Diffserv Working Group
INTERNET-DRAFT

A. Conta (Transwitch)
J. Rajahalme (Nokia)
November 2001

A model for Diffserv use of the IPv6 Flow Label Specification

draft-conta-diffserv-ipv6-fl-classifier-01.txt

Status of this Memo

Abstract

This document specifies a conceptual model for using IPv6 flow labels with Differentiated Services. It also specifies an IPv6 flow label classifier for Diffserv, and a set of rules for using the IPv6 flow label with Diffserv.

Table of Contents

1. Introduction

This document specifies a conceptual model for using IPv6 flow labels with Differentiated Services [Diffserv]. It also defines a IPv6 flow label classifier for Differentiated Services (Diffserv), and a set of rules for the use of the IPv6 flow label with Diffserv. It also suggests an API to be used to set/get IPv6 flow label values for Diffserv. Ultimately, the document provides the specifics of the mechanisms for using IPv6 flow labels with Differentiated Services.

The use of the IPv6 flow label with Diffserv will help achieve, as the flow label is supposed, a more efficient processing of packets in Diffserv quality of service engines in IPv6 forwarding devices.

The keywords MUST, MUST NOT, MAY, OPTIONAL, REQUIRED, RECOMMENDED, SHALL, SHALL NOT, SHOULD, SHOULD NOT are to be interpreted as defined in [KEYWORDS].

2. Diffserv use of IPv6 Flow Label - Conceptual Model

The Differentiated Services QoS model (Diffserv) can be used anywhere in a network. A particular case is the use of Diffserv in a IPv6 access network. In such a case, Diffserv would provide Quality of Service (QoS) functions for IPv6 traffic generated by users and carried by the access network towards the final destinations. It would also provide QoS functions for IPv6 traffic generated by remote sources and carried by various networks, including the access network, which is the last hop network. Part of these QoS functions are those that are typically employed at the edge of a Diffserv network. In fact, the nature of the contractual agreements between the users and the access network providers - service level and traffic conditioning agreements (SLAs and TCAs) -- are such that Diffserv could be easy to use, more efficient, and practical than other models. Obviously, in such a case, the Diffserv QoS functions would begin or would include at the edge routers traffic classification, which could be multi-field packet classification. The IPv6 flow label can play a major role in this multi-field packet classification. The IPv6 flow label, along with other fields in the main IPv6 header, can constitute the elements of the classifiers that are used to differentiate packets belonging to various traffic flows. We call such a classifier a "Diffserv IPv6 flow label classifier".

The "Diffserv IPv6 flow label classifier" is basically a 3 element tuple: source and destination IPv6 addresses, and the IPv6 flow label. The flow label classifier is an alternative to the 5 element tuple (addresses, ports, and protocol). It provides, higher

efficiency for packet classification in quality of service engines in forwarding devices. This is particularly so in the case of IPv6 packets, because it is eliminating the examining or lookup of extension and transport headers during the multi-field packet classification.

As a key element of the IPv6 flow label classifier, the IPv6 flow label value, is matched along with the values of the other elements of the classifier against the packet headers fields.

The value of the IPv6 flow label that is set in a IPv6 header can be chosen by various means (further discussed later).

However, in the case of Diffserv, whatever the means used to chose a value, the "flow_label value" or range of values MUST be known, and agreed by two sides:

- the network client/user, which needs, pays and expects certain QoS from the network, for the traffic it sends into the network, and

- the network provider, which implements mechanisms to satisfy the user's needs regarding the user's traffic sent into and carried by the provider's network.

On the network user side, "flow label values" are set by traffic sources and carried in the IPv6 flow label field of the IPv6 packet headers. We call these "host flow label values".

On the network provider side, "flow label values" of flow label classifiers are configured into the providers' routers that forward the user's traffic into the network. We call this "router flow label values".

These flow label values are captured in contractual agreements between users and network providers, the so called Service Level and Traffic Conditioning Agreements and Specifications (SLAs, SLSs, TCAs, TCSs).

For a further description and understanding of the mechanisms employing the IPv6 Flow label for Diffserv, we are considering a Diffserv conceptual model, which can be further broken apart into two conceptual models, one for entities that are the source and destination of packets, i.e. hosts, and one for entities that are forwarding the packets, i.e. routers:

- the conceptual model for the use of the IPv6 flow label with Diffserv in a host, and

- the conceptual model for the use of the IPv6 flow label with Diffserv in a router.

2.1. Host Conceptual Model for the Diffserv Flow Label

The conceptual model for the use of the IPv6 flow label with Diffserv in a host describes the mechanisms that are implemented and employed in hosts. The hosts are the source and final destination of the IPv6 packets. These mechanisms select what we called above the "host flow label values", and then fill in those selected values into the main IPv6 headers of packets that are sent into the network.

For the use of the IPv6 flow labels with Diffserv, there is one major requirement for selecting the values to set into the flow label: the values must be conforming to the contractual agreements between users and network providers (SLAs, SLSs, TCAS, TCSs). The selecting of the values can be done by any possible means as long as the above requirement is followed strictly. As examples, we have considered the following possible means for selecting the host flow label values:

a. arbitrary number,
b. random number,
c. IANA number of some sort,

A further discussion of these follows:

2.1.1. Selecting the Host Flow Label Value

a. As an "arbitrary number", the "host flow label value" can be any value between 1 and the maximum value allowed for a IPv6 flow label used by Diffserv. If the arbitrary selected flow label value is intended for use with Diffserv in traffic sent over a set of networks, that value must be specified in the contractual agreements between the user and the network providers that will carry the user's traffic. Obviously, the arbitrary value will be used, and will not change for a particular flow or set of flows, as long as the contractual agreements are valid. This can be days, weeks, months or longer. The arbitrary value can be stored on the host, in a system, group, individual user, or application data base. It can be also stored in a network distributed data base.

b. As a "random number", the "host flow label value" can be any value between 1 and the maximum value allowed for a IPv6 flow label. Once the random generated number is selected as a flow label value, as it is intended for use with Diffserv in traffic sent over a set of networks, it must be specified in the contractual agreements between

the user and the network providers that will carry the user's traffic. Obviously, the randomly selected value will be used, and will not change for a particular flow or set of flows, as long as the contractual agreements are valid. This can be days, weeks, months or longer. The random number value can be stored on the host, in a system, group, individual user, or application data base.

c. As a "IANA number", the "host flow label value" can be any value between 1 and the maximum value allowed for the IPv6 flow label. Since it is a IANA number, it is set and held by IANA, in association with a certain flow or flows, characterized by various elements that are carried in packet headers, such as addresses, protocol identifiers, ports, etc., and the type of application or applications exchanging packets of that flow or those flows, or the type of service, or services which the application(s) are delivering to users. As a IANA number, this association has a permanent character. Once a IANA number has been selected as a flow label value, as it is intended for use with Diffserv in traffic sent over a set of networks, it must be specified in the contractual agreements between the user and the network providers that will carry the user's traffic. Obviously, the "IANA number" value will be used, and will not change for a particular flow or set of flows, as long as the contractual agreements are valid. The IANA number value can be stored on the host, in a system, group, individual user, or application data base.

Note: it is not in the scope of this document to specify a certain IANA number, or family of numbers. The IANA number mechanism is mentioned only as one of the possible ways in which a host flow label value can be selected.

2.1.2. Setting the Host Flow Label Value

The host flow label value selected by mechanisms described in the previous section, can be stored on a host, in a system, group, individual user, or application data base. It can also be stored in a network distributed data base. A host can fetch the host flow label value from the data base, and cache it or configure it in on-line memory:

- In the IPv6 protocol stack information base, or
- In a application information base.

A host flow label value cached in the IPv6 protocol stack information base is filled in the IPv6 main header of every packet that belongs to a flow or set of flows that the host flow label value is associated with. This association is defined by a set of elements

such as source and destination addresses, protocol identifiers, TCP, or UDP port numbers, etc... Typically, the host flow label value would be cached adjacent to the source and destination addresses in a structure such as a protocol control block, or in a flow label member of a structure pointed by the protocol control block. A host could cache also a default flow label value for a certain flow or set of flows. Normally the default value would be zero (0).

A host flow label value cached in a application information base is communicated to the IPv6 protocol stack as part of a communication setup, in case of a "connected type of communication". Typically, such a communication is a TCP connection, or a UDP communication in which the source and destination ports and source and destination addresses are setup for the entire duration of the communication, before any packet transmission takes place. In case of a "un-connected type of communication", the host flow label value can be communicated to the IPv6 protocol stack with each message that is passed to the stack. Typically, such a communication is a UDP communication in which the source address and the destination address and port are specified with each message being transmitted.

2.1.2.1. Host Flow Label Value API

The API to be used for setting a Diffserv host flow label value is the IPv6 API [Basic-Socket]. The field used to pass the value is the "sin6_flowinfo" which is a member of the IPv6 socket address structure "sockaddr_in6".

For the "connected communications", the host flow label value can be passed in a "bind", or "connect" call. The host flow label value associated with a connected communication can be extracted from the IPv6 protocol stack using a "accept", or "getpeername", and "getsockname".

For "un-connected communications", the host flow label value can be passed in a "sendto", "sendmsg" or "writev" call.

The host flow label value associated with a cun-connected ommunication can be extracted from the IPv6 protocol stack using a "recvfrom", "recvmsg", "readv", or "getpeername", and "getsockname"

2.2. Router Conceptual Model for the Diffserv Flow Label

The conceptual model for the use of the IPv6 flow label with Diffserv in a router describes the mechanisms that are implemented and employed by the routers that are forwarding the IPv6 packets in the

network towards their destinations. These mechanisms consist basically in configuring or setting up flow label classifiers (classification rules), and the classification processing done by the classification engines

2.2.1. IPv6 Flow Label Diffserv Classifier

The IPv6 Flow Label is an additional element that MAY be included in a Diffserv classifier [Diffserv]. A precise representation or expression of a Diffserv classifier, including the IPv6 flow label classifier, is given in [Diffserv-MIB], specifically "DiffservMultiFieldClfrEntry", and "DiffservMultiFieldClrfrFlowId".

A flow label classifier can be described as a tuple "C" that contains the following:

```
C = (Source Address, Source Address Prefix,
     Destination Address, Destination Address Prefix Length,
     Flow-Label)
```

Another representation of a flow label classifier can be:

```
Flow-label-classifier:
Type:                   IPv6-3-tuple
IPv6DestAddrValue:      IPv6 address
IPv6DestPrefixLength:   byte value
IPv6SrcAddrValue:       IPv6 address
IPv6SrcPrefixLength:    byte value
IPv6FlowLabel:          20 bit value
```

The values set in the fields of the classifiers (or classification rules) are strictly according to the contractual agreements between users and network providers: SLAs, SLSs, TCAs, TCSs.

The mechanisms used in a router to setup the classifiers can be manual configuration, dynamic scripts, NMS provisioning, COPS provisioning, or others. These mechanisms are beyond the scope of this specification.

The classification engines in a QoS engine or engines in routers would match packet header information, which includes the "host flow label value" to flow label classification rules, which includes the "router flow label value" as follows:

```
            Incoming packet header (Source Address,
                                    Destination Address,
                                    Flow Label)
          Match
            Classification rules table entry (C)
```

From the classification step, the Diffserv processing continues the same way as for any other MF Classifier [Diffserv-Model].

2.3. Applicability of Flow Label Classifiers - examples of use

Differentiated Services QoS model (Diffserv) can be used anywhere in a network. Three examples will be presented, in which, the use of flow label classifiers seem to be useful. In each example, Diffserv would provide Quality of Service (QoS) functions for both outgoing and incoming IPv6 traffic, but each example will discuss particularly incoming or outgoing traffic.

2.3.2.1. Example 1 - User Outgoing Traffic

The first example is an access network, and the Diffserv QoS applied to the user outgoing traffic. The outgoing traffic is generated by the direct users of the access network. The access network carries the users' traffic towards various downstream networks which will further carry the traffic towards the final destinations. Part of the QoS functions performed by the access network include those that are typically employed at the edge of a Diffserv network. The contractual agreements between the users and the access network providers -_ service level and traffic conditioning agreements (SLAs, TCAs, SLSs, TCSs) - specify the set of parameters for the QoS, as well as the parameters for particular functions. One particular function that this specification is focusing on for this example is classification.

If the users' hosts are setting the DSCP field in the IPv6 headers of the packets sent, with values that are according to the SLA/TCA/SLS/TCS, than access network Diffserv routers can perform directly DSCP based packet classification.

If the users' hosts are not setting the DSCP field, then a multi-field packet classification can be employed at the edge of the access network. The IPv6 flow label can play a major role. User hosts would set the "host flow label value" according to the SLAs/TCAs/SLSs/TCSs. The access network edge routers would be configured with flow label classification rules, which would contain "router flow label values" according to the SLAs/TCAs/SLSs/TCSs.

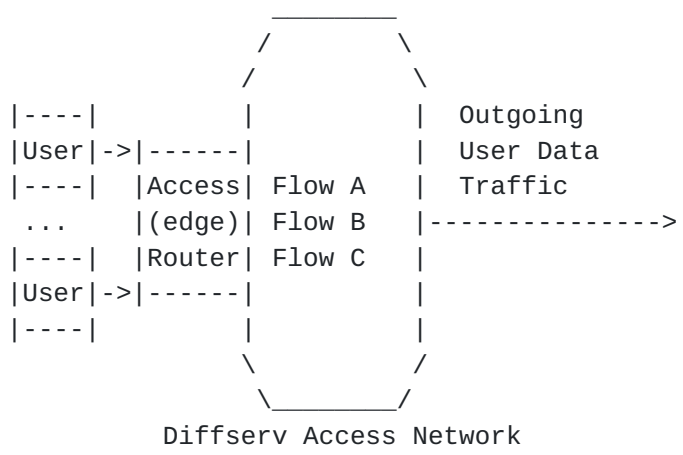

Figure 1. Use of flow label classifier in a Diffserv Network

2.3.2.2. Example 2 - Access Networks Downstream Traffic

Another example is a carrier network that performs QoS functions for downstream traffic from several access networks.

The incoming traffic is generated by direct users of access networks, which are customers of the carrier network. The carrier network, carries the users' traffic towards various downstream networks which will further carry the traffic towards the final destinations. Part of the QoS functions performed by the carrier network include those that are typically employed at the edge of a Diffserv network.

The contractual agreements between the access network providers and the carrier network provider -_ service level and traffic conditioning agreements (SLAs, TCAs, SLSs, TCSs) - are specifying the set of parameters for the QoS, as well as the parameters for particular functions, such as classification.

If the exit routers from the access networks are setting the DSCP field in the IPv6 headers of the packets sent, to values that are according to the SLA/TCA/SLS/TCS, than carrier network Diffserv routers can perform directly DSCP based packet classification.

If the access networks exit routers are not setting the DSCP field, then a multi-field packet classification can be employed at the edge of the carrier network. The carrier network edge router could be configured with flow label classification rules, which would contain "router flow label values" according to the SLAs/TCAs/SLSs/TCSs between the access networks and the carrier network. These values would be a reflection of the "router flow label values" agreed upon in the SLAs/TCAs/SLSs/TCSs between users and access network

providers, and set by users as "host flow label values" in packets.

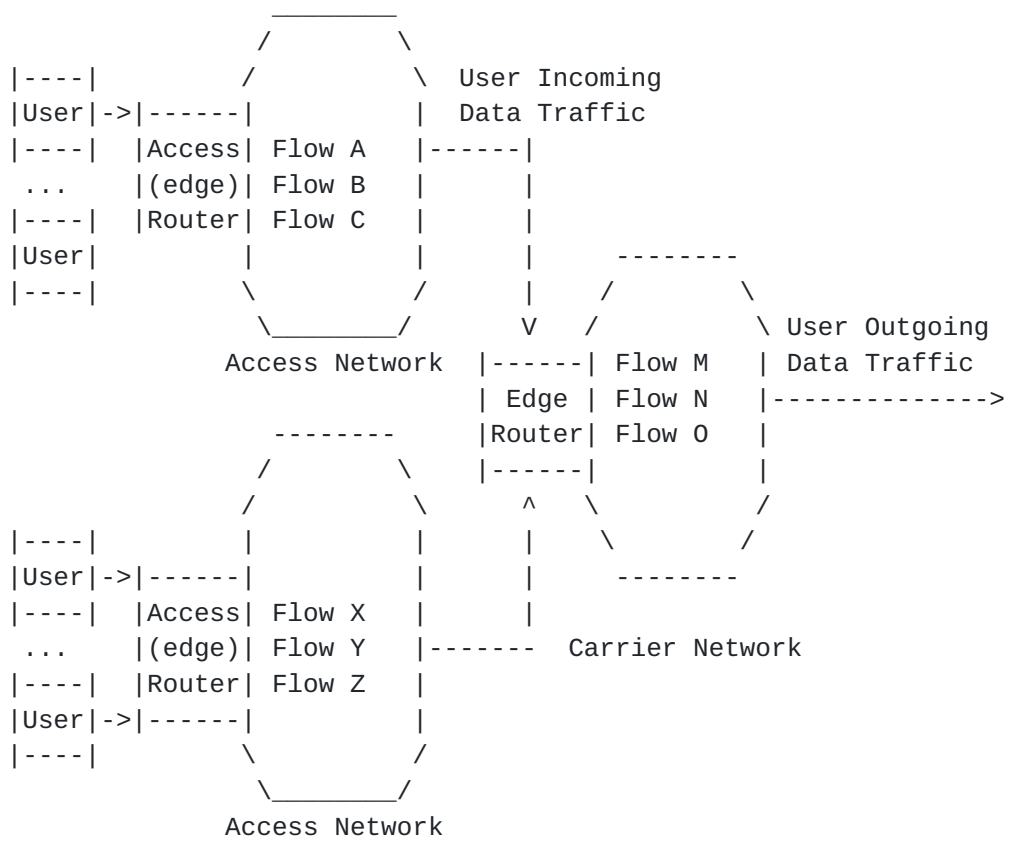

Figure 2. Use of flow label classifier in a Diffserv Network

2.3.2.3. Example 3 - Network Incoming Traffic

The last example is an access network, and the Diffserv QoS applied to the network incoming (upstream) traffic on its way to the users.

The incoming traffic is generated by remote sources and carried by various networks all the way to the access network, which is the last hop network, before the final destination, the user. As typically, the DSCP field value changes in transit, the flow label classifier can play a particularly useful role in restarting the Diffserv QoS machinery in the access network for incoming traffic, that is, differentiate packets belonging to various incoming traffic flows. A access network user, which is a final destination of the incoming traffic, can have a contractual agreement with the access network provider - service level and traffic conditioning agreement or specification (SLA, TCA, SLS, TCS) that specifies certain type of QoS for incoming traffic from certain sources, or range of sources. Please see Figure 3.

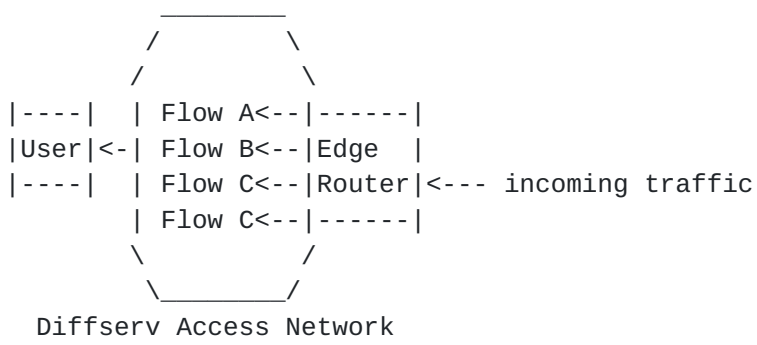

Figure 3. Use of flow label classifier in a Diffserv Network

3. Rules for using the IPv6 Flow Label with Diffserv.

The rules for using the IPv6 flow label with Diffserv are as follows:

(a) A flow is uniquely identified by the combination of source address, destination address and a non-zero flow label. Diffserv flows MAY be aggregated by specifying a range of addresses and/or a range of flow labels (see further in (e)).

(b) A flow label of zero means that the flow label has no significance, the field is unused, and therefore has no effect on, or for the packet processing by forwarding, QOS, or filtering engines.

(c) A flow label is assigned to a flow by the flow's source node. It can be changed en-route, with the condition that its original significance be maintained, or restored, when necessary. For instance if the source of the flow intended that the flow label has a certain significance to the destination end-node, than the nodes en-route, that process and eventually change the value of the flow label, should make sure, in conjunction with the destination end-node, that even when the value or significance has changed en-route, the original information and significance is restored when or before the packet arrives to its destination.

If the action to be performed on a particular flow label value in context with the packet's source and destination addresses is not known, a router MUST not change the value of that flow label.

(d) The flow label must have a value between 1 and FFFFF in hex. It is used with other fields in the header to identify a flow. It is a preset value. No particular method is preferred for choosing the value. However, the value MUST satisfy the following requirements:

It can be configured, uploaded, or transmitted to a router or a group of routers in any possible way, as long as it can be stored in the classification rules tables of the forwarding engines of routers along the path of the flow to the final destination. The flow label values are preset or agreed upon, and specified in a Service Level Agreement (SLA), Service Level Specification (SLS), Traffic Conditioning Agreement (TCA), or Traffic Conditioning Specification (TCS) [Diffserv]. This model is typical of Differentiated Services.

(e) In general, all packets belonging to the same flow are sent with the same source address, destination address, and flow label. However, flows can be trunked, or aggregated in macro-flows. The flows, members of a macro-flow, may have different source or destination addresses. The trunking, or aggregation of flows is achieved by simply wildcarding some bits or all bits in some of the fields of the flow label classification rules, which contain source address, destination address, and flow label. In other words range addresses and/or flow labels can be used.

4. The Diffserv Flow Label Conceptual Model: Conclusions

The general Diffserv flow conceptual model described in the above sections draws the very basic mechanisms. In relationship with these basic mechanisms, one important question can be raised: could the flow label numbering space be shared, regardless of which specific QoS model, Intserv, or Diffserv, a flow or aggregation of flows is using? The answer is YES, It is possible.

There is no difference in the use of TCP or UDP port numbers with Intserv and Diffserv classifiers, in that, the TCP and UDP headers, which are the Intserv classifiers' "host" elements, do not carry any information indicating that they are Intserv classifier elements, or not. Therefore it can be inferred that a model in which there is no difference between Intserv flow labels and Diffserv host flow labels is valid. This means that the IPv6 flow label numbering space can be shared by Intserv and Diffserv.

The acceptance of such a model, in which a flow label value does not

carry information regarding its use by Intserv, or Diffserv has an important consequence. It allows a certain "host flow label value" be used with both Intserv, and Diffserv, on distinct Intserv, and/or Diffserv QoS segments of a flow's or aggregation of flows' path.

Note 1: a segment is a number of interconnected routers; at minimum a segment has one router.

Note 2: It is obvious that in order to use the same host flow label value, as both Intserv, and Diffserv value, the selection of the value MUST be done in such a way that it does not preclude the use of the flow label with one model or the other. Using pseudo-random numbers that are generated on the fly and are short lived, for instance, regenerated each time an application starts execution and establishes a communication, will certainly prohibit the use of the flow label with Diffserv.

The "host flow label value" for the use with Intserv is going to be included in a "filter spec" and signaled with RSVP messages to the Intserv/RSVP routers. The Intserv/RSVP routers will include the "filter spec" in the classification rules tables used by the Intserv traffic classifiers.

The same "host flow label value" is going to be included in SLSs/TCSs, SLAs/TCAs and configured manually, or via automated scripts, or dynamically via COPS provisioning into the Diffserv routers' classification rules tables used by Diffserv flow label classifiers.

Careful consideration must be given to situations in which the same "host flow label value" in the same source and destination address context is used for both Intserv, and Diffserv on the same router, or rather on the same interface. We call this "sharing the flow label numbering space between Intserv and Diffserv". Obviously, to avoid confusion/collisions, a choice could be to not allow sharing the flow label numbering space for Intserv, and Diffserv, but that IS NOT a method that this specification recommends. If the EXACTLY same classifier rule with the same source and destination addresses, and the same flow label value is set in an SLA/TCA for Diffserv, and in the same time, the same flow label value, and source and destination addresses are signaled through RSVP for Intserv, it definitely seems possible that the router would configure in the forwarding plane only the Intserv filter rule, and restore the Diffserv rule when the Intserv flow expires or the reservation is torn down. This could be thought of as giving precedence to the most dynamic method of setting up the flow state. It is also possible that the Diffserv rule is more generic (i.e. matches address prefixes instead of complete addresses). In this case it would be possible to keep both rules in

the forwarding path classifier, but to arrange a "longest-prefix match", so that the most specific filter rule matching a given packet would take precedence.

4.1 IPv6 Flow Label in Mixed Intserv, Diffserv Networks

To illustrate the elements being discussed in the previous section, let's consider a network which includes a Diffserv QoS domain adjacent to a domain supporting the Intserv model - see Figure 4. This example has some similarities with the example give in [Intserv-Diffserv].

The Diffserv domain contains a mesh of routers, and attached hosts. Additional to Diffserv, the hosts support also RSVP/Intserv. A domain adjacent to the Diffserv domain (Intserv region) contains a mesh of routers and attached hosts, which support RSVP/Intserv.

This model network can be extended in two opposite directions. At one extreme the Diffserv domain is pushed all the way to the periphery, with hosts alone having full RSVP/Intserv capability. At the other extreme, Intserv is pushed all the way to the other end, with no Diffserv region.

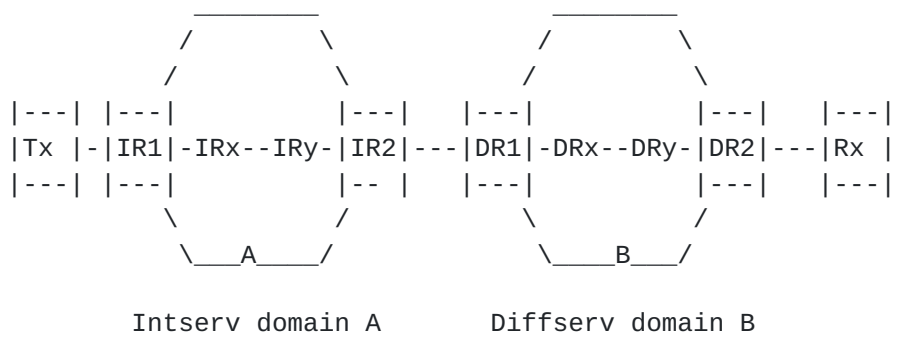

Figure 4. Intserv-Diffserv Network

Explanations to Figure 4:

For simplicity we consider a single QoS sender, Tx, which is part of the Intserv Domain A, and is communicating across the network composed of A, and B regions with a single QoS receiver, Rx, which is part of the B region, but also supports Intserv and RSVP.

The Intserv domain A is adjacent to the Diffserv domain B. Its edge router IR2 is direct neighbor to the Diffserv domain B edge router DR1. Routers IRx, IRy are core routers in domain A. IR1 is a next hop router to host TX.

The Diffserv network domain B supports aggregate (BA) traffic control in the core, in routers DRx, and DRy, and is performing MF flow label classification, policing, and shaping at the edges in routers DR1, and DR2. DR2 is a next hop router to host Rx. If devices in the Diffserv domain are not RSVP aware, they will pass RSVP messages transparently. The Diffserv network domain B provides Diffserv levels of service, to the Intserv region. Diffserv QoS information (SLSs/TCSs, SLA/TCAs) is configured in routers (DR1, DRx, DRy, DR2) manually, or via automated scripts, or dynamically using COPS provisioning.

There is no signaling between the Diffserv network domain B and network elements outside it. IR2 optionally can be configured with the information represented by the SLS/TCS, SLA/TCA and as such, it is able to act as an admission and traffic control agent for the Diffserv network domain B. Such configuration does not readily support dynamically changing SLS/TCSs, SLA/TCAs since IR2 requires reconfiguration each time the SLS/TCS, SLA/TCA changes.

Intserv service requests specify an Intserv service type, a set of quantitative parameters known as a "flowspec", and a set of identifiers for the flow (RSVP session, sender_template, filter_spec). The filter_spec is flow label based. As at each hop in the Intserv network A, the Intserv service requests are interpreted in a form meaningful to the specific link layer medium. Requests for Intserv services must be mapped onto the underlying capabilities of the Diffserv network domain B, analogous to the Intserv mapping into 802.1p capable switched segments described in [RFC 2815].

The Diffserv network domain B is statically provisioned via manual configuration of routers, via automated scripts, or dynamically using COPS provisioning. The customer(s) of the Diffserv network region, which includes network domain A, and the user owner of host Rx, and the owner/provider of the Diffserv network domain B, have negotiated static contracts (service level agreement or specification, SLA or SLS) for the transmit capacity to be provided at each of a number of Diffserv service levels. The "transmit capacity" may be simply an amount of bandwidth or it could be a more complex "profile" involving a number of factors such as burst size, peak rate, time of day etc.

The Diffserv edge routers DR1, and DR2 do a MF flow label classification, policing and scheduling of traffic according to the levels negotiated in the SLS/TCSs, SLA/TCAs.

The following sequence illustrates the process by which an application obtains Intserv end-to-end QoS when RSVP is used by the hosts.

4.1.1. Intserv Control Plane Processing

4.1.1.1. Intserv Control Plane Processing in Domain A

1. The QoS process on the sending host Tx generates an RSVP PATH message that describes the traffic offered by the sending application. The sender template contains a flow label based filter spec that identifies the traffic flow.

2. The PATH message is carried toward the receiving host, Rx. In the network domain A, to which the sender is attached, standard RSVP/Intserv processing is applied at capable network elements, including IR1, IRx, IRy, IR2, which install Intserv/RSVP state in the routers.

3. At the edge router IR2, the PATH message is sent onward to the Diffserv network region.

4.1.1.2. Diffserv Control Plane Processing in Domain B

4. The PATH message is ignored by routers in the Diffserv network domain B.

5. The Diffserv QoS information (SLSs/TCSs) has been configured in edge and core routers manually, or via automated scripts, or dynamically using COPS provisioning.

4.1.1.3. Intserv Control Plane Processing in Host Rx

6. When the PATH message reaches the receiving host Rx, the RSVP module, and the networking stack in the operating system of the receiving host RX generates an RSVP RESV message, indicating interest in offered traffic of a certain Intserv service type.

7. The RESV message is carried back by the Diffserv network domain B towards network domain A, and the sending host Tx. Consistent with standard RSVP/Intserv processing, the RESV message may be rejected at any RSVP-capable node in the path if resources are deemed insufficient to carry the traffic requested.

4.1.1.4. Diffserv Control Plane Processing in Domain B

8. The Diffserv network domain B ignores the RESV message.

4.1.1.5. Intserv Control Plane Processing in Domain A

9. In IR2, the RESV message triggers admission control processing, if it was configured so. IR2 compares the resources requested in the RSVP/Intserv request to the resources available in the Diffserv network domain at the corresponding Diffserv service level. The corresponding service level is determined by the Intserv to Diffserv mapping discussed previously. The availability of resources is determined by the capacity provisioned in the SLS/TCS, SLA/TCA. IR2 may also apply a policy decision such that the resource request may be rejected based on the customer's specific policy criteria, even though the aggregate resources are determined to be available per the SLS/TCS, SLA/TCA.

10. If IR2 approves the request, the RESV message is admitted and is allowed to continue upstream towards the sender. If IR2 rejects the request, the RESV is not forwarded and the appropriate RSVP error messages are sent. If the request is approved, IR2 updates its internal tables to indicate the reduced capacity available at the admitted service level on its transmit interface.

11. The RESV message proceeds through the network domain A, through routers IRy, IRx, and IR1, undergoing RSVP processing, towards the sender Tx. Any RSVP node in this domain may reject the reservation request due to inadequate resources or policy. If the request is not rejected, the RESV message will arrive at the sending host, Tx.

11. At Tx, the QoS process receives the RESV message. It interprets receipt of the message as indication that the specified traffic flow has been admitted for the specified Intserv service type (in the Intserv-capable nodes). Tx can now send traffic to Rx.

4.1.2. Intserv-Diffserv Data Plane Processing

1. Tx sends packets to Rx. The traffic from Tx through IR1, IRx, IRy, to IR2, in domain A, is processed for QoS purposes according to Intserv - it is flow label classified, policed, and shaped/scheduled.

2. The DR1 applies MF flow label classification, policing, marking, shaping/scheduling to the TX->RX traffic according to the SLS/TCS, SLA/TCA.

3. Routers DRx, DRy in the core of the Diffserv domain (region) B, and DR2, apply Diffserv QoS to the TX-> traffic. This includes BA classification, policing, shaping/scheduling.

4. The traffic arrives from DR2 to Rx.

In this manner, we obtain end-to-end QoS through a combination of networks that support flow label RSVP/Intserv and networks that support flow label Diffserv.

5. Security Considerations

This document introduces no new security concerns. The security concerns are essentially identical to those concerning the diffserv field (traffic class) itself, as outlined in [DSCP-Def], {Diffserv], and [Differv-Tun].

When IPv6 packets are encrypted using ESP Transport or Tunnel Mode [IPSec-ESP], the port and protocol numbers are hidden, but the flow label is not. Thus MF classification remains possible even for encrypted traffic.

6. IANA Considerations

No IANA considerations need be made.

7. Acknowledgments

The discussion on the topic in the IPv6 WG mailing list has been instrumental for the definition of this specification. The authors want to thank Steve Blake, Jim Bound, Francis Dupont, Robert Elz, Tony Hain, Christian Huitema, Frank Kastenholz, Hesham Soliman, Michael Thomas, Jun-ichiro itojun Hagino, and others that unintentionally perhaps were omitted, for their tireless contributions on the list.

Very special thanks to Brian Carpenter for his patient participation, guidance, sharing of ideas, and highly principled actions.

8. References

[IPv6] S. Deering, R. Hinden, "Internet Protocol Version 6 Specification", RFC 2460, December 1998.

[Diffserv] M. Carlson, W. Weiss, S. Blake, Z. Wang, D. Black, and E. Davies, "An Architecture for Differentiated Services", RFC 2475, December 1998

[DSCP-Def] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

[PHB-ID] D. Black, S. Brim, B. Carpenter, F. Le Faucheur, "Per Hop Behavior Identification Codes", RFC 3140, June 2001.

[Diffserv-Tun] D. Black, "Differentiated Services and Tunnels", RFC 2983, October 2000.

[Diffserv-PIB] M. Fine, K. McCloghrie, J. Seligson, K. Chan, S. Hahn, A. Smith, "Differentiated Services Policy Information Base", Work in Progress.

[DiffServ-MIB] F. Baker, K. Chan, A. Smith "Management Information Base for the Differentiated Services Architecture", Work in Progress.

[Diffserv-model] Y. Bernet, S. Blake, D. Grossman, A. Smith "An Informal Management Model for Diffserv Routers", Work in Progress.

[IPv6-flow-label]J. Rahajalme, A. Conta, " An IPv6 Flow Label Specification Proposal", Work in Progress.

[KEYWORDS] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[CONS] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, October 1998.

[Basic Socket] R. Gilligan, S. Thomson, J. Bound, W. Stevens. "Basic Socket Interface Extensions for IPv6," RFC 2533 March 1999.

[Intserv-Diffserv] Y. Bernet, P. Ford, R. Yavatkar, F. Baker, L. Zhang, M. Speer, R. Braden, B. Davie, J. Wroclawski, E. Felstaine. "A Framework for Integrated Services Operation over Diffserv Networks", RFC 2998, November 2000

9. Authors' Addresses

Alex Conta
Transwitch Corporation
3 Enterprise Drive
Shelton, CT 06484
USA
Email: aconta@txc.com

Jarno Rajahalme
Nokia Research Center
P.O. Box 407
FIN-00045 NOKIA GROUP
Finland
E-mail: jarno.rajahalme@nokia.com

1239