

A newsletter published
three times per year by
Isocore



www.isocore.com

Isonotes

Volume 4, Issue 2

July 2005

In this issue:

- MPLS/GMPLS Perspectives
- From the Editor
- High Availability: The Key to Service Convergence
- What's Happening with IETF?
- QoS in Service Provider Networks
- Contributing Writers Biographies
- Calendar of Events

MPLS/GMPLS Perspectives

Kireeti Kompella, Juniper Networks

The Maturation of MPLS

Over the past several years, MPLS has gone from being a "wannabe" technology, to the new ATM, to a must-have, to "What? Haven't you deployed MPLS yet??" To some extent, this is because MPLS has proven itself flexible and capable of solving diverse problems:

- ◆ MPLS has delivered on its promises (VPN services, Quality of Service, Fast Restoration, Traffic Engineering, High Availability, etc.);
- ◆ and, more importantly, it has demonstrated its *economic* value: attracting new revenue

(IP VPNs, Virtual Private LAN Service) as well as providing cost savings (network convergence, lower cost resilience).

The question is no longer "Will MPLS ever be ready?" or even "When will MPLS be ready?", but "When will legacy services be fully migrated to MPLS?" More soberly, "What is missing to make MPLS the default technology?"

Fortunately, MPLS isn't resting on its laurels. Innovation hasn't ceased. Some of the last missing pieces are being filled in. Long requi-

sites in carrier networks, i.e. OAM, User-to-Network, and Network-to-Network Interfaces, have now come to MPLS (LSP ping, BFD for MPLS, MPLS UNI and NNI). New ideas, such as "point-to-multipoint" LSPs (i.e., MPLS multicast) and inter-provider LSPs, are being finalized. MPLS is also proving to be a perfect complement to Ethernet in metro area networks. New applications, such as "multi-segment pseudo-wires" and MPLS-based DSLAMs, are being considered. Perhaps most important of all, there are good management tools for MPLS networks for pro-

(Continued on page 3)

From the Editor



REGISTER NOW for the 8th Annual International Conference on MPLS and related technologies. MPLS 2005 will be held in Washington, DC from October 16-19, 2005, and is supported by key leading Service Providers in the industry including NTT, MCI, AT&T, KDDI, DT, Acreo, France Telecom, and BTNAccess. It will be followed by the 6th Public Interoperability Demonstration on October 20. The 4-day conference, co-

MPLS 2005 dates are approaching fast!

Register NOW for early bird discounts!

chaired by Adrian Farrel (Co-chair, IETF CCAMP WG) and Dave McDysan (MCI), features:

- Tutorials, technical sessions, panel discussions, and exhibits
- Presentations by industry-leading experts focusing on the latest developments of IP, MPLS, GMPLS, and related networking technologies
- Panel Sessions led by the service provider and vendor communities encourage audience participation in a lively discussion and provide valuable insight into industry trends.

The conference Exhibit Hall showcases the latest prod-

ucts and services by key vendors in the industry – including ADC, Alcatel, Avici Systems, Chiaro Networks, Ciena Corporation, Cisco Systems, ECI Telecom (formerly Laurel Networks), Extreme Networks, Isocore Corporation, IXIA, Juniper Networks, Lucent Technologies, Mangrove Systems, Nortel Networks, Quarry Technologies, Redback Networks, Spirent Communications, Tellabs, and WANDL. Attendee registration is now open, with substantial early-bird discounts through July 22. Please visit www.mpls2005.com for an update on the conference program.

Kavita Khanna, Isocore
Khanna@isocore.com
703-860-9275

High Availability: The Key to Service Convergence

Eric Brendel, Chiaro

Providers and customers for years have been asking for more reliable equipment and services. Now, as more traditional circuit-based services are being placed onto IP/MPLS based networks the demand for high availability (HA) has reached a new importance. With services such as VOIP, VPN, and IPTV being transplanted onto IP/MPLS networks, the expectation of continuous operation is real. Interruptions to these services, even short ones, are not as easily tolerated as interruptions to traditional Internet based traffic.

The answer is not simply to build more reliable components. Nor is it to simply add redundancy, either at the network level or in the hardware. Technology must develop to facilitate non-stop, continuous operation in light of planned and unplanned outages. The result will be devices that can be upgraded and operated on

without disrupting the data flow. They must also be capable of protecting and rerouting traffic during failures...something that is common in telecom equipment, but not yet in routers. In fact, today, most networks must still interrupt network operations to perform simple capacity expansions and upgrades.

Vendors are responding. Technology now exists to provide continuous operation during planned and unplanned outages. Some of these technologies include:

Protection Against Control Plane Resets. When routers are reloaded due to planned maintenance or failures, the routing and signaling state must be re-obtained. During this time routers typically clear the forwarding plane resulting in an outage. With the current large size of BGP tables along with its slow convergence characteristics, it can

potentially take minutes to reestablish the network state. Two approaches have been developed to address control plane resets: Stateful Protocol Protection and Graceful Restart techniques with non-stop forwarding.

- Stateful Protocol Protection protects the routing and signaling protocols by maintaining each routing protocol state and underlying communication state during a switchover from a primary route processor to a secondary route processor. Because the routing protocols are never interrupted, routing protocol state is maintained and traffic forwarding is not interrupted.
- Graceful Restart techniques also protect against control plane resets by extending the protocols so they recover gracefully from a switchover. With grace-

ful restart techniques, routing protocol sessions fail. However, forwarding tables are locked while the routing protocols re-establish communication and re-synchronize.

Both techniques solve the same problem and are complementary, however, there has been significant debate as to which performs and scales better. It is safe to say that both techniques will likely be found within networks and that individual implementations will need to be tested for scale and performance to determine whether the technology is effective in real world scenarios. For reliable service, it is also important that all protocols are protected and not just a select set.

Fast Reroute and Convergence. To provide uninterrupted traffic flow during a facilities or complete

(Continued on page 3)

What's Happening with IETF?

Ron Bonica, Juniper Networks

The L3VPN working group met in Minneapolis on March 7. Although the working group has only produced three RFCs to date, it has been very productive. Two additional working group drafts will become RFCs very soon (at about the same time that this note publishes). These include a general framework for L3 VPNs and a security framework for PPVPNs. Three more documents are awaiting publication in the RFC Editor's queue: a management framework document, the BGP/MPLS base document (also known as 2547bis), and its' associated applicability statement.

Recently, textual conventions for VPN management and a MIB for MPLS/BGP L3VPNs have been approved by the IESG. These drafts will also be placed in the RFC Editor's queue soon.

Nine additional drafts have passed the working group's last call and are in the IESG review process. These include the CE-based and virtual router-based VPN architecture documents and several implementation alternatives for BGP/MPLS-based VPNs, including alternative encapsulations, constrained route distribution, and distribution of IPv6 VPN routes.

New work addresses multicast in the BGP/MPLS VPN environment. Thomas Morin and his colleagues have submitted a document addressing multicast requirements. The authors of three previous multicast solution drafts have combined their talents to produce a single draft to be progressed by the working group.

The L3VPN working group does not anticipate taking on any new work items. Although the working group still has several work items in progress, multicast is the only remaining work item which requires significant additional new work.

Isocore provides a program of cutting-edge technology validation in the critical areas of next generation Internet and wireless networking.

One of the objectives of Isocore, which is achieved through its Internetworking Lab, is to advance Internet-working technology in the areas of GMPLS and MPLS.

To find out more about participating in our Internetworking Lab, email us at: info@isocore.com

(Perspectives continued from page 1)

visioning, monitoring and troubleshooting – a sure sign of maturity.

So, What About GMPLS?

GMPLS encompasses all of MPLS RSVP-TE and more, even though most people associate it with signaling in optical networks. GMPLS has contributed much to MPLS, for example, LSP hierarchy (a building block for inter-provider LSPs and a tool for scaling), signaling over unnumbered and bundled links, SRLGs for improved protection and LSP crankback. GMPLS UNI and NNI are complete. However, while there is much interest in GMPLS and some deployment, GMPLS doesn't have the uptake in optical networks that MPLS enjoys today in packet networks.

To be sure, GMPLS is about three years younger than MPLS; it may just be that we need a bit of patience. However, two non-technological factors affect GMPLS's success: the "transport mentality" and economics. Transport mentality refers to the general reluctance that transport network operators have to using IP-based protocols to control their networks; it's not clear whether this is due to a lack of confidence in IP protocols or a concern about network dynamism. Perhaps more important is the economics: few are yet convinced of the value of deploying GMPLS, either in terms of capex/opex savings, or in terms of new revenue. One hopeful avenue is "Layer 1 VPNs" that optical carriers can offer to their customers to lower the cost of connectivity. An-

other is the notion of "Layer 2 connections" signaled with GMPLS. Whatever turns out to be the igniting factor, GMPLS is ready – it is just waiting for that one spark to burst forth and transform transport networks.

(High Availability continued from page 2)

node failure, traffic needs to be rerouted down alternate paths. For RSVP-TE based LSPs, this involves MPLS Fast Reroute that redirects traffic down other LSPs during failures. For IP and LDP based LSPs, the IETF is investigating IP Fast Reroute, a way to pre-calculate backup paths during a re-convergence event with the internal routing protocol.

Equipment Protection.

For years, routers have been built with redundancy for virtually every component. However, having redundant components is not enough. The equipment must be capable of detecting failures quickly and the switchover to the redundant component must occur without losing traffic. This capability is not only beneficial for failures, but it also provides the basis for performing in-service maintenance.

As routers become highly available through the implementation of new HA technology, providers will be able to provide services typically delivered over carrier-class transport networks across their converged highly available IP/MPLS core.

QoS in Service Provider Networks

Azhar Sayeed, Cisco Systems

Business applications stretch network capabilities and resources, but also complement, add value, and enhance every business process. Networks must therefore provide secure, predictable, measurable, and sometimes guaranteed services to these applications. Achieving the required Quality of Service (QoS)-by managing the delay, delay variation (jitter), bandwidth, and packet loss parameters on a network, while maintaining simplicity, scalability, and manageability-is critical to running an infrastructure that truly serves the entire business.

QoS is widely deployed today in Service Provider Networks today. The majority of the providers are offering 3-4 classes of service using the IETF DiffServ model for their customers. This is a provisioned QoS model where packets are marked

with certain values and are treated within the network according to that marking. For example voice traffic is treated with a higher priority than data traffic. It is assumed that the network core has plenty of bandwidth and the congestion is largely at the edge of the network. This is mostly true with core links being OC-48 and OC-192 and edge or access links being anywhere from few Kbps up to several Mbps. Traffic must be prioritized, marked, queued, and scheduled through network elements in a deterministic manner to conform to traffic contracts.

Service providers have largely figured out how to provision, manage, and offer multiple classes of services on a single network. However, challenges remain when it comes to offering QoS across multiple provider networks. The challenges have largely to

do with the lack of a proper standard with respect to how SP1 treats QoS packets marked with a certain code point versus SP2. This results in a QoS offering that can be very different in SP1 versus SP2. If the Enterprise is connecting to multiple providers then the enterprise can experience very different QoS behavior from each of them. This poses a challenge to the enterprise when it comes to QoS network design when deploying VoIP in this network. Standards bodies are tackling with this issue of how to define a consistent QoS behavior and ensure SPs comply to this behavior. Interprovider QoS is a hot topic in the industry circles and is also a topic for the most recently formed IP-Sphere forum. Interprovider QoS connections exist today but these are relatively small and are mutually agreed capabilities between two providers and not a

universal set that is applicable to any provider. QoS mapping must be done at provider boundaries to achieve the desired behavior. Billing, Accounting and provisioning is another challenge when the QoS mapping is different. Preserving QoS transparency across provider boundaries is important. Until such time that either providers either agree or there is an agreed upon Interoperable implementation or standard for provisioning, accounting, managing and billing for such Interprovider services it will remain a challenge to obtain uniform QoS based services from multiple providers for IP connectivity. This topic is hotly being pursued and discussed at the MFA (MPLS, Frame and ATM) forum and IPSphere forum meetings and we can see more in the upcoming months.

www.isocore.com

12359 Sunrise Valley Dr.
Suite 100
Reston, Virginia 20191

Phone: 703-860-1777
FAX: 703-860-1778

Email: info@isocore.com
www.isocore.com

To submit an article for
Isonotes contact...

editor@isocore.com

Upcoming Issues of Isonotes...

Publication Date	Deadline Date
Nov. 2005	10/24/05
Mar. 2006	2/22/06
July 2006	06/23/06



Kireeti Kompella is a Juniper Fellow at Juniper Networks. His current interests are all aspects of MPLS including Traffic Engineering, Generalized MPLS, VPNs, VPLS and the application of MPLS to Metro Ethernet. Dr. Kompella is a co-chair of the CCAMP Working Group and the author of several Internet Drafts and RFCs.



Eric Brendel has been involved with network architecture and design for over 10 years. At Chiaro, Eric evaluates next-generation network technologies to further the advancement of Chiaro's Enstara IP/MPLS routing platform. In addition to being a network architect, he has also served as a network security consultant as well as a product architect for a large-scale switch start-up. Eric is currently working on his dissertation towards a PhD in Electrical Engineering.



Ron Bonica is an engineer, employed by Juniper Networks' software development group. His specialties are secure network operation and Layer 3 Virtual Private Networking. Prior to joining Juniper Networks, Mr. Bonica was Sr. Manager of engineering for MCI's vBNS+ network. Mr. Bonica currently serves as co-chairman of the IETF Layer 3 VPN working group and is a past chairman of the IETF Common Control and Measurement Plane (CCAMP) working group.

*Photo not
available
at time of
publication*

Azhar Sayeed has more than 15 years of networking and communications industry experience. As product line manager with the IOS Technologies Division, Cisco Systems, Mr. Sayeed is responsible for product management and rollout of IP Routing, Quality of Service (QoS) and Multi Protocol Label Switching (MPLS) features in Cisco's Internet Technologies Division.

Calendar Of Events

OCTOBER 16-19
MPLS 2005 International Conference
Washington, DC

©2005 Isocore All Rights Reserved

ISOCORE

12359 Sunrise Valley Dr.
Suite 100
Reston, Virginia 20191