

Network Working Group
Internet Draft
Expiration Date: August 2000

Kireeti Kompella (Juniper Networks)
Yakov Rekhter (Cisco Systems)
Daniel O. Awduche (UUNET)
Alan Hannan (Global Crossing)
Gisli Hjalmtysson (AT&T Labs)
Joe Lawrence (Level 3 Communications)
Satoru Okamoto (NTT)
Debashis Basak (Marconi)
Greg Bernstein (Ciena Corporation)
John Drake (Chromisys)
Near Margalit (New Access Communications)
Ed Stern (Sirocco Systems)

Extensions to IS-IS/OSPF and RSVP in support of MPL(ambda)S

[draft-kompella-mpls-optical-00.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Internet Draft [draft-kompella-mpls-optical.txt](#)

Februrary 2000

[2.](#) Abstract

This document specifies extensions to the IS-IS, OSPF, and RSVP in support of Multiprotocol Lambda Switching (MPL(ambda)S). For motivations and the overall architecture of MPL(ambda)S see [Awduche-1]. For the set of required enhancements to IS-IS, OSPF, and RSVP see [Basak].

[3.](#) IS-IS/OSPF enhancements

In this section we describe enhancements to IS-IS and OSPF in support of MPL(ambda)S. These enhancements are in addition to the extensions required to support MPLS Traffic Engineering [Smit], [Katz].

[3.1.](#) Link Type TLV

A network may have links with different characteristics. For example, a node may not be able to demultiplex the data it receives on a given link on a packet-by-packet basis, but it may be to multiplex or demultiplex channels within a SONET payload.

The Link Type sub-TLV allows to identify a particular type of a link.

In IS-IS, this TLV is a sub-TLV of the Extended IS Reachability TLV, with type 19. In OSPF, this TLV is a sub-TLV of the Link TLV, with type 1 (this sub-TLV currently exists in OSPF; here, we extend the range of values, and introduce a further sub-TLV.)

The length is 1 plus the total length of all sub-TLVs. The value has a fixed length field of one octet, followed by optional sub-TLVs. The value is taken from the following list:

Value	Link Type
1	Packet-Switch Capable-1 (PSC-1)
2	Packet-Switch Capable-2 (PSC-2)
3	Packet-Switch Capable-3 (PSC-3)
4	Packet-Switch Capable-4 (PSC-4)
200	Time-Division-Multiplex Capable (TDM)
210	Lambda-Switch Capable (LSC)
220	Fiber-Switch Capable (FSC)
240	Forwarding Adjacency PSC (FA-PSC)

245	Forwarding Adjacency TDM (FA-TDM)
250	Forwarding Adjacency LSC (FA-LSC)

If a link is of type PSC-1 through PSC-4, that means that the node receiving data over this link can demultiplex (switch) the received

data on a packet-by-packet basis. The various levels of PSC establish a hierarchy of LSPs tunneled within LSPs.

If a link is of type TDM, that means that the node receiving data over this link can multiplex or demultiplex channels within a SONET payload.

If a link is of type LSC, that means that the node receiving data over this link can recognise and switch individual lambdas within the link (fiber).

If a link is of type FSC, that means that the node receiving data over this link (fiber) can switch the entire contents to another link (fiber) (without distinguishing lambdas, channels or packets).

If a link is of type FA-PSC, that means that the link is a forwarding adjacency (see [section 3.5](#)). Furthermore, the egress node of the corresponding FA-LSP is capable of switching the data received over the FA-LSP on a packet-by-packet basis.

If a link is of type FA-TDM, that means that the link is a forwarding adjacency. Furthermore, the egress node of the corresponding FA-LSP is capable of multiplexing and demultiplexing the SONET channels received over the FA-LSP.

If a link is of type FA-LSC, that means that the link is a forwarding adjacency. Furthermore, the egress node of the corresponding FA-LSP is capable of recognising and switching the lambdas received over the FA-LSP.

[3.2](#). Link Media Type TLV

A link may support a set of media types, where by media type we mean both the bit encoding format as well as the bit transmission rate. By support we mean that the link has the ability to carry and switch

a signal of one or more of these media types depending on the resource availability and capacity of the link. The link media type should not be confused with the physical link encoding or multiplex structure (if any). These properties are of local significance between neighbor nodes. The types of media a link can support (and its associated node switch) are of overall concern to routing as well as information on link capacity (bandwidth).

A typical example of multiple transport methods with the same routing representation currently occurs in the optical domain. Currently there are systems where four OC-48s are switched and transported over a single fiber via wavelength division multiplexing (WDM)

technology. In addition there are systems where four OC-48s are transported in a "transparent" OC-192 time division multiplex (TDM) structure, i.e., all the overheads of the OC-48's are persevered. In both these cases the essential information from a routing perspective is that the links can both transport media of type OC-48.

The Link Media Type TLV is a list of supported media types for a link. In IS-IS, this TLV is a sub-TLV of the Extended IS Reachability TLV with type 20. In OSPF, this TLV is a sub-TLV of the Link TLV, with type 10.

The length is the length of the list of Link Media Types in octets. The value is a list of two-tuples of which the first field is a two-octet value which defines the media type, and the second field is a one-octet value which defines the lowest priority at which that media type is available. Media types values are taken from the following list:

Value	Bit Rate	Encoding	
<n>	OC-<n>	SONET	1 <= <n> <= 3072
3072 + <n>	OC-<n>	SDH	1 <= <n> <= 3072
6144 + <n>	OC-<n>	Clear	1 <= <n> <= 3072
9217	DS0	DS0	
9218	DS1	DS1	
9219	E1	E1	
9220	DS2	DS2	
9221	E2	E2	
9222	DS3	DS3	
9223	E3	E3	

9224	J3	J3
9225	DS4	DS4
9226	E4	E4
9227	J4	J4
9228	1Gbps	GigE
9229	10Gbps	10GigE

Link Media Types (LMTs) present a new constraint for LSP path computation. Specifically, when a sequence of links traversed by an LSP includes one or more subsequence of links that carry the LMT TLV, then for all the links within each such subsequence their bit rate (as specified in the LMT TLV) has to be the same and at least the LSP's desired bandwidth, and the encoding (as specified in the LMT TLV) has to be the same.

On a bundled link we assume that either the whole link is configured with the Link Media Types, or each of its component links are configured with the Link Media Types. In the latter case, the Link

Media Type of the bundled link is set to the set union of the Link Media Types for all the component links.

[3.2.1.](#) Path sub-TLV

When an LSP advertises a forwarding adjacency into an IGP (IS-IS or OSPF), it may be desirable to carry the information about the path taken by this adjacency. This information may be used for path calculation by other LSRs. This information is carried in the Path sub-TLV, which is a sub-TLV of the Link Type TLV.

In both IS-IS and OSPF, this sub-TLV is encoded as follows: the type is 1, the length is 4 times the path length, and the value is a list of 4 octet IPv4 addresses identifying the links in the order that they form the path of the forwarding adjacency.

[3.2.2.](#) Identifying links that can't terminate an LSP

As discussed in [[Basak](#)], some links may not be capable of terminating LSPs, and this property of the links needs to be taken into account

when performing LSP path computation as well as RSVP signalling. Such links are identified by using the Link Type sub-TLV.

Note that the notion of 'Termination Capable' has been generalized to 'packet-switch capable', 'TDM capable', 'lambda-switch capable', and 'fiber-switch capable'.

3.3. Shared Risk Link Group TLV

A set of links may constitute a 'shared risk link group' (SRLG) if they share a resource whose failure may affect all links in the set. For example, two fibers in the same conduit would be in the same SRLG. A link may belong to multiple SRLGs. Thus the SRLG TLV describes a list of SRLGs that the link belongs to. An SRLG is identified by a 32 bit number that is unique within an IGP domain.

The SRLG of a path is the union of the SRLGs of the links in the path. If an LSP is required to have multiple diversely routed paths, the path computation should attempt to route the paths so that they do not have any links in common, and such that the path SRLGs are disjoint.

In IS-IS, the SRLG TLV is a sub-TLV of the Extended IS Reachability TLV with type 21. In OSPF, this TLV is a sub-TLV of the Link TLV, with type 11. The length is the length of the list in octets. The

value is an unordered list of 32 bit numbers that are the SRLGs that the link belongs to.

3.4. Control channels, data channels, and IP links

A pair of OXCs are neighbors from the MPLS point of view if they are connected by both one or more fibers and one or more (logical) control channels. In what follows, we consider only those fibers that it is desired to announce into IS-IS/OSPF for LSP establishment. (The administrator may decide to reserve some fibers, for example for manual provisioning; these fibers are not to be announced into IS-IS/OSPF.)

For a given pair of OXCs connected by a set of fibers and one or more

(logical) control channels, both of the OXCs must associate the fibers in the set with the control channels and form IP links in a consistent fashion. One way to accomplish this is via configuration. While other mechanisms to accomplish this may be possible, the specification of such mechanisms is outside the scope of this document.

If all the fibers can share the same TE characteristics, then a single control channel suffices. From the IS-IS/OSPF point of view this control channel, together with all the fibers, forms a single IP link. This IP link could be either unnumbered, or numbered.

There are cases where all the fibers between a pair of OXCs cannot be announced as one IS-IS/OSPF link. This may be because not all the fibers can share the same TE characteristics (for example, some are Termination Capable and others are not). Or this may be because it is administratively so desired, for example, some fibers have metric 10 and others metric 20. In such a case, the fibers must be divided into sets that share all TE characteristics. Corresponding to each set, there must be a logical control channel. Each set of fibers, together with the corresponding logical control channel forms an IP link. Each such link should be numbered. All the multiple logical control channels could be realized via the common control channel.

When an adjacency is established over a (logical) control channel that is part of an IP link formed by the channel and a set of fibers, this link is announced into IS-IS/OSPF as a "normal" link; the fiber characteristics are represented as TE parameters of that link. If there are more than one fiber in the set, the set is announced using the "bundling" technique described in [[Kompella](#)].

Note that while there may be as many adjacencies established as there are logical control channels, for flooding purposes, it suffices to

flood over a single logical control channel.

[3.4.1](#). Excluding data traffic from control channels

The control channels between OXCs, or between an OXC and a router are generally meant for low-bandwidth control and administrative traffic. These control channels are advertised into IS-IS/OSPF as normal IP

links as mentioned in the previous section; this allows the routing of (for example) RSVP messages and telnet sessions. However, if routers on the edge of the optical domain attempt to forward data traffic over these channels, the channel capacity will quickly be exhausted.

If one assumes that data traffic is sent to BGP destinations, and control traffic to IGP destinations, then one can exclude data traffic from the control plane by restricting BGP nexthop resolution. (It is assumed that OXCs are not BGP speakers.) Suppose that a router R is attempting to install a route to a BGP destination D. R looks up the BGP nexthop for D in its IGP's routing table. Say R finds that the path to the nexthop is over interface I. R then checks if it has an entry in its Link State database associated with the interface I. If it does, and the link is not packet-switch capable, R installs a discard route for destination D. Otherwise, R installs (as usual) a route for destination D with nexthop I. Note that R need only do this check if it has packet-switch incapable links; if all of its links are packet-switch capable, then clearly this check is redundant.

Other techniques for excluding data traffic from control channels may also be needed.

[3.5](#). Forwarding adjacencies

An LSR at the head end of an LSP may advertise this LSP as a link into a link-state IGP (e.g., IS-IS, OSPF). When this link is advertised into the same instance of the IGP as the one that determines the route taken by the LSP, we call such a link a "forwarding adjacency". We refer to the LSP as the "forwarding adjacency LSP", or just FA-LSP.

In general, creation/termination of a forwarding adjacency and its FA-LSP could be driven either by mechanisms outside of MPLS (e.g., via configuration control on the LSR at the head-end of the adjacency), or by mechanisms within MPLS (e.g., as a result of the LSR at the head-end of the adjacency receiving LSP setup requests originated by some other LSRs).

Forwarding adjacencies could be represented as unnumbered links. In this case the link IP addresses of forwarding adjacencies are the router IDs on the two ends of the link. Alternatively, forwarding adjacencies could be represented as numbered links. In this case the link IP addresses of forwarding adjacencies could be addresses assigned to some "virtual" interfaces on a router (it is assumed that a router may have multiple virtual interfaces).

If there are multiple LSPs that all originate on one LSR, and all terminate on another LSR, then at one end of the spectrum all these LSPs could be merged (under control of the head-end LSR) into a single forwarding adjacency, while at the other end of the spectrum each such LSP could be advertised as its own adjacency.

When a forwarding adjacency is created under administrative control (static provisioning), the following parameters can be configured for the FA-LSP: the head-end address (if left blank, this will default to the head-end LSR's Router ID); the tail-end address (this must be configured, and must be either the Router ID of the tail-end LSR of the forwarding adjacency, or an interface address on the tail-end LSR); bandwidth and resource colors constraints. The path taken by the FA-LSP may be computed by the Constrained SPF (CSPF) mechanism of MPLS TE, or by explicit configuration; this choice is determined by configuration.

When a forwarding adjacency is created dynamically, its parameters are inherited from the LSP which induced its creation. Note that the bandwidth of the FA-LSP must be at least as big as the LSP that induced it, but may be bigger if only discrete bandwidths are available for the FA-LSP.

When the FA-LSP is established, it is advertised as a forwarding adjacency into IS-IS/OSPF. The Link Type associated with the forwarding adjacency depends on the link type of the last link in the FA-LSP; if this link is packet-switch capable, the Link Type of the forwarding adjacency is 'FA-PSC'; if this link is TDM capable, the Link Type is 'FA-TDM'; if this link is lambda-switch capable, the Link Type is 'FA-LSC'. Some of the attributes of this link can be derived from those of the FA-LSP, but others will need to be configured. Configuration of the attributes of statically provisioned forwarding adjacencies is straightforward; for dynamically provisioned forwarding adjacencies, a policy-based mechanism may be needed to associate attributes to forwarding adjacencies. The path taken by the forwarding adjacency may be advertised using the Path sub-TLV.

The Link Media Type of the forwarding adjacency is the most restrictive of the link media types of the components links of the forwarding adjacency. If none of the components links have a link media type, the link associated with the forwarding adjacency doesn't have a link media type either.

This document restricts the priority of the forwarding adjacency to 0, regardless of how the adjacency is created.

By default the maximum reservable bandwidth and the maximum LSP bandwidth for all priorities of a link associated with a forwarding adjacency is set to the bandwidth of the FA-LSP. These may be overridden via configuration at the head-end of the forwarding adjacency (note that the maximum LSP bandwidth at any priority should be no more than the bandwidth of the FA-LSP).

By default the TE metric on the forwarding adjacency is set to $\max(1, (\text{the TE metric of the FA-LSP path}) - 1)$ so that it attracts traffic in preference to setting up a new LSP. This may be overridden via configuration at the head-end of the forwarding adjacency.

It is expected that forwarding adjacencies will not be used for establishing IGP peering relation between the routers at the ends of the adjacency. It is further expected that forwarding adjacencies will be used only in CSPF computation. In IS-IS, this can be accomplished by setting the default metric of the extended IS reachability TLV the link to infinity ($2^{24} - 1$), and not advertising an IS reachability TLV for the link. In OSPF, this can be accomplished by not advertising the link as a regular LSA, but only as a TE opaque LSA.

If forwarding adjacencies are to be established only under administrative control, it is important that LSRs would not attempt to establish LSPs that have to traverse the optical domain, but for which no FA-LSP is established. To accomplish this, all the optical links (fibers) within the optical domain have to be excluded from CSPF by any LSR outside the domain. This can be achieved by coloring optical links with a specific color, and making sure that all LSPs originating outside the optical domain exclude this color.

[3.6](#). Two-way connectivity check suppression

CSPF should not perform two-way connectivity check on the links used by CSPF. This is because some of these links may be associated with forwarding adjacencies, and therefore are unidirectional.

Februrary 2000

4.1. Label request object modification

The first field in the Label Request object is changed to contain the Link Media Type. When an LSR receives a Path message, it must allocate a label, lambda or channel that is compatible with the Link Media Type. The values for Link Media type as the same as for the two-octet media type values in IS-IS/OSPF Link Media Type TLV.

[illegible]

The information carried in the Label Object should be sufficient to control the connectivity matrix on an OXC, even when there are multiple fibers between a pair of OXCs, and (at least some, or even all of) these fibers can't be used for exchanging MPLS control information between these two OXCs, and these fibers are combined into bundled links.

That is, the Label Object should identify both a specific fiber (e.g., as an Interface Index), and a specific wavelength within that fiber, or a specific channel within that fiber. Which of the two is intended should be obvious to the receiver of the Resv message. The special value of 0xffff for the lambda or channel means the whole

domain, and start and end at the edges of that domain. The edge LSRs then use these LSPs for (a) creating forwarding adjacencies, and for (b) nesting of other LSPs that are established and terminated at LSRs beyond the domain and its edges.

Note that a 'normal' packet-switched LSP may only be nested in an FA-LSP of type FA-PSC; a SONET trail may only be nested in an FA-LSP of type FA-TDM; and an optical trail may only be nested in an FA-LSP of type FA-LSC.

Establishment/termination of the LSPs that span a domain could be triggered either by mechanisms outside of MPLS (e.g., via administrative control), or by mechanisms within MPLS (e.g., as a result of the LSR at the edge of a domain domain receiving LSP setup requests originated by some other LSRs beyond the domain and its edges. Procedures described in [Section 4.3.3](#) applied to both cases.

Procedures described in [Section 4.3.4](#) apply only to the latter case.

[4.3.3](#). Common procedures

For the purpose of processing the ERO in a Path message of an LSP that is to be tunneled over a forwarding adjacency, an LSR at the head-end of the FA-LSP views the LSR at the tail of that FA-LSP as adjacent (one IP hop away).

If the LSR at the tail of the FA-LSP is capable of packet switching, the Path message for the tunneled LSP can itself be tunneled over the FA-LSP. If the encapsulation on the carrier LSP can distinguish IP from MPLS, the Path message is sent as a plain IP packet. Otherwise, the Path message is sent with a label of 0, meaning "pop the label and treat as IP".

If the LSR at the tail of the FA-LSP is not capable of packet switching, the Path message is unicast over the control plane to the tail of the carrier LSP, without the Router Alert option.

The Resv message back to the head-end of the FA-LSP (PHOP) cannot be sent over the same FA-LSP as it is unidirectional. The Resv message can either take any LSP whose end-point is the head-end of the FA-LSP, or be unicast over the control plane to the head-end.

When an LSP is tunneled through an FA-LSP, the LSR at the head-end of the FA-LSP subtracts the LSP's bandwidth from the unreserved bandwidth of the forwarding adjacency.

4.3.4. Specific procedures

When an LSR receives a Path message, the LSR determines whether it is at the edge of a domain with respect to the ERO carried in the message. The LSR does this by looking up the link types of the previous hop and the next hop in its IGP database, and comparing them using the relation defined in [section 4.3.1](#). If the LSR is not at the edge of a domain, the procedures in this section do not apply.

If the LSR is at the edge of a domain, it must then determine the other edge of the domain with respect to the ERO, again using the IGP database. The LSR then extracts the strict hop subsequence from itself to the other end of the domain.

The LSR then compares the strict hop subsequence with all existing FA-LSPs originated by the LSR; if a match is found, and that FA-LSP

has enough unreserved bandwidth for the LSP being signaled, the LSR uses that FA-LSP as follows. The Path message for the original LSP is sent to the egress of the FA-LSP, not to the next hop along the FA-LSP's path. The PHOP in the message is the address of the LSR at the head-end of the FA-LSP. Before sending the Path message, the ERO in that message is adjusted by removing the subsequence of the ERO that lies in the FA-LSP, and replacing it with just the end point of the FA-LSP.

Otherwise (if no existing FA-LSP is found), the LSR sets up a new FA-LSP. That is, it initiates (using RSVP) a new LSP setup just for the FA-LSP.

After the LSR establishes the new FA-LSP, the LSR announces this LSP into IS-IS/OSPF as a forwarding adjacency.

The unreserved bandwidth of the forwarding adjacency is computed by subtracting the bandwidth of sessions pending the establishment of

the FA-LSP associated from the bandwidth of the FA-LSP.

An FA-LSP could be torn down by the LSR at the head-end of the FA-LSP as a matter of policy local to the LSR. It is expected that the FA-LSP would be torn down once there are no more LSPs carried by the FA-LSP. When the FA-LSP is torn down, the forwarding adjacency associated with the FA-LSP is no longer advertised into IS-IS/OSPF.

5. Bandwidth accounting on OXCs

As mentioned above, an IP link consists of a control channel and a set of fibers. Initially, for each fiber the administrator sets up how many lambdas could be reserved at each priority on that fiber, and the Link Media type(s) supported by the fiber. For an IP link, the unreserved bandwidth at each priority level is determined by adding up the bandwidths of all the lambdas of all the fibers in the IP link at that priority level. The Link Media Type of the IP link is the set union of the Link Media Types of each fiber in the link. Note that while the IP link contains the aggregated information, RSVP needs to track the unaggregated information, i.e., the available lambdas for each fiber.

Suppose, in response to a previously sent Path message, a Resv message is received at an ingress or intermediate OXC of an optical trail, along with the lambda and fiber to use, say $\langle L, F \rangle$. Say lambda L has bandwidth B and Link Media Type M. RSVP must remove L from the list of available lambdas for fiber F. Furthermore, the unreserved bandwidth at the holding priority of the optical trail is

reduced by B. Also, if after lambda L is removed, there are no more available lambdas with Link Media Type M in the fiber set that F belongs to, M is removed from the list of Link Media Types for the link containing F. Finally, the bandwidth accounting procedures for link bundles [[Kompella](#)] may also apply.

When an optical trail that was using lambda L of fiber F at an ingress or intermediate OXC is torn down, that OXC needs to mark L as available. The unreserved bandwidth (at the holding priority of the optical trail) of the IP link to which fiber F belongs is increased by the bandwidth of lambda L, and the Link Media Type of lambda L is

added to the list of Link Media Types for the IP link if not already present. Finally, the accounting procedures for link bundles may also apply.

6. Security Considerations

This document raises no new security issues for IS-IS, OSPF or RSVP.

7. Acknowledgements

We would like to thank Der-Hwa Gan, Jim Gibson, Robert Goguen, George Swallow, and Quaizar Vohra for their review and comments.

8. References

[Awduche-1] Awduche, D., Rekhter, Y., Drake, J., Coltun, R., "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Crossconnects", [draft-awduche-mpls-te-optical-01.txt](#)

[Awduche-2] Awduche, D., Berger, L., Gan, D-H., Li, T., Swallow, G., Srinivasan, V., "Extensions to RSVP for LSP Tunnels", [draft-ietf-mpls-rsvp-lsp-tunnel-04.txt](#)

[Basak] Basak, D., Awduche, D., Drake, J., Rekhter, Y., "Multi-protocol Lambda Switching: Issues in Combining MPLS Traffic Engineering Control With Optical Crossconnects", [draft-basak-mpls-oxc-issues-00.txt](#)

[Smit] Smit, H., Li, T., "IS-IS extensions for Traffic Engineering", [draft-ietf-isis-traffic-01.txt](#)

[Katz] Katz, D., Yeung, D., "Traffic Engineering Extensions to OSPF", [draft-katz-yeung-ospf-traffic-01.txt](#)

9. Author Information

Kireeti Kompella
Juniper Networks, Inc.
385 Ravendale Drive
Mountain View, CA 94043
email: kireeti@juniper.net

Daniel O. Awduche
UUNET (MCI Worldcom)
22001 Loudoun County Parkway
Ashburn, VA 20147
email: awduche@uu.net

Gisli Hjalmtysson
AT&T Labs - Research
180 Park Avenue
Florham Park, NJ 07932
email: gisli@research.att.com

Satoru Okamoto
NTT Network Innovation Laboratories
Room 807A, 1-1 Hikari-no-oka,
Yokosuka-shi, Kanagawa 239-0847 Japan
email: okamoto@exa.onlab.ntt.co.jp

Greg Bernstein
Ciena Corporation
10201 Bubb Road
Cupertino, CA 95014
email: GregB@ciena.com

Near Margalit
New Access Communications
71 Vista Montana
San Jose CA 95134
email: margalit@new-access.com

Yakov Rekhter
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
email: yakov@cisco.com

Alan Hannan
Global Crossing
141 Caspian Court
Sunnyvale, CA 94087
email: alan@gblx.net

Joe Lawrence
Level 3 Communications
1025 Eldorado Boulevard
Broomfield CO, 80021
email: joe.lawrence@level3.com

Debashis Basak
Marconi
1000 FORE Drive
Warrendale, PA 15086-7502
email: dbasak@fore.com

John Drake
Chromisys
1012 Stuart Drive
Sunnyvale, CA 94086
email:jdrake@chromisys.com

Ed Stern
Sirocco Systems
95 Barnes Rd.
Wallingford, CT 06492
email:ed.stern@siroccosystems.co