Internet Draft
Expiration: August 1999
File: draft-ietf-diffserv-rsvp-02.ps

Y. Bernet, Microsoft
R. Yavatkar, Microsoft
P. Ford, Microsoft
F. Baker, Cisco
L. Zhang, UCLA
K. Nichols, Cisco
M. Speer, Sun Microsystems
R. Braden, ISI

Interoperation of RSVP/Int-Serv and Diff-Serv Networks

February 26, 1999

Status of Memo

Abstract

Differentiated Services (diff-serv) and RSVP/Integrated Services (RSVP/int-serv) provide complementary approaches to the problem of providing QoS for Internet end systems. These approachs must be able to coexist and effectively interoperate. This document outlines one important model for such interoperation, in which diff-serv is used by transit networks in the core of the Internet while hosts and edge networks use RSVP/int-serv. It also contains a brief discussion of some alternative models for interoperation.

1. Introduction

Work on QoS-enabled IP networks has led to two distinct approaches: the Integrated Services (int-serv) architecture [12] and its signaling protocol RSVP [1], and the Differentiated Services (diff-serv) architecture [10].

RSVP enables applications to signal per-flow QoS requirements to the network, with explicit admission control. Int-serv uses RSVP signaling to request tight QoS with quantitative parameters. Int-serv also imposes fine-grain policing and scheduling of traffic, to ensure that admitted flows receive their service requests in strict isolation from each other and from best-effort traffic. RSVP signaling configures packet classifiers in the int-serv capable routers along the path of the flow. These classifiers perform a fine-grain or `MF' [10] classification of packets, using on IP addresses and port numbers for example.

Some basic limitations to the RSVP/int-serv model have impeded its deployment in the Internet at large.

- The use of per-flow state and per-flow processing raises scalability concerns for large networks.
- Only a small number of hosts currently generate RSVP signaling. While this number is expected to grow dramatically, some applications may never generate RSVP signaling.
- Some applications require a form of QoS that cannot be expressed using the int-serv model.
- The necessary policy control mechanisms -- access control, authentication, and accounting -- are not available.

The market is pushing for immediate deployment of a QoS solution that addresses the needs of the Internet as well as enterprise networks. This push led to the development of Differentiated Services. In contrast to the per-flow orientation of int-serv and RSVP, diff-serv networks classify packets into one of a small number of aggregated flows or "classes", based on bits set in the TOS field of each packet's IP header. This is known as `BA' classification [10]. In addition to eliminating the requirement for per-flow state, diff-serv QoS can initially be deployed using long-term provisioning rather than short-term reservations established by end-to-end signaling.

We view int-serv and diff-serv as complementary tools in the pursuit of end-to-end QoS. For many applications, the loose or "qualitative" QoS provided by diff-serv will be adequate. However, some

applications will require the tight quantitative end-to-end QoS assurance provided by int-serv and RSVP. Current examples of applications that need tight QoS include IP-telephony, video-on-demand, and various non-multimedia mission-critical applications, and such applications may proliferate in the future. The diff-serv mechanisms that are deployed must be able to interoperate effectively with hosts and networks that provide per-flow QoS using int-serv models.

There are several different models for coexistence and interoperation between RSVP/int-serv and diff-serv. This draft is primarily concerned with one important model, although Section 5 presents a brief look at other models. Under our model, diff-serv mechanisms are used within transit networks in the `core' of the network, while RSVP/int-serv mechanisms are used within stub networks at the 'edge'. From the int-serv viewpoint, the diff-serv transit network is treated as a virtual link connecting int-serv/RSVP capable routers. This model builds upon work in progress on RSVP aggregation [8, 15].

This model will provide a framework that will allow deployment of diff-serv networks and deployment of RSVP/int-serv networks to proceed at their own pace, providing immediate incremental benefits in areas of the network in which one or the other is deployed and additional benefits where both are deployed. Ultimately, we want RSVP/int-serv and diff-serv mechanisms to interact seamlessly. Network administrators should be able to determine for their own networks the degree to which diff-serv capabilities are pushed towards the edge of their networks, or the degree to which RSVP/int-serv capabilities are pushed towards the core of the Internet.

Section 2 provides an overview of our model for interoperation between int-serv and diff-serv, and discusses some of the assumptions. Section 3 presents the model in more detail, while Section 4 discusses its implications for diff-serv. Finally, Section 5 briefly lists some other possible models for interoperation. Appendix A contains a list of some important terms used in this document.

Even though one of the goals of this draft is to describe a framework for co-existence of RSVP/int-serv with diff-serv, the draft currently does not address the issues specific to IP multicast flows. See Section 5.

2. Overview of the Model

This section examines the issues in providing tight quantitative end-to-end QoS over end-to-end paths that includes both int-serv networks and diff-serv networks, and introduces our model.

2.1 Quantitative End-to-End QoS

The primary focus of this document on end-to-end quantitative QoS. Although quantitative QoS applications may generate only a small fraction of all traffic, servicing this traffic may comprise a significant fraction of the revenues associated with QoS. In addition, while qualitative QoS applications can be satisfied by conventional diff-serv alone, quantitative QoS applications require additional support.

Diff-serv is expected to define some well-defined edge-to-edge services, which will be formed by concatenation of the `per-hop-behaviors' (PHBs [10]) that are being defined for internal diff-serv routers, possibly with some defined shaping and/or policing at the ingress. Our model assumes that it will be possible to map the quantitative QoS services defined by int-serv into these diff-serv services within the diff-serv network, in order to satisfy the end-to-end requirement of a quantitative QoS application.

2.2 Packet Marking

Within the diff-serv regions of the network, traffic is allotted service based on the contents of the DS-field in the IP headers. Setting the DS-field is referred to as `marking' the packet. QoS applications must be able to effect the correct marking of DS-fields before their packets enter a diff-serv network. There are two choices for accomplishing this.

Host Marking
: Hosts may directly mark DS-fields in the packets transmitted by QoS applications. Such marking may be based on host configuration or on local communication between QoS applications and the host operating system.

Int-serv Router Marking
: Routers external to the diff-serv network may mark DS-fields on behalf of QoS applications, based on MF classification. The MF classifier might be dynamically configured by RSVP signaling between QoS applications, or it might be controlled statically by manual configuration or automated configuration scripts.

MF classification is expected to be limited to examination of the network and transport-layer (port) fields of a packet. An advantage of host marking is that it allows marking to depend upon application-specific information that cannot be deduced from MF classification. For example, consider the need to give

preferential service to a website's home page (over other, less important pages at the site) or the case of encrypted traffic flows (IPSEC).

The information required to express useful mappings of application traffic flows to service levels is likely to be quite complex and to change frequently. Thus, manual configuration is likely to impose a significant management burden. If the configuration information is very simple and does not change over time, the management burden may be relatively minor; however, this means that the granularity of allotting service levels to flows will be sub-optimal. These considerations argue for host-based marking or for dynamic configuration of a router's classifier/marker in response to application requests.

2.3 Granularity of Allotment

The term `granularity' is used here to refer to the degree of specificity that is available in allotting a specific service level to a specific traffic flow. There are two measures of allotment granularity: granularity of packet classification and temporal granularity.

Fine grain classification might implement the following requirement: "Telephony traffic from user X is allotted service level A, while telephony traffic from user Y is allotted service level B, and web traffic from any user is allotted service level C." Coarse grain classification might be limited to something of the form: "All traffic from subnet 1.0.0.0 receives service level A, while all traffic from subnet 2.0.0.0 receives service level B."

Temporal granularity determines the frequency with which the service allotted to a flow may change. A temporally fine grain system would allow immediate changes in allotment of service levels to traffic flows, with times of the order of seconds or less. A temporally coarse-grained system might have service levels set by static provisioning with time frames of days or weeks.

2.4 Policing

It may be necessary to protect the network by policing at various points, within the stub network and/or at the interface to the transit network. For example, within the stub network, first-hop routers may police the aggregate traffic coming from a host to ensure that the host is not exceeding its traffic limit.

It should be noted that some diff-serv PHBs (e.g., a "billing" PHB [14]) may not require any policing at all at any point in the network.

2.5 Admission Control

Under RSVP/int-serv, quantitative QoS applications use RSVP to submit QoS requests to explicit admission control at each hop of the end-to-end path. Integrated Services admission control (ISAC) may be avoided only on hops that are known to be sufficiently over-provisioned to satisfy the service requirements. When a request is rejected, the host application may choose to try again with a smaller request or to accept the best-effort service available everywhere along the path, or it may simply avoid sending traffic. These mechanisms protect traffic on flows that were previously admitted.

In diff-serv regions of the network, admission control may be provided implicitly by policing at ingress points, based on provisioning. However, to support end-to-end tight QoS, explicit admission control must be applied to the virtual hop represented by the diff-serv transit network. An diff-serv service level used by the int-serv traffic is provisioned for some maximum level of traffic. The admission control function must ensure that this limit is not exceeded by the total int-serv traffic submitted for this class.

2.6 Policy Control

QoS support provides preferential treatment to particular traffic flows. As a result, admission control must be based on policy as well as on resource availability.

In an int-serv network, resource-based admission control is handled by RSVP enabled routers (and SBMs [2]), and is typically at the granularity of individual users. Policy based admission control is handled by RSVP-capable policy servers. These assure that int-serv network resources are allotted to users according to policy specific to the int-serv network. In addition, policy servers within the int-serv network must assure that appropriate policy is applied when diff-serv resources are allotted to int-serv users.

In a diff-serv network, resource and policy-based admission control are handled by entities such as bandwidth brokers. Policy decisions made within the diff-serv network are likely to be at the granularity of peer networks. In general, the diff-serv network may make multiple service levels available to a single

peer int-serv network.

3. Description of Model

We envision an internet that consists of RSVP/int-serv capable stub networks interconnected by diff-serv capable transit networks. The simplest example of this model is a diff-serv capable transit network and two RSVP/int-serv capable stub networks, as shown in Figure 1. The transit network contains a mesh of routers, at least some of which are diff-serv capable. The stub networks contain meshes of routers, at least some of which are int-serv capable.

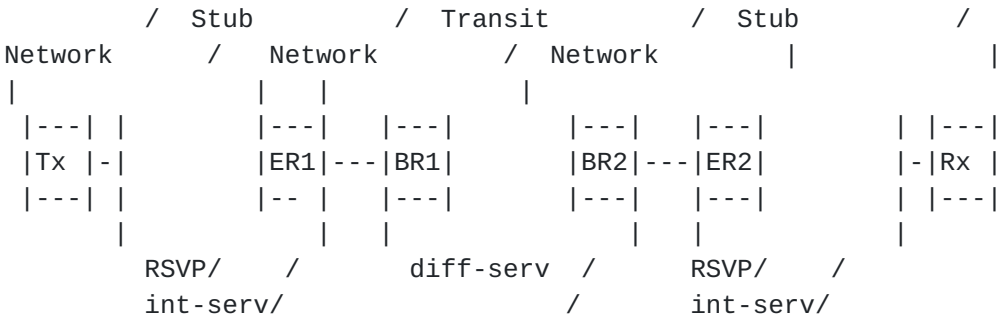

Figure 1: Sample Network Configuration

In the interest of simplicity, Figure 1 shows a single QoS sender Tx on one of the stub networks and a single QoS receiver Rx on the other. The edge routers (ER1, ER2) within the RSVP/int-serv networks interface to the border routers (BR1, BR1) of the diff-serv network.

From an economic viewpoint, we may consider that the transit network sells service to the stub networks, which in turn sell service to end systems. Thus, we may think of the stub networks as customers of the transit network. In the following, we use the term "customer" for each of the stub networks in Figure 1.

3.1 Components of the Model

We now define the major components of the proposed model.

3.1.1 Hosts

Both sending and receiving hosts use RSVP to communicate the quantitative QoS requirements of QoS-aware applications running on the host. Typically, a QoS process within the host operating system generates RSVP signaling on behalf of the applications; this process may also invoke local traffic control.

Traffic control in the host may mark the DS-field in transmitted packets, and it may shape transmitted traffic to the requirements of the int-serv service in use. Alternatively, the first-hop router within the int-serv network may provide these traffic control functions.

3.1.2 End-to-End RSVP Signaling

We assume that RSVP signaling messages travel end-to-end between hosts Tx and Rx to support int-serv reservations in the stub networks. We require that these end-to-end RSVP messages be tunneled transparently across the diff-serv transit network. Mechanisms for this purpose are proposed in [8]; they do not require the routers in the transit network to understand/interpret RSVP messages and do not adversely impact the transit network.

3.1.3 Edge Routers

We choose to place the boundary between the RSVP/int-serv region and the diff-serv region of the network within the edge routers. It is helpful to think of an edge router as consisting of two halves: a standard RSVP half, which interfaces to a stub network, and a diff-serv half, which interfaces to the transit network. The RSVP half has full RSVP capability. It is able to do MF classification, if required, and it is able to police traffic that will be passed to the border router.

The diff-serv half of the edge router provides an interface to the diff-serv network's admission control function, which we refer to as as `DSAC' (Diff-Serv Admission Control).

The customer(s) (owner(s) of the stub networks) and the carrier owning the transit network will negotiate a contract for the capacity to be provided at each of a number of standard diff-serv service levels. If the service agreement between the stub networks and the transit networks is statically provisioned, then the DSAC can be simply based upon a table that specifies capacity at each service level. If the service agreement is dynamic, the DSAC may communicate with counterparts within the diff-serv network (such as a bandwidth broker [4]) in order to make admission control decisions based on provisioned limits as well as the topology and the capacity of the diff-serv network.

Since the individual int-serv flows are policed according to int-serv rules within the stub network, the edge router needs to shape only the aggregate stream, not the individual flows.

3.1.4 Border Routers

BR1 and BR2 are diff-serv capable border routers, and are not required to run RSVP. They are expected to implement the policing function of diff-serv ingress routers, based on the results of a BA classifier. They are required neither to provide MF classification nor to mark the DS-field (with the possible exception of differential marking to indicate out-of-profile traffic [10, 8]).

3.1.5 Stub Networks

A stub network consists of int-serv capable hosts and some number of routers. These routers may reasonably be assumed to be RSVP/int-serv capable, although this might not be required for a small over-provisioned stub network. If they are not int-serv capable, we assume that they are not capable of per-flow classification, signaling, or admission control and will pass RSVP messages unhindered.

3.1.6 Transit Network

The transit network is not typically capable of per-flow classification, signaling, and admission control. It provides two or more levels of service based on the DS-field in the headers of carried packets (diff-serv capable). Furthermore, the transit network is able to carry RSVP messages transparently, with minimal or no impact on its performance (see [8]). The transit network may include multiple carrier networks.

3.1.7 Service Mapping

RSVP signaling requests carry an int-serv service type and a set of quantitative parameters known as a "flowspec"; these describe the QoS expected from the int-serv regions of the network. At each hop in an int-serv network, the generic int-serv service requests are interpreted in a form meaningful to the specific link layer medium. For example, at an ATM hop, a VC of the correct type (CBR, ABR or VBR) is established [13]. At an 802.1 hop, the int-serv service type is mapped to an appropriate 802.1p priority level [5].

In our model, the entire diff-serv network plays the role of a single virtual link layer as far as RSVP/int-serv are concerned. Therefore, the int-serv service request must be mapped to the DS-field when the packets enter the diff-serv cloud. The requested int-serv service must be mapped to a

diff-serv service level that can reasonably extend the int-serv service type requested by the application. The edge router can then provide admission control to the diff-serv network by accepting or rejecting the request based on the capacity available at the requested diff-serv service level.

One of two schemes may be used to map int-serv service types to diff-serv service levels.

Default

In this scheme, there is some standard, well-known mapping from int-serv service type to a PHB that will invoke the appropriate behavior in the diffserv network.

To improve the quality of the mapping, it may prove necessary to add additional information to an int-serv QoS request. For example, consider QoS requests for two different flows, one interactive voice traffic and the other latency-tolerant traffic. They may both have the same int-serv parameters (especially using the Controlled Load service), but they are likely to map to different diff-serv services. For this reason, we suggest adding a qualifier to the int-serv service type indicating its relative latency tolerance (high or low). The qualifier would be defined as a standard object in int-serv signaling messages.

Customer-Specified

In this scheme, the edge routers in the customer (stub) networks are allowed to modify the service mapping. RESV messages originating at hosts will carry the usual int-serv service type (perhaps with a qualifier, as described above). When a RESV message arrives at the edge router from which the traffic enters the diff-serv region (e.g., router ER1 in Figure 1), the edge router determines the PHB code point that should be used to obtain the corresponding diff-serv service level. This information is appended to the RESV message by ER1 and carried to the sending host. When the RESV message arrives at the sending host, the sender (or intermediate int-serv routers) start marking outgoing packets with the indicated PHB code point.

A decision to override the well-known service mapping at the edge router may be based on configuration and/or a policy decision. For example, when a reservation request arrives at

the ingress to a diff-serv network, if accepted reservations have already reached the pre-negotiated capacity limit at the corresponding service level then the edge router may decide to use a PHB corresponding to a different service level, based on an administratively-set policy.

3.2 Example: Obtaining End-to-End QoS

The following sequence illustrates the process by which an application obtains end-to-end QoS.

1. The QoS process on the sending host Tx generates an RSVP PATH message that describes the traffic offered by the sending application.

2. The PATH message is carried toward the receiving host Rx. In the sender's stub network, standard RSVP processing is applied at RSVP capable nodes (routers, SBMs, etc).

3. At the edge router ER1, the PATH message is subjected to standard RSVP processing and PATH state is installed in the router. The PATH message is sent onward, to the transit network.

4. The PATH message is carried transparently through the transit network, and then processed in stub router ER2 according to standard RSVP processing rules.

5. When the PATH message reaches the receiving host Rx, its QoS process generates an RSVP RESV message, indicating interest in the offered traffic at a certain int-serv service level.

6. The RESV message is carried back towards the sending host. Consistent with standard RSVP processing, it may be rejected at any RSVP node in the receiver's stub network if resources are deemed insufficient to carry the traffic requested.

7. At ER2, the RESV message is subjected to standard RSVP processing. It may be rejected if resources on the downstream interface of ER2 are deemed insufficient to carry the resources requested. If it is not rejected, it will be carried transparently through the transit network, arriving at ER1.

8. In ER1, the RESV message triggers DSAC processing. The DSAC compares the resources requested to the resources available at the corresponding diff-serv service level, in the diff-serv enabled transit network. If the RESV message is

admitted, the DSAC updates the available capacity for the service class, by subtracting the approved resources from the available capacity.

9. Assuming the available capacity is sufficient, the RESV message is admitted and is allowed to continue upstream towards the sending host. If the available capacity is insufficient, the RESV message is rejected and the available capacity for the service class remains unchanged.

10. The RESV message proceeds through the sender's stub network. RSVP nodes in the sending stub network may reject it. If it is not rejected, it will arrive at the sender host Tx.

11. At Tx, the QoS process receives the RESV message. It interprets receipt of the message as an indication that the specified traffic has been admitted for the specified int-serv service type (in the RSVP enabled regions of the network) and for the corresponding diff-serv service level (in the diff-serv enabled regions of the network).

    Tx begins to set the DS-field in the headers of transmitted packets to the value which maps to the Intserv service type specified in the admitted RESV message.

In this manner, we obtain end-to-end QoS through a combination of networks that support RSVP style reservations and networks that support diff-serv style prioritization. The successful arrival of RESV messages at the original sender indicates that admission control has succeeded both in the RSVP regions of the network and in the diff-serv regions of the network.

3.3 Variations of the Model

It is useful to consider some variations of the model just presented.

3.3.1 Moving the Boundary

We have assumed that the boundary between the RSVP/int-serv network and the diff-serv network lies in the edge routers. Alternative models could shift this boundary to the left or to the right in Figure 1. It could for example, be placed within the border routers in the transit network. In this case, the DSAC would be implemented entirely within the diff-serv network (and would essentially be the bandwidth broker proposed in [4]); however, it would require that the diff-serv border routers be RSVP capable.

Alternative, the boundary could be shifted all the way to the end hosts. This would mean that the traffic was using diff-serv mechanisms in the stub networks as well as the transit network, while the int-serv mechanisms would be only in the host. The QoS-aware application could still use RSVP within the lost to signal its needs. The host would implement per-flow policing, the DSAC function, and packet marking.

3.3.2 Service Agreements

- Statically-Provisioned Service Agreements

  In the simplest model, service agreements are negotiated statically between stub networks and transit networks. A service agreement consists of a table of capacities available to a stub network, at each diff-serv service level. In this case, DSAC functionality is provided at the edge routers in the stub networks. The `diff-serv half' of these routers appear to the `RSVP half' as a sending interface with resource limits defined by the service agreement table. While there may be bandwidth brokers and dynamic provisioning within the transit networks, these are not coupled with the int-serv stub networks, and admission control in the two regions of the network is completely independent.

- Dynamic Service Agreements

  In a more sophisticated model, service agreements between customer stub networks and carrier transit networks are more dynamic. Customers may be able to dynamically request changes to the service agreement. In this case, a statically provisioned edge router cannot provide the required DSAC functionality. Instead, DSAC functionality must be provided by coupling the stub network's admission control with the transit network's admission control.

  The two admission control mechanisms meet at the boundary between the diff-serv network and the int-serv network. This boundary may be implemented at the edge router (in the stub network), at the border router (in the transit network), or at the bandwidth broker for the int-serv network.

  Note that coupling int-serv and diff-serv admission control does not imply that each int-serv admission control request will result in DSAC processing. Int-serv admission control requests may be aggregated at the

boundary between the int-serv and the diff-serv network. For example, int-serv admission control requests may trigger DSAC requests to bandwidth brokers only when some high-water or low-water resource threshold is crossed. Separate high-water and low-water thresholds can provide hysteresis to prevent thrashing.

%.cm In the latter case, any MF classification on %.cm behalf of the diff-serv ingress point is provided as a service to the %.cm customer and goes beyond policing requirements).

3.3.3 Setting the DS-field

Allowing hosts to set the DS-field directly may alarm network administrators. The obvious concern is that hosts may attempt to 'steal' resources. In fact, hosts may attempt to exceed the negotiated capacity at a particular service level regardless of whether they invoke this service level directly (by setting the DS-byte) or indirectly (by submitting traffic that classifies in an intermediate router to a particular diff-serv PHB).

In summary, the security concerns of marking the DS-field at the edge of the network can be dismissed since each carrier will have to do some form of policing (per-flow or per-host) at their boundary anyway. Furthermore, this approach reduces the granularity at which border routers must police, thereby pushing finer grain shaping and policing responsibility to the edges of the network, where it scales better. The carriers are thus focused on the task of protecting their transit networks, while the customers are focused on the task of shaping and policing their own traffic to be in compliance with their negotiated token bucket parameters.

It is also possible to mark the DS-field at intermediate routers rather than at the host and still realize many of the benefits of our approach. In this case, intermediate routers may use the RSVP signaling to configure an MF classifier and marker. Then the configuration of MF classifiers and markers would be dynamic (minimizing the management burden), and full resource- and policy-based admission control could be applied.

The disadvantages of marking the DS-field at intermediate routers (instead of the host) are that full MF classifiers are required at the intermediate nodes and that responsibility for

traffic separation is shifted away from the host.

Nonetheless, marking at intermediate routers will be necessary to support those hosts which support RSVP signaling but are incapable of marking the DS-field. In addition, there may be cases in which the network administrators wish to shift the responsibility for traffic separation away from the hosts. In particular, we expect that there will continue to be a need for top-down provisioned MF classification, especially for qualitative (as opposed to quantitative) QoS applications. See Section 5.2.

4. Implications for Diff-Serv

We have described a framework for end-to-end QoS in which a diff-serv network can be included as a segment of an int-serv path. This section discusses some of the implications of this framework for diff-serv.

4.1 Requirements for Diff-Serv

In order to use a diff-serv network as described in this draft, the diff-serv network must satisfy the following requirements.

1. A diff-serv network must be able to provide standard QoS services between its border routers, and such a service must be selectable by specifying a standard code in a (PHB) sub-field of the DS-field of a packet.

2. There must be appropriate service mappings from int-serv services into these diff-serv services.

3. Diff-serv networks must provide admission control information to the int-serv network. This information can be provided by a dynamic protocol or, at the very least, through static service level agreements.

4. Diff-serv networks must be able to transparently carry RSVP messages, in such a manner that they can be recovered at the egress point from the diff-serv network.

4.2 End-to-End Integrity of the DS-field

Our model assumes that int-serv networks uses a code point of the DS-field in order to specify the desired PHB within the transit network. It also assumes that packets submitted to the transit network specifying a certain DS-field will receive a standard service throughout the transit network. Strictly speaking, this

does not dictate that the transit network must leave the DS-field field intact. For example, the border router may map a DS-field value set by the host or edge router to a different value before forwarding the data packets.

However, we see little value in allowing the PHB field to be altered within the network. This is likely to perpetuate local and incompatible interpretations of the field.

4.3 Policing and Shaping

We are assuming that border routers will police in aggregate. As a result, the customer cannot rely on border routers to provide traffic isolation between the customer's flows, when policing or shaping. Instead, it is the customer's responsibility to ensure that the customer's flows are properly shaped and policed within the customer's sending network. Overall, this approach simplifies border routers and still allows protection against misbehaving hosts/users.

Ideally, hosts should provide per-flow shaping at their sending interfaces. However, there is always a chance that the customer's traffic will become distorted as it nears the boundary between the customer and the carrier. In this case, the customer should do per flow policing (or even re-shaping) at the egress point from the customer's network unless the policing agreement at the other side is known to accommodate the transient bursts that can arise from adding the flows.

4.4 Managing Resource Pools

Network administrators must be able to share diff-serv network resources between three types of traffic:

a. Quantitative (explicitly signaled) QoS application traffic

b. Qualitative (implicitly signaled) QoS application traffic

c. All other (best-effort) traffic

These pools must be isolated from each other by the appropriate configuration of policers and classifiers at ingress points to the diff-serv network, and by appropriate provisioning within the diff-serv network. To provide protection for quantitative QoS traffic in diff-serv regions of the network, we suggest that the DS codepoints allotted to such traffic must not overlap the codepoints assigned to other traffic (qualitative QoS and best-

effort traffic).

5. Other Models

5.1 RSVP and Diff-Serv

Since RSVP was originally designed to support int-serv, we use the term "RVSP/int-serv" as the complement to diff-serv. However, RSVP and int-serv are separable, and RSVP may be used as a general-purpose QoS signaling protocol. For example, RSVP might be used for dynamic provisioning and admission control in the diff-serv region of the network. Routers in the diff-serv region would continue use the DS-field in the IP header to identify and offer services to flow aggregates.

5.2 Qualitative QoS

This document has focused largely on the class of applications that use RSVP to explicitly signal per-flow QoS requirements and that expect end-to-end tight QoS assurances. We have been referring to these applications as `quantitative QoS applications'. Suitable end-to-end services must also be available to qualitative QoS applications. The services that these applications require are generally less demanding.

Qualitative services can be obtained from the diff-serv regions of the network with loose top-down provisioning. Network managers can configure classifiers at the ingress to the diff-serv network to recognize traffic requiring specific qualitative service levels. Thus, these classification fields are used as a form of implicit signaling. In the int-serv portion of the network, qualitative QoS applications can get best-effort service, which may be good enough.

There would be no explicit admission control for such qualitative QoS applications. Therefore, it is difficult to assure that the total traffic offered at an ingress point will not exceed the provisioned capacity for a particular service level. When the traffic exceeds the allowed limit, there is only indirect feedback to the applications, in the form of packet loss or an Congestion Experienced mark from Explicit Congestion Notification (ECN). Thus, traffic from qualitative applications can be offered only loose QoS.

5.3 Multicasting

<To be written>

6. Security Considerations

We are proposing that RSVP signaling be used to obtain resources in both the diff-serv and int-serv regions of the network. Therefore, all RSVP security considerations apply [11]. In addition, network administrators are expected to protect network resources by configuring secure policers at interfaces with untrusted customers.

7. Acknowledgments

Authors thank the following individuals for their comments that led to improvements to the previous version(s) of this draft: David Oran, Andy Veitch, Curtis Villamizer, Walter Weiss, and Russel white.

Many of the ideas in this document have been previously discussed in the original int-serv architecture document [12].

8. References

[1] Braden, R., Zhang, L., Berson, S., Herzog, S. and Jamin, S., "Resource Reservation Protocol (RSVP) Version 1 Functional Specification", RFC 2205, Proposed Standard, September 1997

[2] Yavatkar, R., Hoffman, D., Bernet, Y., Baker, F. and Speer, M., "SBM (Subnet Bandwidth Manager): A Protocol For RSVP-based Admission Control Over IEEE 802 Style Networks", Internet Draft, March 1998

[3] Berson, S. and Vincent, R., "Aggregation of Internet Integrated Services State", Internet Draft, December 1997.

[4] Nichols, K., Jacobson, V. and Zhang, L., "A Two-bit Differentiated Services Architecture for the Internet", Internet Draft, December 1997.

[5] Seaman, M., Smith, A. and Crawley, E., "Integrated Services Over IEEE 802.1D/802.1p Networks", Internet Draft, June 1997

[6] Elleson, E. and Blake, S., "A Proposal for the Format and Semantics of the TOS Byte and Traffic Class Byte in Ipv4 and Ipv6 Headers", Internet Draft, November 1997

[7] Ferguson, P., "Simple Differential Services: IP TOS and Precedence, Delay Indication, and Drop Preference", Internet Draft, November 1997

[8] Guerin, R., Blake, S. and Herzog, S.,"Aggregating RSVP based QoS Requests", Internet Draft, November 1997.

[9] Nichols, Kathleen, et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[10] Blake, S., et al., "An Architecture for Differentiated Services." RFC 2475, December 1998.

[11] Baker, F., "RSVP Cryptographic Authentication", Internet Draft, August 1997

[12] Braden, R., Clark, D. and Shenker, S., "Integrated Services in the Internet Architecture: an Overview", Internet RFC 1633, June 1994

[13] Garrett, M. W., and Borden, M., "Interoperation of Controlled-Load Service and Guaranteed Service with ATM", RFC2381, August 1998.

[14] Weiss, Walter, Private communication, November 1998.

[15] Berson, S. and Vincent, S., "Aggregation of Internet Integrated Services State", Internet Draft, August 1998.

APPENDIX A. Terminology

The following terms were used in this draft.

Int-serv
: The part of an internet that uses per-flow classification, signaling, and admission control to deliver per-flow QoS guarantee.

[Diff-serv region (or diff-serv capable network)] The part of an internet that provides aggregate QoS services

Quantitative
: Application for which QoS requirements are readily quantifiable, and which relies on these QoS requirements to function properly.

Qualitative
: Applications for which relative, but not readily quantifiable, QoS requirements exist.

QoS Application that requires some form of QoS, either qualitative or quantitative.

LooseQoS assurances that are relative, rather than absolute, or generally not quantifiable.

TightQoS assurances which are quantifiable, though they may or may not provide 100% guarantee.

Top-down
Traditional provisioning methods that configure network capacities using heuristics and experience, typically from a console, based upon traffic predictions.

Author's Addresses

Yoram Bernet
Microsoft
One Microsoft Way,
Redmond, WA 98052
Phone: (425) 936-9568
Email: yoramb@microsoft.com

Raj Yavatkar
Intel Corporation, JF3-206
2111 NE 25th. Avenue,
Hillsboro, OR 97124
Phone: (503) 264-9077
Email: raj.yavatkar@intel.com

Peter Ford
Microsoft
One Microsoft Way,
Redmond, WA 98052
Phone: (425) 703-2032
Email: peterf@microsoft.com

Fred Baker
Cisco Systems
519 Lado Drive,
Santa Barbara, CA 93111
Phone: (408) 526-4257
Email: fred@cisco.com

Lixia Zhang
UCLA
4531G Boelter Hall
Los Angeles, CA 90095
Phone: +1 310-825-2695
Email: lixia@cs.ucla.edu

Kathleen Nichols
Cisco Systems
Email: kmn@cisco.com

Michael Speer
Sun Microsystems, Inc
901 San Antonio Road UMPK15-215
Palo Alto, CA 94303
phone: +1 650-786-6368
Email: speer@Eng.Sun.COM

Bob Braden
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
phone: 310-822-1511
Email: braden@isi.edu