

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

Martin Tatham
BT

Thomas Telkamp
David Cooper
Global Crossing

Jim Boyle
PDNets

Waisum Lai, Co-editor
Luyuan Fang
Jerry Ash
AT&T

Pete Hicks
Core Express

Angela Chiu
Celion Networks

William Townsend
Tenor Networks

Darek Skalecki
Nortel Networks

IETF Internet Draft

Expires: December, 2002

Document: [draft-ietf-tewg-diff-te-reqts-05.txt](#)

June 2002

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

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."

Le Faucheur, et. al

1

Requirements for Diff-Serv-aware Traffic Engineering June 2002

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>.

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

Abstract

This document presents the Service Provider requirements for support of Diff-Serv aware MPLS Traffic Engineering (DS-TE).

Its objective is to provide guidance for the definition, selection and specification of a technical solution addressing these requirements. Specification for this solution itself is outside the scope of this document.

A problem statement is first provided. Then, the document describes example applications scenarios identified by Service Providers where existing MPLS Traffic Engineering mechanisms fall short and Diff-Serv-aware Traffic Engineering is required. The detailed requirements that need to be addressed by the technical solution are also reviewed. Finally, the document identifies the evaluation criteria that should be considered for selection and definition of the technical solution.

1. Introduction

1.1. Problem Statement

Diff-Serv is becoming prominent in providing scalable network designs supporting multiple classes of services.

In some Diff-Serv networks where optimization of transmission resources on a network-wide basis is not sought, MPLS Traffic Engineering (TE) mechanisms may simply not be used.

In other networks, where optimization of transmission resources is sought, Diff-Serv mechanisms [[DIFF-MPLS](#)] need to be complemented by existing MPLS Traffic Engineering mechanisms [[TE-REQ](#)] [[ISIS-TE](#)] [[OSPF-TE](#)] [[RSVP-TE](#)] [[CR-LDP](#)] which operate on an aggregate basis across all classes of service. In this case, Diff-Serv and MPLS TE both provide their respective benefits.

Where fine-grained optimization of transmission resources is sought, it is necessary to perform traffic engineering at a per-class level instead of an aggregate level, in order to further enhance networks in performance and efficiency as discussed in [[TEWG-FW](#)]. By mapping the traffic from a given class of service on a separate LSP, it allows this traffic to utilize resources available to the given class on both shortest path and non-shortest paths and follow paths

Le Faucheur et. al

2

Requirements for Diff-Serv-aware Traffic Engineering June 2002

that meet engineering constraints which are specific to the given class. This is what we refer to as "Diff-Serv-aware Traffic Engineering (DS-TE)".

This document focuses exclusively on the specific environments which would benefit from DS-TE. Some examples include:

- networks where bandwidth is scarce (e.g. transcontinental networks)
- networks with significant amounts of delay-sensitive traffic
- networks where the relative proportion of traffic across classes of service is not uniform

This document focuses on intra-domain operation. Inter-domain operation is not considered.

1.2. Definitions

For the convenience of the reader, relevant Diffserv ([[DIFF-ARCH](#)], [[DIFF-NEW](#)]) definitions are repeated herein.

Behavior Aggregate (BA): a collection of packets with the same (Diff-Serv) codepoint crossing a link in a particular direction.

Per-Hop-Behavior (PHB): the externally observable forwarding behavior applied at a DS-compliant node to a Diff-Serv behavior aggregate.

PHB Scheduling Class (PSC): A PHB group for which a common constraint is that ordering of at least those packets belonging to the same microflow must be preserved.

Ordered Aggregate (OA): a set of BAs that share an ordering constraint. The set of PHBs that are applied to this set of Behavior Aggregates constitutes a PHB scheduling class.

We also repeat the following definition from [\[TE-REQ\]](#):

Traffic Trunk: an aggregation of traffic flows of the same class which are placed inside a Label Switched Path.

In the context of the present document, "flows of the same class" is to be interpreted as "flows from the same Forwarding Equivalence Class which are to be treated equivalently from the DS-TE perspective".

[1.3.](#) Mapping of traffic to LSPs

A network may have multiple classes of traffic it wishes to service. Recalling from [\[DIFF-MPLS\]](#), there are several options on how traffic from multiple OAs of a given Forwarding Equivalence Class can be

split into Traffic Trunks for mapping onto LSPs when running MPLS Traffic Engineering.

One option is to not split traffic from the different OAs so that each Traffic Trunk combines traffic from all OAs. This option is typically used when aggregate traffic engineering is deployed using current MPLS TE mechanisms. In that case, traffic from all OAs is routed collectively according to a single shared set of constraints and will follow the same path. Note that the LSP transporting such a Traffic Trunk is, by definition, an E-LSP as defined in [\[DIFF-MPLS\]](#).

Another option is to split traffic from the different OAs into multiple Traffic Trunks. In other words, traffic from a given node to another given node, is split into multiple Traffic Trunks which are transported over separate LSPs, which can potentially follow a different path through the network. DS-TE precisely takes advantage of this fact and indeed computes a separate path for each LSP. In so doing, DS-TE can take into account the specific requirements of the Traffic Trunk transported on each LSP (e.g. bandwidth requirement, preemption priority). Moreover DS-TE can take into account specific engineering constraints to be enforced for these sets of Traffic Trunks (e.g. limit all Traffic Trunks transporting a given set of OAs to x% of link capacity). In brief, DS-TE achieves per LSP constraint based routing with paths that tightly match the specific objectives of the traffic forming the corresponding Traffic Trunk.

2. Application Scenarios

2.1. Scenario 1: Limiting Proportion of Classes on a Link

An IP/MPLS network may need to carry a significant amount of VoIP traffic compared to its link capacity. For example, 10,000 uncompressed calls at 20ms packetization result in about 1Gbps of IP traffic, which is significant on an OC-48c based network. In case of topology changes such as link/node failure, VoIP traffic levels can even approach the full bandwidth on certain links.

For delay/jitter reasons it is undesirable to carry more than a certain percentage of VoIP traffic on any link. The rest of the available link bandwidth can be used to route other classes corresponding to delay/jitter insensitive traffic (e.g. Best Effort Internet traffic). The exact determination of this "certain" percentage is outside the scope of this requirements draft.

During normal operations, the VoIP traffic should be able to preempt other classes of traffic (if these other classes are designated as preemptable and they have lower preemption priority), so that it will be able to use the shortest available path, only constrained by the maximum defined link utilization ratio/percentage of the VoIP class.

Existing TE mechanisms only allow constraint based routing of traffic based on a single bandwidth constraint common to all classes, which does not satisfy the needs described here. This leads to the requirement for DS-TE to be able to enforce a different bandwidth constraint for different classes of traffic. In the above example, the bandwidth constraint to be enforced for VoIP traffic may be the "certain" percentage of each link capacity, while the bandwidth constraint to be enforced for the rest of the classes might have their own constraints or have access to the rest of the link capacity.

2.2. Scenario 2: Maintain relative proportion of traffic classes

Suppose an IP/MPLS network supports 3 classes of traffic. The network administrator wants to perform Traffic Engineering to distribute the traffic load. Assume also that proportion across traffic classes varies significantly depending on the source/destination POPs.

With existing TE mechanisms, the proportion of traffic from each class on a given link will vary depending on multiple factors including:

- in which order the different TE-LSPs are routed
- the preemption priority associated with the different TE-LSPs
- link/node failure situations

This may make it difficult or impossible for the network administrator to configure the Diff-Serv PHBs (e.g. queue bandwidth) to ensure that each traffic class gets the appropriate treatment. This leads again to the requirement for DS-TE to be able to enforce a different bandwidth constraint for different classes of traffic. This could be used to ensure that, regardless of the order in which tunnels are routed, regardless of their preemption priority and regardless of the failure situation, the amount of traffic of each class routed over a link matches the Diff-Serv scheduler configuration on that link for the corresponding class (e.g. queue bandwidth).

As an illustration of how DS-TE would address this scenario, the network administrator may configure the service rate of Diff-Serv queues to (45%,35%,20%) for classes (1,2,3) respectively. The administrator would then split the traffic into separate Traffic Trunks for each class and associate a bandwidth to each LSP transporting those Traffic Trunks. The network administrator may also want to configure preemption priorities of each LSP in order to give highest restoration priority to the highest priority class and medium priority to the medium class. Then DS-TE could ensure that after a failure, class 1 traffic would be rerouted with first access

at link capacity but without exceeding its service rate of 45% of the link bandwidth. Class 2 traffic would be rerouted with second access at the link capacity but without exceeding its allotment. Note that where class 3 is the Best-Effort service, the requirement

Requirements for Diff-Serv-aware Traffic Engineering June 2002

on DS-TE may be to ensure that the total amount of traffic routed across all classes does not exceed the total link capacity of 100% (as opposed to separately limiting the amount of Best Effort traffic to 20 even if there was little class 1 and class 2 traffic).

In this scenario, DS-TE allowed for the maintenance of a more steady distribution of classes, even during rerouting. This relied on the required capability of DS-TE to adjust the amount of traffic of each class routed on a link based on the configuration of the scheduler and the amount of bandwidth available for each class.

Alternatively, some network administrators may want to solve the problem by having the scheduler dynamically adjusted based on the amount of bandwidth of the LSPs admitted for each class. This is an additional requirement on DS-TE.

2.3. Scenario 3: Guaranteed Bandwidth Services

In addition to the Best effort service, an IP/MPLS network operator may desire to offer a point-to-point "guaranteed bandwidth" service whereby the provider pledges to provide a given level of performance (bandwidth/delay/loss...) end-to-end through its network from an ingress port to and egress port. The goal is to ensure all "guaranteed" traffic within a subscribed traffic contract, will be delivered within stated tolerances.

One approach for deploying such "guaranteed" service involves:

- dedicating a Diff-Serv PHB (or a Diff-Serv PSC as defined in [\[DIFF-NEW\]](#)) to the "guaranteed" traffic
- policing guaranteed traffic on ingress against the traffic contract and marking the "guaranteed" packets with the corresponding DSCP/EXP value

Where very high level of performance is targeted for the

"guaranteed" service, it may be necessary to ensure that the amount of "guaranteed" traffic remains below a given percentage of link capacity on every link. Where the proportion of "guaranteed" traffic is high, constraint based routing can be used to enforce such a constraint.

However, the network operator may also want to simultaneously perform Traffic Engineering of the rest of the traffic (i.e. non-guaranteed traffic) which would require that constraint based routing is also capable of enforcing a different bandwidth constraint, which would be less stringent than the one for guaranteed traffic.

Again, this combination of requirements can not be addressed with existing TE mechanisms. DS-TE mechanisms allowing enforcement of a different bandwidth constraint for guaranteed traffic and for non-guaranteed traffic are required.

3. Detailed Requirements for DS-TE

This section specifies the functionality that the above scenarios require out of DS-TE implementations. Actual technical protocol mechanisms and procedures to achieve such functionality are outside the scope of this document.

3.1. DS-TE Compatibility

While DS-TE is required in a number of situations such as the ones described above, it is important to keep in mind that using DS-TE may impact scalability (as discussed later in this document) and operational practices. DS-TE should only be used when existing TE mechanisms combined with Diff-Serv cannot address the network design requirements. Many network operators may choose to not use DS-TE, or to only use it in a limited scope within their network.

Thus, the DS-TE solution must be developed in such a way that:

- (i) it raises no interoperability issues with existing deployed

- TE mechanisms.
- (ii) it allows DS-TE deployment to the required level of granularity and scope (e.g. only in a subset of the topology, or only for the number of classes required in the considered network)

3.2. Class-Types

The fundamental requirement for DS-TE is to be able to enforce different bandwidth constraints for different sets of Traffic Trunks.

[TEWG-FW] introduces the concept of Class-Types when discussing operations of MPLS Traffic Engineering in a Diff-Serv environment.

We refine this definition into the following:

Class-Type (CT): the set of Traffic Trunks crossing a link, that is governed by a specific set of Bandwidth constraints. CT is used for the purposes of link bandwidth allocation, constraint based routing and admission control. A given Traffic Trunk belongs to the same CT on all links.

Note that different LSPs transporting Traffic Trunks from the same CT may be using the same or different preemption priorities as explained in more details in [section 3.4](#) below.

Mapping of OAs to Class-Types is flexible. Different OAs can be mapped to different CTs, multiple OAs can be mapped to the same CT and one OA can be mapped to multiple CTs.

For illustration purposes, let's consider the case of a network running 4 Diff-Serv classes of services respectively based on the EF PHB [[EF](#)], the AF1x PSC [[AF](#)], the AF2x PSC and the Default (i.e. Best-Effort) PHB [[DIFF-FIELD](#)]. The network administrator may decide to deploy DS-TE in the following way:

- o from every DS-TE Head-end to every DS-TE Tail-end, split traffic into 4 Traffic Trunks: one for traffic of each Diff-

Serv class

- o because the QoS objectives for the AF1x Traffic Trunks and for the AF2x Traffic Trunks may be of similar nature (e.g. both targeting low loss albeit at different levels perhaps), the same (set of) Bandwidth Constraint(s) may be applied collectively over the AF1x Traffic Trunks and the AF2x Traffic Trunks. Thus, the network administrator may only define three CTs: one for the EF Traffic Trunks, one for the AF1x and AF2x Traffic Trunks and one for the Best Effort Traffic Trunks.

As another example of mapping of OAs to CTs, a network operator may split the EF traffic into two different sets of traffic trunks, so that each set of traffic trunks is subject to different constraints on the bandwidth it can access. In this case, two distinct CTs are defined for EF: one for the EF traffic subject to the first (set of) bandwidth constraint(s), the other for the EF traffic subject to the second (set of) bandwidth constraint(s).

DS-TE must support at least 2 CTs and up to 8 CTs. Those are referred to as CT_c, $0 \leq c \leq \text{MaxCT}-1 = 7$.

In a given network, DS-TE must not require the network administrator to always deploy the maximum number of CTs. The network administrator must be able to deploy only the number of CTs actually utilized.

3.3. Bandwidth Constraints

We refer to a Bandwidth Constraint Model as the set of rules defining:

- the maximum number of Bandwidth Constraints; and
- which CTs each Bandwidth Constraint applies to and how.

By definition of CT, each CT is assigned either a Bandwidth Constraint, or a set of Bandwidth Constraints.

We refer to the Bandwidth Constraints as BC_b, $0 \leq b \leq \text{MaxBC}-1$

Different models of Bandwidth Constraints are conceivable for control of the CTs.

For example, a model with one separate Bandwidth Constraint per CT could be defined. This model is defined by:

- $\text{MaxBC} = \text{MaxCT}$

Requirements for Diff-Serv-aware Traffic Engineering June 2002

- All LSPs supporting Traffic Trunks from CTc use no more than BCc

For illustration purposes, on a link of 100 unit of bandwidth where three CTs are used, the network administrator might then configure BC0=30, BC1= 50, BC2=20 such that:

- All LSPs supporting Traffic Trunks from CT0 use no more than 30 (e.g. Voice \leq 30)
- All LSPs supporting Traffic Trunks from CT1 use no more than 50 (e.g. Premium Data \leq 50)
- All LSPs supporting Traffic Trunks from CT2 use no more than 20 (e.g. Best Effort \leq 20)

As another example, a "Russian Doll" model of Bandwidth Constraints may be defined whereby:

- MaxBC= MaxCT
- All LSPs supporting Traffic Trunks from CTc (with $b \leq c \leq 7$) use no more than BCb

For illustration purposes, on a link of 100 units of bandwidth where three CTs are used, the network administrator might then configure BC0=100, BC1= 80, BC2=60 such that:

- All LSPs supporting Traffic Trunks from CT2 use no more than 60 (e.g. Voice \leq 60)
- All LSPs supporting Traffic Trunks from CT1 or CT2 use no more than 80 (e.g. Voice + Premium Data \leq 80)
- All LSPs supporting Traffic Trunks from CT0 or CT1 or CT2 use no more than 100 (e.g. Voice + Premium Data + Best Effort \leq 100).

Other Bandwidth Constraints model can also be conceived. Those could involve arbitrary relationships between BCb and CTc. Those could also involve additional concepts such as associating minimum reservable bandwidth to a CT.

At the time of writing this document, it is not clear whether a single model of Bandwidth Constraints is sufficient, which one it should be and how flexible this model really needs to be and what are the implications on the DS-TE technical solution and its implementations. Work is currently in progress to investigate the performance and trade-offs of different operational aspects of Bandwidth Constraints models. The DS-TE technical solution must specify one default bandwidth constraint model which must be supported by any DS-TE implementation. However, additional bandwidth constraint models may also be specified. The purpose of such a default model is to ensure that there is at least one common Bandwidth Constraints model implementation across various vendors

equipment and allows for easier deployment of DS-TE. However, this does not preclude a network operator to activate different Bandwidth Constraints models on different links in a network, if he/she wishes to do so.

In the selection of a default model, at least the following criteria are expected to be considered:

Le Faucheur et. al

9

Requirements for Diff-Serv-aware Traffic Engineering June 2002

- (1) addresses the scenarios in [Section 2](#)
- (2) works well under both normal and overload conditions
- (3) applies equally when preemption is either enabled or disabled
- (4) minimizes signaling load processing requirements
- (5) maximizes efficient use of the network

In selection criteria (2), "normal condition" means that the network is attempting to establish a volume of DS-TE LSPs for which it is designed; "overload condition" means that the network is attempting to establish a volume of DS-TE LSPs beyond the one it is designed for; "works well" means that under these conditions, the network should be able to sustain the expected performance, e.g., under overload it is x times worse than its normal performance [BC-MODEL].

These selection criteria will be further discussed and refined as part of the ongoing work on BC model selection. In particular, the applicability of criterion (5) needs to be qualified.

Regardless of the Bandwidth Constraint Model, the DS-TE solution must allow support for up to 8 BCs.

[3.4.](#) Preemption and TE-Classes

[TEWG-FW] defines the notion of preemption and preemption priority. DS-TE must retain full support of such preemption. However, a network administrator preferring not to use preemption for user traffic should be able to disable the preemption mechanisms described below.

The preemption attributes defined in [\[TE-REQ\]](#) must be retained and applicable across all Class Types. The preemption attributes of

setup priority and holding priority must retain existing semantics, and in particular these semantics must not be affected by the Ordered Aggregate transported by the LSP or by the LSP's Class Type. This means that if LSP1 contends with LSP2 for resources, LSP1 may preempt LSP2 if LSP1 has a higher set-up preemption priority (i.e. lower numerical priority value) than LSP2's holding preemption priority regardless of LSP1's OA/CT and LSP2's OA/CT.

We introduce the following definition:

TE-Class: A pair of:

- (i) a Class-Type
- (ii) a preemption priority allowed for that Class-Type. This means that an LSP transporting a Traffic Trunk from that Class-Type can use that preemption priority as the set-up priority, as the holding priority or both.

Note that by definition:

- for a given Class-Type, there may be one or multiple TE-classes using that Class-Type, each using a different preemption priority

Requirements for Diff-Serv-aware Traffic Engineering June 2002

- for a given preemption priority, there may be one or multiple TE-Class(es) using that preemption priority, each using a different Class-Type.

DS-TE must allow all LSPs transporting Traffic Trunks of a given Class-Type to use the same preemption priority. In other words, DS-TE must allow a Class-Type to be used by single TE-Class. This effectively allows the network administrator to ensure that no preemption happens within that Class-Type, when so desired.

As an example, the DS-TE solution must allow the network administrator to define a Class-Type comprising a single TE-class using preemption 0.

DS-TE must allow two LSPs transporting Traffic Trunks of the same Class-Type to use different preemption priorities, and allow the LSP with higher (numerically lower) set-up priority to preempt the LSP

with lower (numerically higher) holding priority when they contend for resources. In other words, DS-TE must allow multiple TE-Classes to be defined for a given Class-Type. This effectively allows the network administrator to enable preemption within a Class-Type, when so desired.

As an example, the DS-TE solution must allow the network administrator to define a Class-Type comprising three TE-Classes; one using preemption 0, one using preemption 1 and one using preemption 4.

DS-TE must allow two LSPs transporting Traffic Trunks from different Class-Types to use different preemption priorities, and allow the LSP with higher setup priority to preempt the one with lower holding priority when they contend for resources.

As an example, the DS-TE solution must allow the network administrator to define two Class-Types (CT0 and CT1) each comprising two TE-Classes where say:

- one TE-Class groups CT0 and preemption 0
- one TE-Class groups CT0 and preemption 2
- one TE-Class groups CT1 and preemption 1
- one TE-Class groups CT1 and preemption 3

The network administrator would then, in particular, be able to :

- transport a CT0 Traffic Trunk over an LSP with setup priority=0 and holding priority=0
- transport a CT0 Traffic Trunk over an LSP with setup priority=2 and holding priority=0
- transport a CT1 Traffic Trunk over an LSP with setup priority=1 and holding priority=1
- transport a CT1 Traffic Trunk over an LSP with setup priority=3 and holding priority=1.

The network administrator would then, in particular, NOT be able to :

- transport a CT0 Traffic Trunk over an LSP with setup priority=1

- and holding priority=1
- transport a CT1 Traffic Trunk over an LSP with setup priority=0 and holding priority=0

DS-TE must allow two LSPs transporting Traffic Trunks from different Class-Types to use the same preemption priority. In other words, the DS-TE solution must allow TE-classes using different CTs to use the same preemption priority. This effectively allows the network administrator to ensure that no preemption happens across Class-Types, if so desired.

As an example, the DS-TE solution must allow the network administrator to define three Class-Types (CT0, CT1 and CT2) each comprising one TE-Class which uses preemption 0. In that case, no preemption will ever occur.

Since there are 8 preemption priorities and up to 8 Class-Types, there could theoretically be up to 64 TE-Classes in a network. This is felt to be beyond current practical requirements. The current practical requirement is that the DS-TE solution must allow support for up to 8 TE-classes. The DS-TE solution must allow these TE-classes to comprise any arbitrary subset of 8 (or less) from the (64) possible combinations of (8) Class-Types and (8) preemption priorities.

As with existing TE, an LSP which gets preempted is torn down at preemption time. The Head-end of the preempted LSP may then attempt to reestablish that LSP, which involves recomputing a path by Constraint Based Routing based on updated available bandwidth information and then signaling for LSP establishment along the new path. It must be noted that there may be cases where the preempted LSP cannot be reestablished (e.g. no possible path satisfying LSP bandwidth constraints as well as other constraints). In such cases, the Head-end behavior is left to implementation. It may involve periodic attempts at reestablishing the LSP, relaxing of the LSP constraints, or other behaviors.

3.5. Mapping of Traffic to LSPs

The DS-TE solution must allow operation over E-LSPs onto which traffic from a single OA is transported.

The DS-TE solution must allow operation over L-LSPs.

The DS-TE solution may allow operation over E-LSPs onto which traffic from multiple OAs is transported, under the condition that traffic from those OAs can effectively be treated by DS-TE as a single atomic traffic trunk (in particular this means that traffic from those multiple OAs is routed as a whole based on a single

collective bandwidth requirement, a single affinity attribute, a single preemption level, a single Class-Type, ...). In that case, it is also assumed that OAs are grouped together in a consistent manner throughout the DS-TE domain (e.g. if OA1 and OA2 are transported together on an E-LSP, then there will not be any L-LSP transporting OA1 or OA2 on its own and there will not be any E-LSP transporting OA1 or OA2 with another OA).

3.6. Dynamic Adjustment of Diff-Serv PHBs

As discussed in [section 2.2](#), DS-TE may support adjustment of Diff-Serv PHBs parameters (e.g. queue bandwidth) based on the amount of TE-LSPs established for each OA/Class-Type. Such dynamic adjustment is optional. It is a local matter to the LSR and as such is outside the scope of this specification.

Where this dynamic adjustment is supported, it must allow for disabling via configuration (thus reverting to PHB treatment with static scheduler configuration independent of DS-TE operations). It may involve a number of configurable parameters which are outside the scope of this specification. Those may include configurable parameters controlling how scheduling resources (e.g. service rates) need to be apportioned across multiple OAs when those belong to the same Class-Type and are transported together on the same E-LSP.

The dynamic adjustment must take account of the performance requirements of each class when computing required adjustments.

3.7. Overbooking

Existing TE mechanisms allow overbooking to be applied on LSPs for Constraint Based Routing and admission control. Historically this has been achieved in TE deployment through factoring overbooking ratios at the time of sizing the LSP bandwidth and/or at the time of configuring the Maximum Reservable Bandwidth on links.

DS-TE must also allow overbooking and must effectively allow different overbooking ratios to be enforced for different CTs.

DS-TE should optionally allow the effective overbooking ratio of a given CT to be tweaked differently in different parts of the

network.

3.8. Restoration

With existing TE, restoration policies use standard priority mechanisms such as, for example, the preemption priority to effectively control the order/importance of LSPs for restoration purposes.

DS-TE must ensure that similar application of the use of standard priority mechanisms for implementation of restoration policy are not

Requirements for Diff-Serv-aware Traffic Engineering June 2002

prevented since those are expected to be required for achieving the survivability requirements of DS-TE networks.

Further discussion of restoration requirements are presented in the output document of the TEWG Requirements Design Team [[SURVIV-REQ](#)].

4. Solution Evaluation Criteria

A range of solutions is possible for the support of the DS-TE requirements discussed above. For example, some solutions may require that all current TE protocols syntax (IGP, RSVP-TE, CR-LDP) be extended in various ways. For instance, current TE protocols could be modified to support multiple bandwidth constraints rather than the existing single aggregate bandwidth constraint. Alternatively, other solutions may keep the existing TE protocols syntax unchanged but modify their semantic to allow for the multiple bandwidth constraints.

This section identifies the evaluation criteria that should be used to assess potential DS-TE solutions for selection.

4.1. Satisfying detailed requirements

The solution must address all the scenarios described in [section 2](#) and satisfy all the requirements listed in [section 3](#).

4.2. Flexibility

- number of Class Types that can be supported, compared to number identified in Requirements section
- number of Classes within a Class-Type

4.3. Extendibility

- how far can the solution be extended in the future if requirements for more Class-Types are identified in the future.

4.4. Scalability

- impact on network scalability in what is propagated, processed, stored and computed (IGP signaling, IGP processing, IGP database, TE-Tunnel signaling ,...).
- how does scalability impact evolve with number of Class-Types/Classes actually deployed in a network. In particular, is it possible to keep overhead small for a large networks which only use a small number of Class-Types/Classes, while allowing higher number of Class-Types/Classes in smaller networks which can bear higher overhead)

4.5. Backward compatibility/Migration

- backward compatibility/migration with/from existing TE mechanisms
- backward compatibility/migration when increasing/decreasing the number of Class-Types actually deployed in a given network.

5. Security Considerations

The solution developed to address the requirements defined in this document must address security aspects. DS-TE is not expected to add specific security requirements beyond those of Diff-Serv and

existing TE. Networks which employ Diff-Serv techniques might offer some protection between classes for denial of service attacks. Though depending on how the technology is employed, it is possible for some (lower scheduled) traffic to be more susceptible to traffic anomalies (which include denial of service attacks) occurring within other (higher scheduled) classes.

References

- [AF] Heinanen, J et al., "Assured Forwarding PHB Group", [RFC2597](#)
- [CR-LDP] Jamoussi et al., "Constraint-Based LSP Setup using LDP", [RFC3212](#), January 2002
- [DIFF-ARCH] Blake et al., "An Architecture for Differentiated Services", [RFC2475](#).
- [DIFF-FIELD] Nichols et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC2474](#).
- [DIFF-MPLS] Le Faucheur et al, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", [RFC3270](#), May 2002.
- [DIFF-NEW] Grossman, " New Terminology and Clarifications for Diffserv ", [RFC3260](#), April 2002.
- [EF] Davie et al., "An Expedited Forwarding PHB (Per-Hop Behavior)", [RFC3246](#), March 2002.
- [ISIS-TE] Smit, Li, IS-IS extensions for Traffic Engineering, [draft-ietf-isis-traffic-04.txt](#), August 2001.
- [OSPF-TE] Katz, Yeung, Traffic Engineering Extensions to OSPF, [draft-katz-yeung-ospf-traffic-06.txt](#), October 2001.

- [RSVP-TE] Awduche et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.

[TEWG-FW] Awduche et al, Overview and Principles of Internet Traffic Engineering, [RFC3272](#), May 2002.

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

[SURVIV-REQ] W.S. Lai, D. McDysan, J. Boyle, M. Carlzon, R. Coltun, T. Griffin, E. Kern, and T. Reddington, "Network Hierarchy and Multilayer Survivability," work in progress, October 2001.

Authors' Address:

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3
400, Avenue de Roumanille
06410 Biot-Sophia Antipolis
France
Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com

Martin Tatham
BT
Adastral Park,
Martlesham Heath,
Ipswich IP5 3RE
UK
Phone: +44-1473-606349
Email: martin.tatham@bt.com

Thomas Telkamp
Global Crossing
Olympia 6
1213 NP Hilversum
The Netherlands
Phone: +31 35 655 651
Email: telkamp@gblix.net

David Cooper
Global Crossing
960 Hamlin Court
Sunnyvale, CA 94089
USA
Phone: (916) 415-0437
Email: dcooper@gblix.net

Jim Boyle

Requirements for Diff-Serv-aware Traffic Engineering June 2002

Protocol Driven Networks, Inc.
1381 Kildaire Farm Road #288
Cary, NC 27511
USA
Phone: (919) 852-5160
Email: jboyle@pdnets.com

Luyuan Fang
AT&T Labs
200 Laurel Avenue
Middletown, New Jersey 07748
Phone: (732) 420-1921
Email: luyuanfang@att.com

Gerald R. Ash
AT&T Labs
200 Laurel Avenue
Middletown, New Jersey 07748
Phone: (732) 420-4578
Email: gash@att.com

Wai Sum Lai
AT&T Labs
200 Laurel Avenue
Middletown, New Jersey 07748
Phone: (732) 420-3712
Email: wlai@att.com

Pete Hicks
CoreExpress, Inc
12655 Olive Blvd, Suite 500
St. Louis, MO 63141
Phone: (314) 317-7504
Email: pete.hicks@coreexpress.net

Angela Chiu
Celion Networks
1 Sheila Drive, Suite 2
Tinton Falls, NJ 07724
Phone: (732) 747-9987

Email: angela.chiu@celion.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

Le Faucheur et. al

17

Requirements for Diff-Serv-aware Traffic Engineering June 2002

Chelmsford, MA 01824
Phone: (978) 244-3051
Email: tnadeau@cisco.com

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

