



# **Description of QoS classes and Service Level Agreements**

**JANET QoS Development Project Report**

**Version 1.6**

**Author:  
Victor Olifer**

January 2005

Abstract .....	2
Introduction.....	2
DiffServ model for a multi-domain environment .....	3
Per Hop Behaviour.....	3
Multi-domain SLAs .....	4
Edge functionality.....	8
SLA types.....	8
IP Premium .....	10
Static IP Premium .....	11
Guaranteed IP Premium with booking services .....	12
Less than Best Effort (LBE).....	15
References:.....	16

## Abstract

This document contains a high-level description of the QoS classes of service which are recommended for development and eventual deployment across JANET. The recommendations have been formulated on the basis of experience gained in the framework of the JANET QoS Development Project.

Two QoS services are described: IP Premium and Less than Best Effort. These complement the regular Best Effort service, which is known on JANET as the Basic IP Transmission service and is currently the only production transport service offered to the JANET community.

## Introduction

QoS uses various mechanisms to decrease the negative effects of congestion, i.e. large queues building up in packet-switched networks. Congestion leads to large and variable packet delays and loss.

QoS improves the basic transport service offered by IP networks (also known as Best Effort service) by processing different traffic classes in different ways, according to their requirements. Modern network traffic can be categorised by the different requirements of the applications concerned, in terms of relative predictability of the data rate (for example, traffic with Constant Bit Rate or Variable Bit Rate), traffic sensitivity to packet delays and their variations, and traffic sensitivity to packet losses or corruption. Even though each application can be described by its own set of requirements (which can be quantitative or qualitative), a more scalable approach is to combine all the applications into a few classes with similar requirements. In very simplistic terms there could be only two classes of applications: time-sensitive and non-time-sensitive (also known as elastic). This simple scheme becomes more elaborate as the sensitivity of traffic to data loss is taken into account, leading to more classes of traffic. Asynchronous Transfer Mode (ATM) takes such an approach with

four predefined class of applications, named A, B, C and D. For each class of traffic there is a list of quantitative criteria which assess QoS for that class. For example, for the time-sensitive class, the criteria include latency and the loss of packets which a network is delivering to an end-user. Note that the loss is due to the queues overloading and there are no other causes, e.g. bit distortion because of electromagnetic noise.

The main idea behind QoS is to provide better than Best Effort treatment to a restricted amount of traffic for a particular class, which needs that kind of improved treatment. For example, a network can transfer some privileged traffic with a very low level of delay only if the amount of that traffic does not exceed a particular level – beyond this level, queues of privileged packets start building up. In other words, when there is too much privileged traffic we don't have privileged traffic at all.

A QoS service is an advanced transport service that results from the systematic deployment of QoS mechanisms across a network. (It is advanced compared to the regular Best Effort service.) Very popular examples of such a service are IP Premium, Platinum, Gold, Silver or Bronze services. These are characterised by performance indicators that are well-known QoS metrics, such as One Way Delay (OWD), IP Packet Delay Variation (IPDV or jitter) and loss. As with any transport service, a QoS service should be characterised not just by the QoS metrics but by generic performance indicators such as availability of service, its scope, schedule and so on.

A QoS service is usually intended to transport particular classes of traffic. Traffic classes can be mapped directly onto QoS, i.e. each traffic class is served by a particular QoS service. Alternatively, several traffic classes can be mapped onto one QoS service so that the number of QoS services is less than number of traffic classes. For example Voice over IP (VoIP), videoconferencing and real-time traffic of computer-aided manufacturing systems could be served by one QoS service class that is often called IP Premium.

The QoS services considered in this document are based on the DiffServ (Differentiated Services) model, which has been developed by the IETF (Internet Engineering Task Force),<sup>1,2</sup> and is deployed by several commercial ISPs and across the pan-European education and research network GÉANT.<sup>3</sup>

The JANET QoS Development Project explored the DiffServ model, and the recommendation for the QoS services in this document is based on experience gained during the project.

## ***DiffServ model for a multi-domain environment***

### ***Per Hop Behaviour***

The DiffServ model is currently considered to be the most scalable QoS model for IP networks as it treats aggregate flows as service objects. Furthermore, every router makes its own decision based on the service provided to a particular aggregate flow. This decision is made independently and need not be co-ordinated with other routers. This approach is called per hop behavior (PHB). Since DiffServ does not consider flows along the entire route from the source to the destination, it is more precise to say that the service objects are these traffic classes rather than aggregate flows.

Accordingly, the DiffServ architecture does not use a signalling protocol, which normally tells routers along a path how much bandwidth should be reserved for each

particular flow. Instead, a value carried by every IP packet in its DiffServ field is used as an attribute to specify that the packet belongs to a particular traffic class. Based on this value, called DSCP (DiffServ Code Point), each network router decides on the service to be provided to that particular packet as opposed to the packets of other classes. Supporting only a few traffic classes instead of a large amount of separate flows makes DiffServ relatively simple in terms of the additional load on routers, as routers do not store and maintain state information about individual flows.

Evidently, if all routers agreed to provide the same treatment for a particular DSCP-marked traffic class then users would have an end-to-end QoS service. However, routers must co-ordinate their efforts in:

- the definition of QoS services;
- marking packets for a particular QoS service with the same DSCP value;
- agreeing levels of maximum rates for privileged QoS services in order not to suppress less privileged services significantly.

The IETF standardised several QoS PHBs and their respective DSCP values:

- Expedited Forwarding PHB<sup>4</sup> with DSCP 46, designed to provide the best QoS service, with strict low limits for QoS metrics such as packets delays, delay variation (jitter), loss and disordering. This guarantees these desired values only when the rate of Expedited Forwarding traffic does not exceed an agreed upper limit. All time-sensitive traffic should be served by Expedited Forwarding PHB.
- Assured Forwarding PHBs, which should be better than Best Effort in terms of QoS metrics, but have less strict limits in comparison to Expedited Forwarding. Assured Forwarding PHBs are designed for guaranteed delivery of traffic with specified limits of rate and loss, but with much wider limits for delays and jitter. Traffic that is sensitive to rate and loss but not so sensitive to delays should be served by Assured Forwarding PHBs. An example of such traffic is videostreaming. Each Assured Forwarding class supports three different DSCPs which define the probability of discarding packets by each class.
- Quite recently, the IETF has published RFC 3662, *A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services* which contains the description for a Less than Best Effort service (LBE), based on the Scavenger service developed by the Internet2 community. LBE has been deployed in some academic and research networks, including GÉANT. The LBE and Lower Effort (LE) PHBs have a DSCP value of 8.

It is worth emphasising that PHB is not an end-to-end QoS service. In fact, routers that support PHB are only really the building blocks for creating an end-to-end QoS service. To provide such a service, all routers between the end-users should be configured according to the same policy framework and behave consistently.

### **Multi-domain SLAs**

The DiffServ model was designed to accommodate the multi-domain nature of modern IP networks seamlessly. The Internet as such consists of a large number of domains; JANET is a three-tier multi-domain network (see Figure 1 below). If the

hierarchical structure of some Regional and Campus Networks is taken into account then the number of JANET tiers is further increased.

DiffServ suggests that every pair of neighbouring domains should negotiate a Service Level Agreement (SLA), which forms the common policy framework mentioned above. The end-to-end service for end-users could be built on the basis of a chain of SLAs between neighbouring domains. As end-users on JANET are usually in Campus Networks, this chain will include two Campus Networks, two Regional Networks and the JANET backbone.

An SLA is a unidirectional agreement in which one domain, e.g. domain J in Figure 1, agrees to accept traffic from another domain, e.g. domain B, and transfers it towards its destination according to a particular level of QoS, for example, with delays of no more than 100 ms for 99% of the packets to be transferred. A domain which agrees to accept and transfer traffic is a domain-provider, and a domain which directs traffic into the domain-provider is a domain-client.

An SLA defines the quantitative or qualitative QoS criteria for transferring traffic between a domain-provider and a third domain, or between a domain-provider and one of its end-users. In Figure 1, SLA B-J means that domain B is a domain-client and domain J is a domain-provider. The direction of the arrow corresponds to the direction of the traffic covered by the SLA. Quantitative guarantees are usually expressed in standard QoS metrics, defined by the IETF's IP Performance Metrics (IPPM) working group<sup>5,6,7</sup> – for example, a percentile of one-way packets which are delayed above some limit.

Qualitative guarantees could be given – for example, if a network experiences periods of congestion, the provider could promise the client a significant decrease in delays during those periods for the IP Premium service compared to the Best Effort service.

As traffic between domains travels in both directions, two SLAs should be negotiated between neighbouring domains, as shown in Figure 1. SLAs between two particular domains may differ because, for example, of the different amounts of traffic that need to travel in each direction.

An SLA could be agreed not only between domains but between a domain and its end-users as well.

An SLA consists of two parts: the administrative/legal part and the Service Level Specification (SLS).<sup>8</sup>

The SLS describes the technical parameters of the SLA, e.g.:

- the values of QoS metrics which the domain-provider guarantees for the client traffic it transfers;
- the values of non-QoS metrics of the service, e.g. availability;
- the scope of the service, i.e. a description of the bounds between which the service defined by the SLA could be provided to clients. For example, a domain can provide the service between all of its edge routers, or only between some particular routers (the same could apply for the end-users' equipment);
- the traffic profile, which describes the parameters of the traffic transferred by the domain-provider. The traffic profile usually limits the sustained rate and

burstiness of traffic; the parameters that are not frequently used are the top and bottom limits of packet size.

The traffic profile is a very important part of an SLS because, without knowing the traffic to be transferred, a domain-provider cannot give any guarantees to the domain-client, whether quantitative or qualitative. To know the traffic sustained rate limit lets a domain-provider plan the bandwidth distribution among its clients, and hence assess the QoS metrics for the traffic to be transferred through the domain.

In fact, the SLS determines two kinds of responsibilities:

- those of the domain-provider to transfer traffic of an agreed profile with an agreed quality;
- those of the domain-client to direct traffic into the domain-provider within an agreed profile.

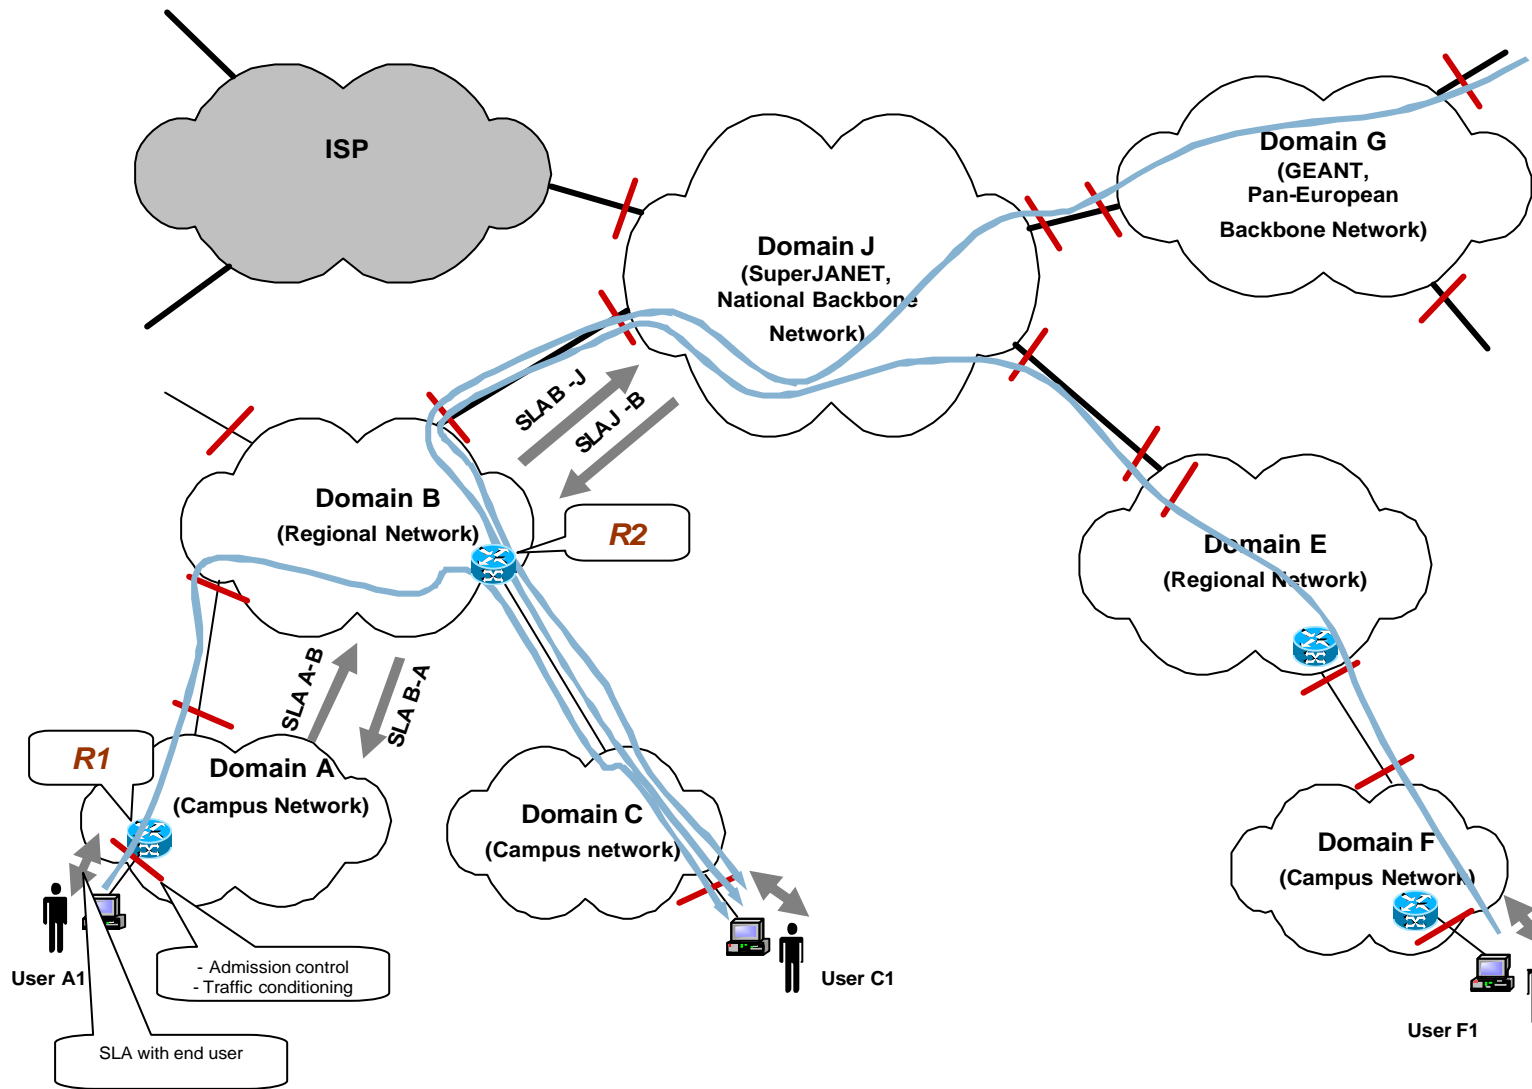


Figure 1. Multi-domain architecture of the DiffServ model.

The administrative/legal part of the SLA defines all the non-technical parameters of the service agreement between the domain-provider and the domain-client. It is quite flexible as any aspects of the agreement can be included here. Important components of the administrative/legal part of the SLA include:

- the service description;
- description of the monitoring tools for checking SLS parameters (they are needed for both the provider and the client);
- service provisioning policy and procedures;
- service support and troubleshooting procedures.

Service provisioning policy and procedures define how the client can request service from the provider. These are not less important than the SLS parameters: for example a potential user could be happy with the QoS metrics guaranteed by a service, but could find the provisioning policy that requires a service to be booked three months in advance absolutely unacceptable.

### ***Edge functionality***

Under the DiffServ model, the edge (border) routers and the core routers serve different QoS functions.

A domain-provider's edge (border) routers act as check-in points for the domain, because they communicate directly with the edge routers of other administrative domains and accept traffic from them for processing. Edge routers check ingress traffic for conformity to the SLS traffic profile. The DiffServ model calls this traffic conditioning, which usually includes admission control, measurement, policing and traffic shaping. Admission control checks whether traffic could in principle use the requested QoS class (according to the DSCP value), and measurement checks the conformance of the traffic to the traffic profile. If the traffic does not conform to its profile then policing takes place, which means that excess packets are dropped by an edge router. Policing could alternatively lead to remarking the DSCP value, usually transforming privileged traffic into Best Effort traffic. Traffic shaping means changing inter-packet intervals to even out the traffic. Usually, shaping is applied to egress traffic in order to decrease queues in downstream routers, resulting in less bursty traffic and hence smaller queues. Shaping is also used to reconcile the parameters of a domain's egress traffic and the policing of a downstream domain's ingress traffic.

After traffic conditioning, the edge routers queue traffic according to QoS class requirements, using priority or weighted queuing techniques.

Figure 1 shows how this situation is applicable to JANET. Traffic conditioning is represented by red lines crossing the ingress interfaces of domain edge routers. As all core routers are under the same administrative control, they can trust the edge routers in terms of the agreed parameters of traffic, and themselves only handle queuing and shaping.

### ***SLA types***

1. **SLA for well-known QoS classes and SLA for individual QoS classes.** Well-known QoS classes have been developed for popular classes of network applications. The most important and elaborated is the IP Premium class, which can also be known by other names. This was developed for the most time-sensitive applications. The typical SLA for IP Premium is also quite well-known, especially the SLS part, usually based on Expedited Forwarding PHB.

A pair of neighbouring domains could also develop its own SLA for a more specific QoS class, which would include individual QoS indicators, and a description of service and service provisioning.

2. **Wholesale and retail SLAs.** Inter-domain SLAs are usually wholesale,<sup>9</sup> i.e. they describe aggregate traffic between domains and do not pay attention to the traffic flows of end-users. For example, the traffic profile of a wholesale SLS could have a single limit for the total amount of IP Premium traffic accepted by a domain through every access interface.

A retail SLA is intended for end-users, not domains, and describes all agreements in terms of individual traffic flows.

3. **Destination-aware and destination-unaware SLAs.** A destination-aware SLA specifies the destination of the traffic. For example, domain J in Figure 1 could agree to accept IP Premium traffic from domain B at up to 10% of the access link's bandwidth if it is directed to domain G, and up to 15% of the access link's bandwidth if it is directed to domain E.

A destination-unaware SLA does not include the traffic destination: e.g. domain J could agree to accept IP Premium traffic from domain B at up to 25% of the access link's bandwidth regardless of its destination.

4. **Edge-to-edge or end-to-end SLAs.** An SLA between domains is normally an edge-to-edge SLA, as the scope of agreement could only cover the edge routers of the neighbouring domains. An end-to-end SLA is a retail-type agreement concluded by the domain operator with its end-users, describing the quality of the service for the users of the domain when they direct traffic to other end-users, who could be users of the same or some other domain. If the latter is the case then there is a problem with formulating the quantitative QoS parameters (if the SLA has such parameters) because they are heavily dependent on a chain of domains leading from the source to the destination of the traffic. Several approaches could be applied:
  - the end-to-end SLS could be qualitative;
  - the end-to-end SLS could be built on the worst case estimation, which is obtained by measuring all possible destinations available for the users of the domain to send traffic to;
  - the end-to-end SLS could combine the worst case estimation with a promise to specify more exact parameters during the service provisioning.

## **IP Premium**

In the framework of the JANET QoS Development project trial, two QoS services were defined and deployed across JANET:

- IP Premium
- LBE

which together with the regular Best Effort service resulted in three different QoS services. Both IP Premium and LBE are well-known QoS services, which are based on the IETF recommendations and deployed by a number of Internet Service Providers.

There is a third well-known QoS service, IP Plus (IP+), based on the IETF Assured Forwarding PHB specification, mentioned above. IP Plus was described along with IP Premium and LBE in the QoS Think Tank<sup>10</sup> – the initiative which accumulated the expertise of JANET community in the QoS area in 2001. IP Plus is an intermediate service between Best Effort and IP Premium, and could be useful for a number of applications that have less strict demands for latency parameters but need bandwidth guarantees (e.g. for streaming audio and video broadcasting). However, IP Plus was not explored in the course of the JANET QoS Development Project in order not to overcomplicate the first large-scale QoS trial across JANET.

This section includes a high-level description of the types of IP Premium service which could be deployed as a JANET production service. They are based on experience gained from the project and the external experience of several sources.

The next section includes a description of the types of LBE service.

All the types of IP Premium should be based on the well-known SLS parameters for this class of service. These parameters are:

- DSCP 46 value for marking IP Premium packets;
- using the following QoS metrics as quality indicators of the service:
  - upper-bounded one-way delay
  - upper-bounded IP Packet Delay Variation (IPDV)
  - no packet loss due to congestion
  - guaranteed capacity.

Guaranteed capacity (i.e. bandwidth) should be at least a quantitative parameter of the IP Premium service. Others QoS metrics could be quantitative or qualitative.

The former would result in a guaranteed IP Premium service with a high level of probability of guarantees (there are other informal names for such a service: ‘hard’, ‘strict’ etc.).

The latter would result in a ‘soft’ or ‘basic’ IP Premium service where QoS metrics such as latency and loss are described in qualitative terms, e.g. as ‘significantly lower than the latency and loss experienced by a Best Effort service’.

The differences between hard and soft IP Premium services, and the procedure for provisioning the type of service, lie in the SLS. Below are some examples of such services.

## **Static IP Premium**

This is a 'soft' version of the IP Premium service based on a 'pure' DiffServ multi-domain model shown in Figure 1. The term 'static' relates to the method used for service provisioning. It means that there is no dynamic booking procedure for end-users of the service. The service has a few well-known users who can send traffic marked as IP Premium without advance booking. This was the scenario that was implemented by the project participants. This model will most probably be used for a few videoconferencing studios and VoIP software switches which communicate only with each other.

Should the policing limits of edge routers be equal to or higher than the total IP Premium traffic of the well-known users, this kind of service might be a guaranteed IP Premium service. However, this scenario is very improbable as it needs very high policing limits. When policing limits are lower than the maximum total amount of the IP Premium traffic of all registered IP Premium users, then this kind of service is 'soft', as a situation where a policer drops some IP Premium traffic as excess traffic cannot be excluded.

A static IP Premium service could be destination-unaware or destination-aware. In the former case, admission control is based only on the source IP addresses of registered IP Premium users. The latter case needs admission control to be based on the source IP address and destination IP address. This is more reliable as it excludes situations where a registered IP Premium user communicates with a non-IP Premium user. Also, the QoS metrics will be lower than if both users were IP Premium users. (However, note that although we use the term 'lower', we cannot estimate a degree of quantitative degradation.)

If trust relationships could exist between domains then admission control and conditioning of individual IP Premium flows could be conducted only by domains of registered IP Premium users. For the JANET community, user domains would normally be Campus Networks and the existence of such trust relationships is very likely. In turn, such a system of trust relationships relieves high-level domains (i.e. Regional Networks and JANET) of having to perform admission control and traffic conditioning at the flow level. Delegating this function to Campus Networks makes the high-level domains more reliable and stable, as their configuration of edge routers supporting only traffic classes is very simple. It is worth emphasising that these will be limited trust relationships as high-level domains will still police aggregate flows of IP Premium traffic, protecting themselves from IP Premium overflow. They will only delegate granular admission control and policing functionality to low-level domains. Misconfiguring equipment in a low-level domain could possibly lead to an IP Premium traffic jam in the output interface of a high-level domain. However, this is more likely to happen in a static destination-unaware model because it will not be misconfiguring but rather normal practice for the destination-unaware service.

The advantage of the 'softness' of Static IP Premium is its easy incorporation into the chain of over-provisioned non-IP Premium domains. On the one hand, overprovisioned domains could provide relatively good QoS metrics without differentiating traffic into classes. On the other hand, an absence of admission control and traffic conditioning operations on the border of such domains does not allow

quantitative guarantees of QoS to users. Therefore, over-provisioned domains cannot be part of a chain for guaranteed IP Premium service but could be part of a 'soft' IP Premium service.

A static destination-unaware model could potentially compromise the IP Premium service, because without any control of the flow paths, many IP Premium service flows go through the same output interface of a router, and the total rate of IP Premium traffic exceeds a technological limit for priority traffic beyond which delay and jitter would be too high for time-sensitive applications. Another unwanted effect of overloading an interface with IP Premium traffic is a low share of bandwidth for other traffic classes, such as Best Effort and LBE in our case. Such a situation is shown in Figure 1, where several IP Premium flows go through the output interface of R2 towards the domain C.

The drawback of a static destination-aware IP Premium model is the potentially ineffective use of reserved bandwidth. The static nature of a reservation means that it is made on the basis of network topology, without taking into account real need in transferring data. For example, imagine that domain B in Figure 1 supports a static destination-aware IP Premium service. Domain B has four external interfaces connecting it to the neighbouring domains, and for each interface the domain should have three separate reservations for ingress IP Premium traffic, leading to three external interfaces. These reservations are permanent even though traffic to some or all of those destinations could be zero for significant periods of time.

### ***Guaranteed IP Premium with booking services***

A guaranteed bookable IP Premium service is very desirable as it offers users the highest level of QoS. It is a dynamic service, as a reservation is created only if traffic is going to be transferred across a network, and this reservation will only exist for a predefined period of time.

The general idea of a guaranteed bookable IP Premium service is illustrated by Figure 2.

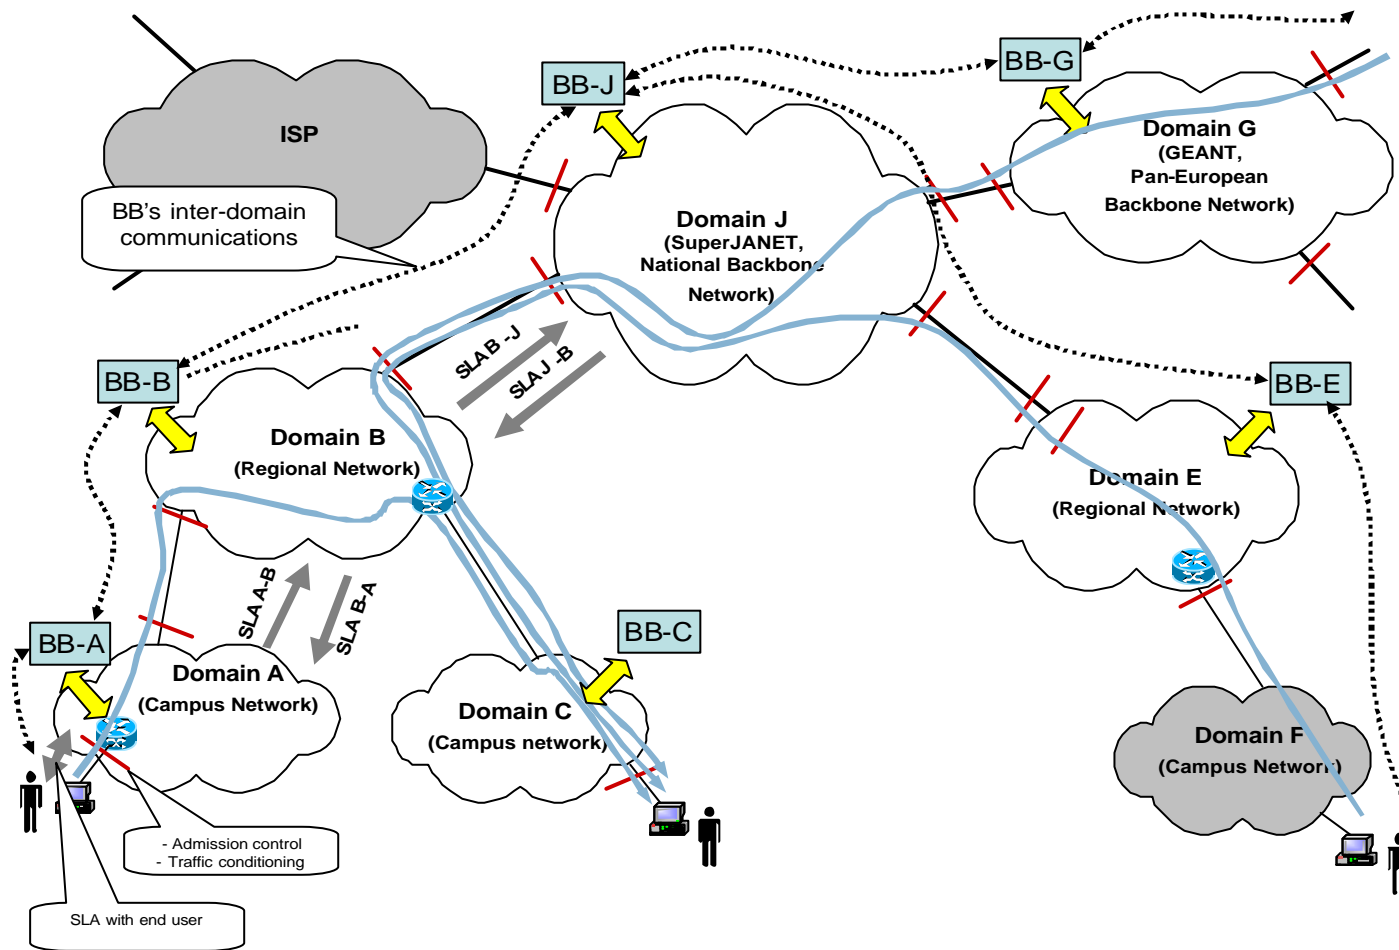


Figure 2. Guaranteed bookable IP Premium service.

Guaranteed IP Premium is always destination-aware, as it needs to provide the necessary guarantees in order to reserve the network resources for the traffic flow from the source address to the destination address. Advance reservation prevents the aggregate IP Premium flow exceeding the policing limit of the rate on the input interface or the technological limit of the rate on the output interface of the downstream routers.

There are two ways of making this reservation. The first uses a signalling protocol on packets sent along the path through domains and networks to inform the routers of the traffic profile parameters. For IP networks, RSVP (Resource Reservation Protocol) is a standard protocol for making reservations. It makes them automatically, but this is currently considered suitable only for enterprise-wide networks. This is because of RSVP's poor scalability, as all routers must keep state information about each individual user flow.

The network diagram shown in Figure 2 illustrates the second approach, based on the Bandwidth Broker (BB) model. A Bandwidth Broker is an additional tool, external to the communication equipment of a domain, i.e. to its routers and switches.

According to this model every domain has its own Bandwidth Broker, which is responsible for processing requests for the IP Premium service within its own domain. The mechanism for a Bandwidth Broker is as follows. The end-user sends a request to the Bandwidth Broker of his/her own domain and, on the basis of the destination and rate of the request and the available resources of the domain, the Bandwidth Broker defines the possibility of allocating resources to this request. It forwards the request to the Bandwidth Broker of the domain downstream (in terms of the user flow), and so on. If all Bandwidth Brokers agree to accept a request then the reservation is made, the respective border routers of all the intermediate domains are reconfigured to permit the new IP Premium flow (or another better than Best Effort service flow, e.g. IP Plus), and the end-user is notified about the successful booking. Usually the Bandwidth Broker is a software system that is implemented and configured on a dedicated server.

The Bandwidth Broker model is quite new and has not yet been deployed widely in a commercial or an academic environment. The basic architectural ideas were explored by the Internet2 community but the service has not been implemented. Commercial ISPs are currently not very motivated in the development and deployment of the model, as they mostly provide advanced QoS services only across their own networks.

Implementing the Bandwidth Broker model does not have to be completely automated. The first stage can be completely manual, where the role of the Bandwidth Broker is performed by the network operations staff of a particular NOC. The automation process could be gradual, by assigning some areas of the Bandwidth Broker's functionality to software (e.g. Web-based interface for booking) while others remain manual.

A guaranteed IP Premium service can have a range of quantitative guarantees. The top-level service ('hard' or advanced IP Premium) should provide quantitative guarantees for all QoS metrics while the basic service could guarantee only a certain amount of bandwidth. It is reasonable to start with basic IP Premium as it is simpler to

implement, and it allows over-provisioned domains to be included in the scope of service.

As well as a Bandwidth Broker, a guaranteed IP Premium service needs authorisation and authentication elements. Every domain could implement these elements in its own way and a trust relationship system can be implemented to delegate some functionality to Campus Networks, as with Static IP Premium.

### ***Less than Best Effort (LBE)***

LBE is based on the Internet2 Scavenger Service.<sup>11</sup> Internet2 defines it in the following way:

‘The scavenger service creates a parallel virtual network with very scarce capacity. This capacity, however, is elastic and can expand into capacity of the normal best-effort class of service whenever the network has spare cycles. The expansion happens with a very high time granularity: everything not used by the default class is available for the scavenger class.’

The LBE service has lower priority than Best Effort. Essentially, it conforms to the DiffServ idea of providing services with *different* QoS levels. It is rather surprising that this kind of QoS service has appeared relatively late.

RFC 3662, *A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services* generalises the LBE approach and states that it:

‘is intended for traffic of sufficiently low value (where “value” may be interpreted in any useful way by the network operator), in which all other traffic takes precedence over LE [Lower Effort] traffic in consumption of network link bandwidth. One possible interpretation of “low value” traffic is its low priority in time, which does not necessarily imply that it is generally of minor importance. From this viewpoint, it can be considered as a network equivalent to a background priority for processes in an operating system. There may or may not be memory (buffer) resources allocated for this type of traffic.’

LBE uses the DSCP value 8 (IP precedence 1) for marking packets, following the practice of the Scavenger service and the recommendations of RFC 3662.

LBE also differs from IP Premium in that it does not need an SLA between domains, because the nature of the service guarantees nothing to its users. However, typical LBE implementation uses Weighted Round Robin queuing in all routers with a very small percentage (1-3%) of the links bandwidth. This is because Weighted Round Robin gives all its classes a provisioned bandwidth minimum during congestion periods, but they can receive up to the total link bandwidth while other classes (including priority ones) temporarily are not using the link. This feature of Weighted Round Robin conforms exactly to the definition of LBE.

The guaranteed small percentage of bandwidth for LBE is intended to keep established TCP sessions going even during congestion periods.

LBE could be of use for bulky Grid applications which need to transfer huge amounts of data, but which at the same time are not time-sensitive and can endure a wide range of delays in the delivery of the data. When served as Best Effort, such bulky traffic can distort regular traffic (the majority of which is Web traffic). When served as LBE,

this huge amount of data could be delivered, say, over several hours, using gaps in the regular (and IP Premium) traffic. This would be an adequate quality of service for this kind of traffic. Note that when the LBE service is under populated and the Best Effort service is congested then the LBE service can provide better than Best Effort, as the relative utilisation of the LBE bandwidth will be lower than the same characteristic of Best Effort.

## References:

---

- <sup>1</sup> RFC2475. An Architecture for Differentiated Service. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss.
- <sup>2</sup> RFC3260. New Terminology and Clarifications for Diffserv. D. Grossman.
- <sup>3</sup> Specification and implementation plan for a Premium IP service. GÉANT, Deliverable D9.1.
- <sup>4</sup> RFC3246 – An Expedited Forwarding PHB (Per-Hop Behavior).
- <sup>5</sup> RFC2330. Framework for IP Performance Metrics, Paxson, V., Almes, G., Mahdavi, J. and M. Mathis.
- <sup>6</sup> RFC2679. A One-way Delay Metric for IPPM. Almes, G., Kalidindi, S. and M. Zekauskas.
- <sup>7</sup> RFC2680. Almes, G., Kalidindi, S. and M. Zekauskas. A One-way Packet Loss Metric for IPPM.
- <sup>8</sup> Service Level Agreements specification for IP Premium Service. SEQUIN. Deliverable D2.1 – Addendum 2 (<http://archive.dante.net/sequin/deliverables/SEQ-01-043.pdf>)
- <sup>9</sup> Service Level Agreements specification for IP Premium Service. SEQUIN. Deliverable D2.1 – Addendum 2.
- <sup>10</sup> Report of Quality of Service Think Tank, July 2001 ([http://www.ja.net/development/qos/qos\\_tt\\_report.pdf](http://www.ja.net/development/qos/qos_tt_report.pdf))
- <sup>11</sup> QBone Scavenger Service (QBSS) (<http://qbone.internet2.edu/qbss>).