MSO	Multiple System Operator
MPLS	Multi-Protocol Label Switching
PHB	Per Hop Behavior
PC	Personal Computer
POTS	Plain Old Telephone System
PBX	Private Branch Exchanges
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RU	Rack Unit
RTP	Real-time Transport Protocol
SLA	Service Level Agreements
SIP	Session Initiation Protocol
SMB	Small-Medium Business
TE	Traffic Engineering
UDP	User Datagram Protocol
VoQ	Virtual Output Queuing
VPN	Virtual Private Networks
VoIP	Voice over IP
WRED	Weighted Random Early Discard

10 References

1. Lawful Intercept and Enablement in High-Performance Networks, http://www.foundrynet.com/pdf/wp-lawful-intercept.pdf
2. MultiService Forum Release 3 Architecture Framework, http://www.msforum.org/techinfo/approved/MSF-ARCH-003.00-FINAL.pdf
3. Lawful Intercept and Enablement in High-Performance Networks, http://www.foundrynet.com/pdf/wp-lawful-intercept.pdf
4. NetIron XMR architecture brief, http://www.foundrynet.com/pdf/wp-ni-xmr-architecture.pdf
5. BigIron RX architecture brief, http://www.foundrynet.com/pdf/wp-bi-rx-architecture-brief.pdf

Author: Rahul Vir
Document version 1.1

Foundry Networks, Inc.
Headquarters
4980 Great America Parkway
Santa Clara, CA 95054-1200
U.S. and Canada Toll-free: (888) TURBOLAN
Direct telephone: +1 408.207.1700
Email: info@foundrynet.com
Web: http://www.foundrynet.com