Francois Le Faucheur
Thomas D. Nadeau
Cisco Systems, Inc.

Angela Chiu
AT&T

William Townsend
Tenor Networks

Darek Skalecki
Nortel Networks

Martin Tatham
BT

IETF Internet Draft
Expires: May, 2001
Document: draft-ietf-mpls-diff-te-reqts-00.txt November, 2000

Requirements for support of Diff-Serv-aware MPLS Traffic Engineering

Status of this Memo

Abstract

This document defines the requirements for support of Diff-Serv-aware MPLS Traffic Engineering on a per-Class-Type basis, as discussed in the Traffic Engineering Working Group Framework

document [TEWG-FW].

Companion documents [DIFF-TE-EXT] [DIFF-TE-OSPF] [DIFF-TE-ISIS] respectively propose actual extensions to RSVP and CR-LDP, to OSPF and to ISIS, in order to meet those requirements.

1. Introduction

As Diff-Serv becomes prominent in providing scalable multi-class of services in IP networks, performing traffic engineering at a per-class level instead of an aggregated level is needed to further enhance networks in performance and efficiency. By mapping a traffic trunk in a given class on a separate LSP, it allows the traffic trunk to utilize resources available on both shortest path(s) and non-shortest paths and follow paths that meet constraints which are specific to the given class. It also allows each class to select the proper protection/restoration mechanism(s) that satisfy its survivability requirements in a cost effective manner.

Besides the set of parameters defined for the general aggregate TE [TE-REQ], a new set of per-class parameters needs to be provided at each LSR interface and propagated via extensions to the IGP (ISIS/OSPF) [TEWG-FW]. Furthermore, the per-class parameters can be aggregated into per-Class-Type parameters. The main motivation for grouping a set of classes into a Class-Type is to improve the scalability of the IGP link state advertisements by propagating information on a per-Class-Type basis instead of on a per-class basis. This approach also has the benefit of allowing better bandwidth sharing between classes in the same Class-Type.

A Class-Type [TEWG-FW] is defined as a set of classes that satisfy the following two conditions:

1) Classes in the same Class-Type possess common aggregate maximum and minimum bandwidth requirements to guarantee the required performance level.

2) There is no maximum or minimum bandwidth requirement to be enforced at the level of an individual class within the Class-Type. One can still implement some "priority" policies for classes within the same Class-Type in terms of accessing the Class-Type bandwidth (e.g. via the use of preemption priorities).

An example of Class-Type comprising multiple Diff-Serv classes is a low-loss Class-Type that includes both AF1-based and AF2-based Ordering Aggregates.

Note that with per Class-Type TE, Constraint-Based Routing is performed with bandwidth constraints on a per Class-Type basis but LSPs may carry a single Diff-Serv class (Ordered Aggregate) with Diff-Serv scheduling (i.e. PHB) performed separately for each class.

In this document, we will only discuss "per Class-Type TE" because "per Class TE" can be viewed as a special case of per Class-Type TE (where each Class-Type is degenerated into a single Diff-Serv class).

This document focuses on intra-domain operations. Inter-domain operations is for further study.

The following sections detail the requirements on OSPF/ISIS, RSVP/CR-LDP, Constraint Based Routing, MPLS MIBs, Diff-Serv Scheduling and Traffic Mapping for support of MPLS Traffic Engineering on a per-Class-Type basis.

2. Requirements for ISIS/OSPF Extensions

[OSPF-TE] and [ISIS-TE] define extensions to OSPF and ISIS for support of (aggregate) MPLS Traffic Engineering. In this section we define the requirements on OSPF and ISIS for support of Diff-Serv Traffic Engineering on a per-Class-Type basis. These requirements are expected to require further extensions to OSPF and ISIS. Such extensions are proposed in [DIFF-TE-OSPF] and [DIFFF-TE-ISIS].

Given that there are hard limits imposed by ISIS/OSPF TLVs, the TLV space must be used frugally. An additional concern is that the amount of information advertised by the IGP directly affects the scalability of the solution. These considerations strongly influence the requirements defined in this section.

As pointed out in [TEWG-FW], the IGP needs to advertise separate "TE information" for each Class-Type. We focus now on detailing what this "TE information" should be.

For Constraint Based Routing to be able to compute paths which satisfy different bandwidth constraints for each Class-Type, the IGP needs to advertise different "Unreserved Bandwidth" information for each Class-Type.

Moreover, we propose that the preemption attribute defined in [TE-REQ] be retained for all Class-Types. We also propose that the preemption attributes of setup priority and holding priority retain existing semantics, and in particular the preemption should work across class-types (i.e. independently of class-types), rather than having preemption levels operating only within each class-type. Thus, the IGP needs to advertise "Unreserved Bandwidth" at each preemption level for each Class-Type. However, we observe that within a Class-Type, the "Unreserved Bandwidth" value is often identical for multiple preemption levels. Firstly, many practical MPLS Traffic Engineering deployments will use less than the maximum 8 possible levels of preemption. In such cases, the Unreserved Bandwidth corresponding to a preemption level which is not used will

always be equal to the Unreserved Bandwidth for the preemption level immediately above (numerically lower). For example, if only preemption levels 0,1 and 4 are used in a network, then the "Unreserved Bandwidth" for preemption levels 2 and 3 are identical to the "Unreserved Bandwidth" for preemption 1, and the "Unreserved Bandwidth" for preemption levels 5, 6 and 7 are identical to the "Unreserved Bandwidth" for preemption 4. Secondly, even when all 8 preemption levels are used in the network, whenever there is no TE Tunnels actually setup for a given preemption level (for that Class-Type) on a given link, the Unreserved Bandwidth corresponding to that preemption level will again be equal to the Unreserved Bandwidth for the preemption level immediately above (numerically lower).

Thus we propose that the IGP extension encoding includes a compression scheme which can efficiently reduce the length of the "Unreserved Bandwidth" advertisement (for each Class-Type) whenever there are duplicate values across multiple preemption levels.

For the bandwidth constraints to be effectively different for each Class-Type, LSRs need to allow configuration for every link of a "Maximum Reservable Bandwidth" for each Class-Type. Clearly, the "Unreserved Bandwidth" advertised for each Class-Type takes into account the "Maximum Reservable Bandwidth" configured for the corresponding Class-Type. Consequently, Constraint Based Routing can compute paths for the different Class-Types without receiving the "Maximum Reservable Bandwidth" for each Class-Type from the IGP. Thus we feel that the IGP need not advertise the Maximum Reservable Bandwidth for each Class-Type. We note that the Maximum Reservable Bandwidth for each Class-Type could have been used by Constraint Based Routing to enhance route computation in some situations (e.g. as a tie breaker), but we feel this does not justify the extra overhead in IGP advertisement.

Current IGP extensions for (aggregate) TE [OSPF-TE][ISIS-TE] specify advertisement of the Maximum Reservable Bandwidth for (aggregate) TE. Note that this document does not propose that this be changed.

Other TE attributes already advertised by the IGP (i.e. administrative group/color, IPv4 interface and neighbor addresses, Maximum Link Bandwidth, TE Metric) need not be advertised per Class-Type as those will be applicable to all Class-Types. We note that a separate TE Metric per Class-Type could be defined in the future if required, for instance to allow Constraint Based Routing to take account of link propagation delay for LSPs from a Class-Type with strict delay requirements.

We propose to begin by allowing a total of 4 Class-Types (i.e., 3 beyond the existing one aka. Class-Type 0). This is expected to be sufficient for practical deployments in the foreseeable future. As an example, a total of three Class-Types already allow support of separate bandwidth control for Real-Time, Low-Loss and Best Effort,

while allowing multiple classes within each Class-Type (e.g. AF1 and AF2 flavors of "Low-Loss"). More Class-Types could be defined in the future if required.

Where a Class-Type is not effectively used in a network, it is recommended that the corresponding sub-TLV is not included in the IS reachability TLV. Therefore, the Class-Types for which "Unreserved Bandwidth" is to be advertised in the IGP should be configurable and the IGP must allow advertisement of "Unreserved Bandwidth" for any subset of the Class-Types.

Implementations of Diff-Serv Traffic Engineering in compliance with this specification MUST support at least a total of 2 Class-Types and MAY support a total of 3 or 4 Class-Types.

2.1. Bandwidth Reservation Scheme

This section discusses the algorithm for computing the "Unreserved Bandwidth" for each Class-Type as a function of:

- the configurable parameters controlling how link bandwidth can be reserved by each Class-Type (in particular in situations where Class-Types compete for bandwidth reservation) such as "Max Reservable Bandwidth for the Class Type",
- the already established LSPs of all Class-Types.

This "Unreserved Bandwidth" for each Class-Type is computed for advertisement in the IGP (and therefore used as the per Class-Type bandwidth constraint for Constraint Based Routing). The same algorithm for computation of "Unreserved Bandwidth" must also be used for admission control of Diff-Serv-aware Traffic Engineering LSPs at establishment time through RSVP or CR-LDP signalling; otherwise persistent deadlock situations could occur whereby Constraint Based Routing believes that a given LSP for a given Class-Type can be routed through a link but local admission control rejects it.

It is desirable to be able to avoid under-utilizing aggregate resource. To achieve this, it is necessary to allow the sum of the configurable Maximum Reservable Bandwidth of all Class-Types to be larger than a configurable Maximum Reservable Aggregate Bandwidth (i.e. aggregate across all Class-Types). At the same time, it is desirable to be able to avoid over-utilizing the aggregate resource. To achieve this, it is necessary to be able to enforce this Maximum Reservable Aggregate Bandwidth; in other words it is necessary to ensure that the sum of all LSPs across all Class-Types never exceeds the Maximum Reservable Aggregate Bandwidth.

For example, a 10Gb/s link may be configured to allow:

- Class-Type 0 (eg: BE) to reserve up to 9 Gb/s
- Class-Type 1 (eg: real time including EF) to reserve up to 5 Gb/s

- Class-Type 2 (eg: low loss including AF1 and AF2) to reserve up to 8 Gb/s

and at the same may be configured to allow:

- on an aggregate basis, the sum of all Class-Types to reserve up to 10 Gb/s.

Therefore, a path computed by the Constraint Based Routing for an LSP of Class-Type N must ensure that this LSP fits within the remaining Class-Type N bandwidth AND that this LSP fits within the remaining Aggregate bandwidth.

One way to achieve this, would be:
- for each Class-Type, that IGP uses the "Unreserved Bandwidth for Class-Type N" to advertise the Class-Type N bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Bandwidth for Class-Type N and the bandwidth reserved by existing Class-Type N LSPs),
- in addition, that IGP separately advertises the "Unreserved

Aggregate Bandwidth" (i.e. the difference between the Maximum Reservable Aggregate Bandwidth and the bandwidth reserved by existing LSPs of all Class-Types)

- have Constraint Based Routing ensure that a new Class-Type N LSP fits both in the received "Unreserved Bandwidth for Class-Type N" and in the "Unreserved Aggregate Bandwidth".

Such an approach has the drawbacks that it would require that N+1 "unreserved bandwidth" information be advertised by the IGP when N Class-Types are supported, and that it requires the node performing Constraint Based Routing to meet a double bandwidth constraints.

Instead we propose that:

- for each Class-Type, that IGP uses the "Unreserved Bandwidth for Class-Type N" to directly advertise the amount of bandwidth that is effectively useable by Class-Type N. This is computed as the smaller of these two values:
    - o The Class-Type N bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Bandwidth for Class-Type N and the bandwidth reserved by existing Class-Type N LSPs).
    - o The aggregate bandwidth currently unreserved (i.e. the difference between the Maximum Reservable Aggregate Bandwidth and the bandwidth reserved by existing LSPs of all Class-Types).
- have Constraint Based Routing ensure that a new Class-Type N LSP simply fits in the received "Unreserved Bandwidth for Class-Type N".

Such an approach only requires that N "unreserved bandwidth" information be advertised by the IGP when N Class-Types are supported, and only requires that the node performing Constraint Based Routing meets a single bandwidth constraint.

It may be desirable to prevent a Class-Type from being starved by others. In the example given above where we defined three Class-Types, it may be useful to be able to always ensure that some amount of Class-Type 0 LSPs can be routed over that link (i.e. to prevent Class-Type 1 LSPs and Class-Type 2 LSPs from reserving up to 100% of the maximum reservable aggregate bandwidth which would result in Class-Type 0 LSPs not having any access to the capacity of that link). Such capability might require the ability from the IGP to advertise an optional "minimum reservable bandwidth" per Class-Type. This is not seen as an immediate requirement but could be defined in the future if required.

[Editor's Note: We are considering a potential enhancement to the Bandwidth Reservation scheme presented above and therefore to the calculation of advertised `unreserved bandwidth' per Class-Type.

This potential enhancement would take into account a Minimum Reservable Bandwidth in addition to the Maximum Reservable Bandwidth and the Maximum Aggregate Reservable bandwidth and could take account of bandwidth `borrowing' between classes. Thus traffic trunks of one Class-Type (say CT-2) might initially reserve all the bandwidth of the link, and implicitly that Class-Type would have borrowed unreserved capacity from other Class-Types (say CT-1, CT-0). However, the LSR would continue to advertise `unreserved bandwidth' for CT-1 and CT-0 based on the maximum bandwidth entitlement of each Class-Type. A reservation may then be received within CT-1 or CT-0, which would preempt an existing reservation in CT-2. This form of preemption between Class-Types would only operate within existing preemption priority levels, and existing rules for preemption across priority levels would still be followed. The potential goals of this enhancement would be :

- to allow enforcement of a Minimum Reservable Bandwidth per Class-Type in addition to the Maximum Reservable Bandwidth while still advertising a single "Unreserved Bandwidth" per Class-Type (for each preemption levels),
- to match the bandwidth "sharing" properties (across Class-Types) of common work-conserving Diff-Serv schedulers (e.g. Weighted Fair Queuing). This could ensure that, even in the case of statically configured Diff-Serv packet scheduler, there is always consistency between :
    * the bandwidth reserved by each Class-Type for Constraint Based Routing purposes, and
    * the bandwidth actually granted to the corresponding Class-Type by the Diff-Serv Packet Scheduler

As discussed in section 6 below, the Bandwidth Allocation Scheme currently proposed requires dynamic adjustement of the Diff-Serv packet scheduler.]

3. Requirements for RSVP/CR-LDP Extensions

[RSVP-TE] and [CR-LDP] define extensions to RSVP and LDP for support of (aggregate) MPLS Traffic Engineering. [DIFF-MPLS] defines the extensions to RSVP and LDP for support of Diff-Serv over MPLS. In this section we define the requirements on RSVP and CR-LDP for support of Diff-Serv Traffic Engineering on a per-Class-Type basis.

These requirements are expected to require further extensions to RSVP and CR-LDP. Such extensions are proposed in [DIFF-TE-EXT].

In order for an LSR to perform resource availability checking for an LSP that belongs to a certain Class-Type, the LSR needs to be made aware through RSVP/CR-LDP signaling of the Class-Type associated with the LSP.

To that end, we propose that RSVP/CR-LDP be extended to be able to signal the Class-Type.

We identify the following backward compatibility requirements for the RSVP/CR-LDP extensions:

- operations in heterogeneous environments need to be supported for smooth migration, where some LSRs are Diff-Serv-TE-capable (as defined in this specification) while some other LSRs are not Diff-Serv-TE-capable (i.e. support (aggregate) TE only)
- in such heterogeneous environments, the solution needs to allow establishment of Class-Type 0 LSPs across paths combining Diff-Serv-TE-capable LSRs and non-Diff-Serv-TE-capable LSRs
- in heterogeneous environments, the solution needs to ensure that a non-Diff-Serv-TE-capable LSR would reject establishment of a Class-Type N (N=1,2,3) LSP.

More generally we identify the following backward compatibility requirements for the RSVP/CR-LDP extensions:

- operations in heterogeneous environments need to be supported for smooth migration, where some Diff-Serv-TE-capable LSRs (as defined in this specification) support N Class-Types while other Diff-Serv-TE-capable LSRs support M Class-Types with M>N.
- in such heterogeneous environments, the solution needs to allow establishment of Class-Type 0,.., N-1 LSPs across paths combining both types of LSRs
- in such heterogeneous environments, the solution needs to ensure that a Diff-Serv-TE-capable LSR supporting only N Class-Types would reject establishment of a Class-Type X LSP if N<=X<=M-1.

The admission control algorithm implemented for LSP establishment must locally maintain different variables which keep track of the currently unreserved bandwidth for each Class-Type. These unreserved bandwidth variables must be updated in accordance with the Bandwidth Reservation Scheme discussed in the previous section.

4. Requirements for Constraint Based Routing Extensions

In order for Constraint Based Routing to support Diff-Serv TE on a per-Class-Type basis, the Constraint Based Routing algorithm need to be capable of taking into account the "Unreserved Bandwidth for Class-Type N" when computing a path for a Class-Type N LSP.

The Constraint Based Routing may also take into account different metric for different Class-Types. As an example, when computing a path for a non-real-time Class-Type, the Constraint Based Routing may use a metric similar to the one currently used by IGP SPF Routing for Best Effort traffic, while when computing a path for a real-time Class-Type, the Constraint Based Routing may use a metric reflecting the link propagation delay.

5. Requirements for MIB Extensions

In order for an LSR to support the configuration and monitoring of Diff-Serv Traffic Engineering certain enhancements to some of the existing MPLS Management Information Bases (MIBs) will be required. [LSRMIB] defines the MPLS Label Switch Router MIB (LSR MIB) which contains objects useful for the management and configuration of MPLS LSPs. [TE MIB] defines the MPLS Traffic Engineering MIB (TE MIB) which contains objects useful for the management and configuration of MPLS Traffic Engineered Tunnels.

In particular, the MIB extensions need to:

- track for each MPLS interface, the Maximum Reservable Bandwidth configured for each Class-Type.
- track for each MPLS interface, the Maximum Reservable Aggregate Bandwidth configured.
- track for each LSP, the Class-Type associated with the LSP. On the Head-End LSRs, the Class-Type is configured as part of the tunnel configuration. On other LSRs, the Class-Type is associated with the LSP at establishment time based on signaled information.

Additional details of these changes will be provided in forthcoming versions of this draft. It is the authors' intent to transfer these MIB requirements to future versions of the MPLS TE and the MPLS LSR MIBs. It is not the intent of this document to define the SMI required for the MIB enhancements; rather, it is to flesh out and define the details of these changes in the context of this document.

6. Diff-Serv Scheduling Requirements

Diff-Serv-aware Traffic Engineering is expected to be deployed in conjunction with Diff-Serv buffer management and differential packet scheduling. Per-class-type performance would be controlled on longer timescales by Diff-Serv-aware Traffic Engineering; this would be complemented by Diff-Serv buffer management and differential packet scheduling mechanisms controlling the per-class performance on packet scheduling timescales.

If the Diff-Serv packet scheduling mechanisms are configured statically, it is possible, however, that the allocation of resources by Diff-Serv-aware Traffic Engineering for reservation purposes will not always correspond with the resources allocated in the packet scheduler. As an example, consider the following 10Mbps link:

	Max. Res'ble Bandwidth	Admitted Load	Pkt Scheduler Bandwidth share
Class-Type 0	10 Mbps	1 Mbps	3 Mbps
Class-type 1	3 Mbps	3 Mbps	3 Mbps
Class-type 2	10 Mbps	6 Mbps	4 Mbps

In the above example, the Maximum Reservable Bandwidth for Class-Type 2 (e.g. AF) has been set above the corresponding scheduler share of 4 Mbps to prevent under-utilisation of the link, in case there was little traffic in Class-Types 0 (e.g. BE) and 1 (e.g. EF). Under actual load, let's assume that 6 Mbps of Class-Type 2 traffic was admitted on the link, possibly due to Class-type 2 LSPs being given a high preemption priority. However, the packet scheduler had been configured to give the classes comprising Class-type 2 an overall share of 4Mbps, reflecting the expected load and taking account of the performance targets of Class-type 2 traffic. Even though, for most link schedulers, resources would be dynamically redirected from Class-Type 0 to Class-Type 2, this would leave Class-type 2 traffic at risk of seeing poor performance, under conditions of link congestion. Assume for instance that :

- there is indeed 3 Mbps worth of actual traffic sent onto the admitted Class-Type 1 LSPs,
- there is indeed 6 Mbps worth of actual traffic sent onto the admitted Class-Type 2 LSPs,
- although only 1 Mbps worth of Class-Type 0 (eg. Best Effort traffic) has been admitted by Diff-Serv-aware TE, 3 Mbps of Class-Type 0 traffic is actually sent onto the Class-Type 0 LSPs.

In that case, the packet scheduler will :

- serve the 3 Mbps of Class-Type 0 traffic
- serve the 3 Mbps of Class-Type 1 traffic
- will only be able to provide 4 Mbps of service rate for the 6 Mbps of Class-Type 2 traffic.

This example reflects the fact that the Bandwidth allocation Scheme described in section 2.1 above may not under all conditions align with the resource allocation approach taken by a statically configured local link packet scheduler.

In general, therefore, the requirements on Diff-Serv packet scheduling are that:

1) An LSR should be capable of dynamically adjusting resource allocation to Classes based on per-class LSP resource requests.

2) The dynamic adjustment should take account of the performance requirements of the classes. This may be achieved using, for example, per-Class maximum allocation multipliers or overbooking factors to adjust scheduling weights.

7. Traffic Mapping Requirements

This section describes the requirement for an LSR which is the Head-end of Diff-Serv-aware Traffic Engineering LSPs to map incoming traffic onto these LSPs.
Each Diff-Serv-aware Traffic Engineering LSP has the following attributes:

- it can transport one (or a set of) Diff-Serv class(es) (Ordered Aggregate) in accordance with [DIFF-MPLS]
- it has been constraint based routed based on the Class-Type which comprises this (or these) class(es), in accordance with the previous sections of this document.

Mapping of incoming traffic onto Diff-Serv-aware Traffic Engineering LSPs is to be performed in accordance with [DIFF-MPLS] so that only packets that belong to the (set of) Behavior Aggregate(s) transported over a given Diff-Serv-aware TE LSP should be mapped to that LSP. In particular, where the Head-end LSR is also the MPLS Edge LSR, determination of the Behavior Aggregate (and thus determination of the egress Diff-Serv-aware TE LSP) is based on the Diffserv Codepoint (DSCP) in the packet header.

8. Security Considerations

The solution developed to address the requirements defined in this document must address security aspects.

9. Acknowledgments

This document has benefited from discussions with Carol Iturralde and Rob Goguen.

References

[TE-REQ] Awduche et al, Requirements for Traffic Engineering over MPLS, RFC2702, September 1999.

[TEWG-FW] Awduche et al, A Framework for Internet Traffic Engineering, draft-ietf-tewg-framework-02.txt, July 2000.

[DIFF-TE-EXT] Le Faucheur et al, Extensions to RSVP and CR-LDP for support of Diff-Serv-aware MPLS Traffic Engineering, draft-ietf-mpls-diff-te-ext-00.txt, November 2000.

[DIFF-TE-OSPF] Le Faucheur et al, Extension to OSPF for support of Diff-Serv-aware MPLS Traffic Engineering, draft-lefaucheur-diff-te-ospf-01.txt, November 2000.

[DIFF-TE-ISIS] Le Faucheur et al, Extension to ISIS for support of Diff-Serv-aware MPLS Traffic Engineering, draft-lefaucheur-diff-te-isis-01.txt, November 2000.

[OSPF-TE] Katz, Yeung, Traffic Engineering Extensions to OSPF, draft-katz-yeung-ospf-traffic-03.txt, September 2000.

[ISIS-TE] Smit, Li, IS-IS extensions for Traffic Engineering, draft-ietf-isis-traffic-02.txt, September 2000.

[RSVP-TE] Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", draft-ietf-mpls-rsvp-lsp-tunnel-07.txt, August 2000.

[DIFF-MPLS] Le Faucheur et al, "MPLS Support of Diff-Serv", draft-ietf-mpls-diff-ext-07.txt, August 2000

[LDP] Andersson et al., "LDP Specification", draft-ietf-mpls-ldp-11.txt, August 2000

[CR-LDP] Jamoussi et al., "Constraint-Based LSP Setup using LDP", draft-ietf-mpls-cr-ldp-04.txt, July 2000

[TEMIB] Srinivansan, C., and A. Viswanathan, "MPLS Traffic Engineering Management Information Base Using SMIv2", draft-ietf-mpls-te-mib-04.txt, July 2000.

[LSRMIB] Srinivansan, C., Viswanathan, A., and T. Nadeau "MPLS Label Switch Router Management Information Base Using SMIv2", draft-ietf-mpls-lsr-mib-06.txt, July 2000.

Authors' Address:

Francois Le Faucheur
Cisco Systems, Inc.
Petra B - Les Lucioles - 291, rue Albert Caquot - 06560 Valbonne - France
Phone: +33 4 92 96 75 64
Email: flefauch@cisco.com

Angela Chiu
AT&T Labs
200 Laurel Ave. Rm A5-1F06
Middletown, NJ 07748, USA

Tel: 1-(732) 420-9057
Email: alchiu@att.com

William Townsend
Tenor Networks
100 Nagog Park
Acton, MA 01720
Phone: +1-978-264-4900
Email: btownsend@tenornetworks.com

Thomas D. Nadeau
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
Phone: +1-978-244-3051
Email: tnadeau@cisco.com

Darek Skalecki
Nortel Networks
3500 Carling Ave,
Nepean K2H 8E9
Phone: +1-613-765-2252
Email: dareks@nortelnetworks.com

Martin Tatham
BT
Adastral Park,
Martlesham Heath,

Ipswich IP5 3RE
UK
Phone: +44-1473-606349
Email: martin.tatham@bt.com