MPLS-DiffServ TE Class-Types

In the context of the MPLS-Diffserv TE Requirements Draft, class-types are as defined in the TEWG Framework draft (http://www.ietf.org/internet-drafts/draft-ietf-tewg-framework-03.txt) as aggregations of individual service classes, as follows:

Instead of having per-class parameters being configured and propagated on each LSR interface, 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 IGP link state advertisements by propagating information on a per-class-type basis instead of on a per-class basis, and also to allow better bandwidth sharing between classes in the same class-type. A class-type is a set of classes that satisfy the following two conditions:

1) Classes in the same class-type have common aggregate maximum requirements (and, if applicable, common minimum bandwidth requirements) to satisfy required performance levels.

2) There is no maximum or minimum bandwidth requirement to be enforced at the level of individual class in the class-type. Note that it is still possible, nevertheless, to implement some priority policies for classes in the same class-type to permit preferential access to the class-type bandwidth through the use of preemption priorities.

An example of the class-type can be a low-loss class-type that includes both AF1-based and AF2-based Ordering Aggregates. With such a class-type, one may implement some priority policy which assigns higher preemption priority to AF1-based traffic trunks over AF2-based ones, vice versa, or the same priority.

Some discussion within AT&T suggests that it might be very useful to generalize class types and consider them to be a combination of

1. QoS classes (e.g., as specified in Draft Recommendation Y.1541),
2. DiffServ/queuing priority classes,
3. MPLS/CAC priority classes,
4. restoration priority classes at both the MPLS-LSP and transport link level, and
5. preemption priority classes.

The above categories are inclusive of the CT definition given in the TE Framework, but go beyond to include considerations of restoration priority. Within this more general context, attached are a few ideas on definitions of 5 user-traffic class types and 1 control-traffic class type.

Admission control priority is a way of giving preference to admit higher priority LSPs ahead of lower priority LSPs. In the Table below we define key, normal, and BE admission control priorities. A key service LSP can be admitted in preference over a normal service LSP by reserving the last-available, minimum-level of bandwidth (called the "reserved bandwidth") for key service LSP admissions vs. normal service LSP admissions. That is, the higher priority LSP gets the reserved bandwidth when that is all the bandwidth left. This concept is widely used in practice today for voice and data services (e.g., "800 gold service", "emergency communication service", key-switched-digital-service, etc.).

In the Table below, we define restoration priority as either premium, basic, or unprotected. Restoration priority is achieved by provisioning sufficient backup capacity, as necessary, and allowing relative priority for access to available bandwidth when there is contention for restoration bandwidth. For example, a premium restoration priority LSP can take precedence over a basic priority LSP in accessing bandwidth, if there is contention. LSP preemption might also be used, such as for premium restoration priority, in which active unprotected priority LSPs could be bumped to provide available bandwidth for LSP restoration of premium priority LSPs. It is useful to minimize preemption so as to not interrupt established service.

Class-Type Characteristics	User-Traffic Class Type					Control-Traffic Class Type
	1	2	3	4	5	1
QoS Class	Y.1541 Class-0 Interactive (Real-time)	Y.1541 Class-1 Interactive (Real-time)	Y.1541 Class 2 Non-interactive (Low loss)	Y.1541 Class 3 or 4 Non-interactive (Low loss)	Y.1541 Class 5 Unspecified (Best-effort)	Y.1541 Class-0Interactive (Real-time)
MPLS CAC priority	Key	Normal	Key	Normal	Best-effort	Key; LSP Preemption allowed
DiffServ Queuing priority	EF	EF	AF1	AF2	BE	EF
MPLS restoration priority (per-MPLS-LSP)	Premium	Basic	Premium	Basic	Unprotected	Premium; LSP Preemption allowed
Transport restoration priority (per-transport-link)	Premium (weighted avg. across MPLS-LSP CTs)	Basic (weighted avg. across MPLS-LSP CTs)	Premium (weighted avg. across MPLS-LSP CTs)	Basic (weighted avg. across MPLS-LSP CTs)	Unprotected	Premium (weighted avg. across MPLS-LSP CTs) ; LSP Preemption allowed

Notes

Five class types are defined for user traffic. Note that the class types are mapped to the corresponding QoS classes as specified in the current Draft Recommendation Y.1541 (not finalized yet), as well as in DiffServ. CAC priority and restoration priority are as defined above.

The inclusion of QoS priority classes, admission-control priority classes, restoration priority, and preemption priority classes into a combined set of 6 CT definitions is intended to minimize the number of possible class types across these 4 levels of priority classes. That is, if the

number of QoS classes = W

number of admission-control priority levels = X
number of restoration priority levels = Y,
number of preemption priority levels = Z,

then the combination of the possible number of CTs when considering all 4 dimensions independently would be

number of independent CTs = WXYZ (could be a large number)

In the Table above we define 6 CTs, which combine the 4 dimensions into a smaller number (6) of combinations.

QoS priority and CAC priority information can be conveyed in LSP setup using existing Priority TLVs already being specified. Conveying the restoration priority information could be done in the Restoration TLV being defined.

A separate class type is defined for network control traffic. This is for control traffic generated by routing protocols, network management, and signaling to support service requests. This class of traffic has the following characteristics:

1. must have a very low drop precedence: e.g., dropped control traffic may cause routing to fail or flap, leading to network instability or even prolonged and widespread network failure,
2. a relatively small bandwidth guarantee under normal network conditions: to minimize overhead (control traffic is usually designed to consume a relatively small amount of bandwidth),
3. high priority: large user packets on slow links may cause significant delays that will reduce the responsiveness needed by control traffic,
4. preemption of all other traffic under overload or failed conditions: to ensure that control traffic gets through with minimal loss and delay.