CS625 : Advanced Computer Networks
Instructor: Prof Bhaskaran Raman braman@
Lecture 14 , Fri, 29 Aug 2003
Scribe : K.Anantha Kiran (Y3111019) ananth@

Differentiated Services (diff serv)

Initially internet delivered one type of service , best-effort , to all traffic. There are two methods of adding a differentiated level of service to IP , each designated by One-bit in the Packet header.

Premium service:

In this a portion of the Network bandwidth is allotted to the requester for a particular flow. This service will be priced as it is creating virtual-leased lines.

It is interesting to note here that whenever allocated capacity is not being used it is avilable to the best-effort traffic. Incontrast to normal best effort traffic which is bursty and requires queue management to deal fairly with congestion control , this Premium traffic by design creates very regular traffic pattrens .

Premium service levels are specified as a desired peak bit-rate for a specific flow (or aggregation of flows) .

The user contract with the n/w is not to exceed the peak rate . The n/w duty is that the contracted bandwidth will be avialable when traffic is sent .

First-hop routers filter the packets entering the network , set the premium bit of those that match a Premium service specification, and perform traffic shaping on the flow that smooths all traffic bursts before they enter the network. This approach requires no changes in hosts. A compliant router along the path needs two levels of priority queueing, sending all packets with the Premium bit set first. Best-effort traffic is unmarked and queued and sent at the lower priority. This results in two "virtual networks": one with buffers designed to absorb traffic bursts; and one where traffic is limited and shaped to a contracted peak-rate, but packets move through a network of queues where they experience almost no queueing delay.

This works within trust boundary. Because intermediate routers assume that all the policing needed was done by First hop router perfectly.

This scheme also requires complicated admission and setup procedures.

Assured Service:

The Assured service implements a service that has the same delay characteristics as (undropped) best effort packets and the firmness of its guarantee depends on how well individual links are provisioned for bursts of Assured packets. On the other hand, it permits traffic flows to use any additional available capacity without penalty.

This follows "expected capacity" usage profiles that are statistically provided. The assurance that the user of such a service receives is that such traffic is unlikely to be dropped as long as it stays within the expected capacity profile. An Assured service traffic flow may exceed its Profile, but the excess traffic is not given the same assurance level.

RED queue management is used for this type of traffic. But Assured traffic has lower packet dropping probability than the Best-effort traffic. Here also admission control involves lots of mathematics making it complicated.

Advantages of having two levels of services:

Premium service is a conservative allocation of resources, unused bandwidth that had been allocated to Premium might provide some "headroom" for under allocated or burst periods of Assured traffic or for best effort. This service would seem to have a strong appeal for commercial applications, video broadcasts, voice-over-IP, and VPNs. On the other hand, this service may prove both too restrictive (in its hard limits) and overdesigned (no overallocation) for some applications.

Assured service allows bursts to happen in their natural fashion . As there appear to be a number of advantages to an architecture that permits these two types of service , and they can be made to share many of the same mechanisms using two bit patterns from the IP header precedence field. Those two bits are Premium or P-bit , Assured or A-bit.

Two bit differentiated services architecture:

The forwarding path mechanisms can be broken down into those that happen at the input interface, before packet forwarding, and those that happen at the output interface, after packet forwarding. Intermediate routers only need to implement the post packet forwarding functions, while leaf and border routers must perform functions on arriving packets before forwarding. We describe the mechanisms this way for illustration; other ways of composing their functions are possible.

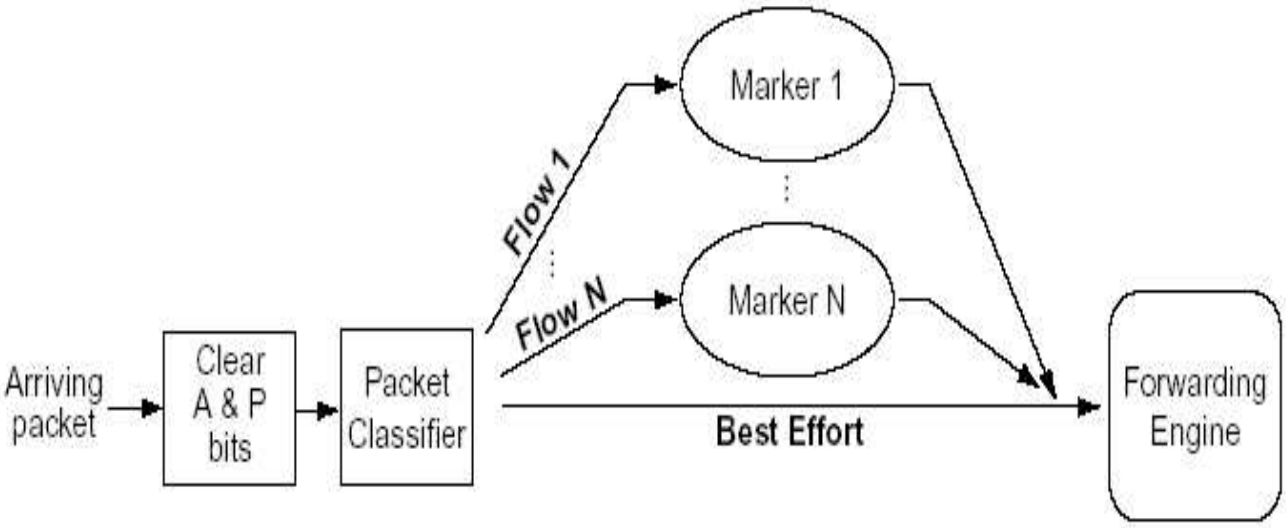

Leaf router input functionality

All arriving packets must have both the Abit and the P-bit cleared after which packets are classified on their header. If the header does not match any configured values, it is immediately forwarded. Matched flows pass through individual Markers that have been configured from the usage profile for that flow: service class (Premium or Assured), rate (peak for Premium, "expected" for Assured), and permissible burst size (may be optional for Premium). Assured flow packets emerge from the Marker with their A-bits set when the flow is in conformance to its Profile, but the flow is otherwise unchanged. For a Premium flow, the Marker will hold packets when necessary to enforce their configured rate. Thus Premium flow packets emerge from the Marker in a shaped flow with their

P-bits set. (It is possible for Premium flow packets to be dropped inside of the Marker when burst is enough to overflow the holding queue). Packets are passed to the forwarding engine when they emerge from Markers. Packets that have either their P or A bits set we will refer to as Marked packets.

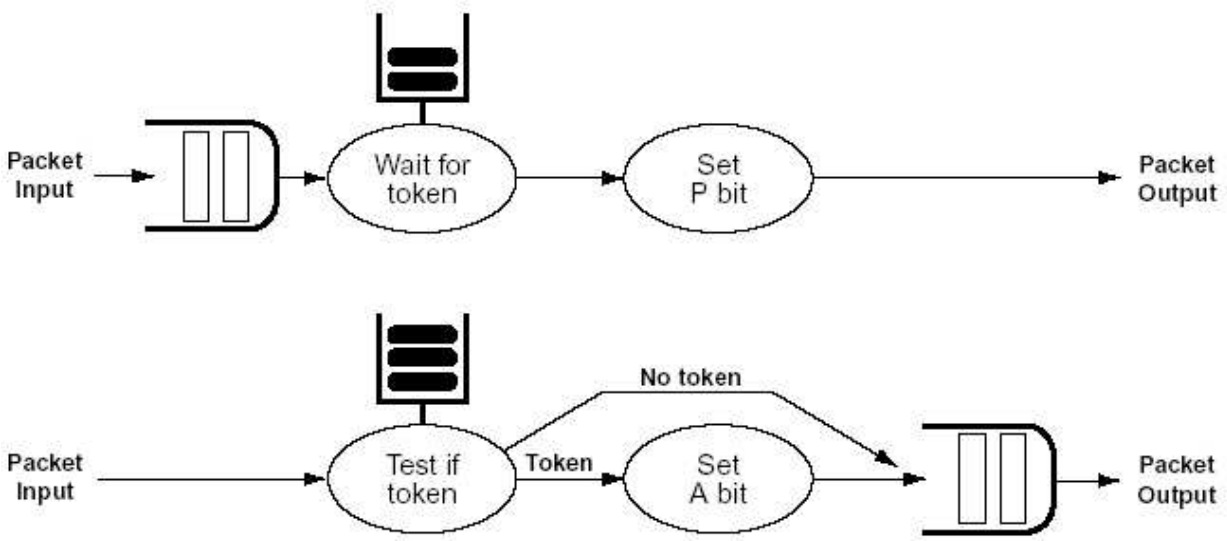

Markers Block Diagram

Here is the working of Marker. For Premium, the bucket can only be allowed to fill to the maximum packet size; while Assured may fill to the configured burst parameter. Premium traffic is held until a sufficient byte credit has accumulated and this holding buffer provides the only real queue the flow sees in the network. For Assured, traffic, we just test if the bytes in the bucket are sufficient for the packet size and set A if so. If not, the only difference is that A is not set. Assured traffic goes into a queue following this step and potentially sees a queue at every hop along its path.

SERVICE	Token PRESENT	Token not PRESENT
Premium service	P bit is set	Wait until required no. of tokens are accumilated
Assured service	A bit is set	Forward traffic as if best-effort

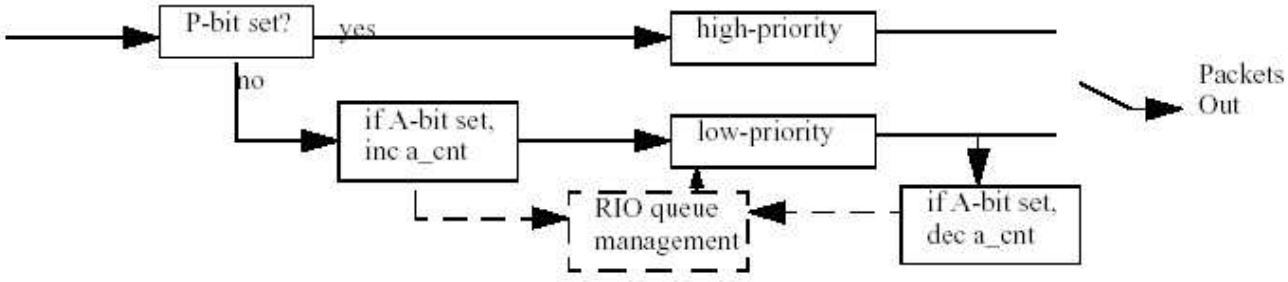

Router output interface

Each output interface of a router must have two queues and must implement a test on the P-bit to select a packet's output queue. The two queues must be serviced by simple priority, Premium packets first. Each output interface must implement the RED-based RIO mechanism described in on the lower priority queue.RIO uses two thresholds for when to begin dropping packets, a lower one based on total queue occupancy for ordinary best effort traffic and one based on the number of packets enqueued that have their A-bit set. This means that any action preferential to Assured service traffic will only be taken when the queue's capacity exceeds the threshold value for ordinary best effort service. In this case, only unmarked packets will be dropped (using the RED algorithm) unless the threshold value for Assured service is also reached. Keeping an accurate count of the number of A-bit packets currently in a queue requires either testing the A-bit at both entry and exit of the queue or some additional state in the router.

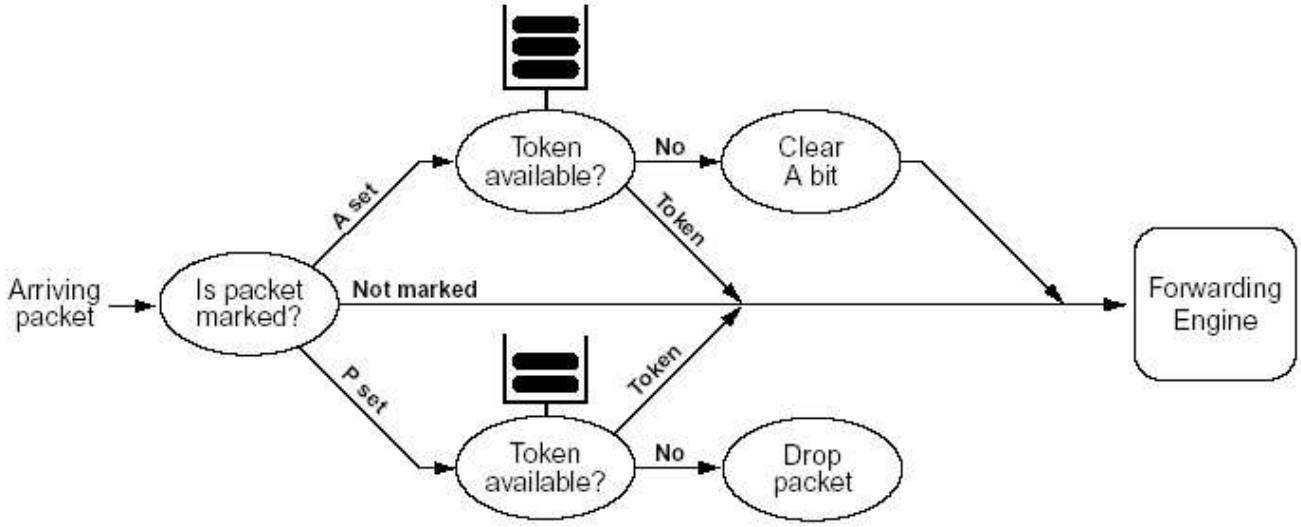

Border routers input interface Profile Meters

Border routers input interface will require Profile Meters . The bilateral agreement between Adjacent administrative domains must specify a peak rate on all P traffic and a rate and burst for A traffic (and possibly a start time and duration). A Profile Meter is required at the ingress of a trust region to ensure that differentiated service packet flows are in compliance with their agreed-upon rates. Non-compliant packets of Premium flows are discarded while non-compliant packets of Assured flows have their A-bits reset.

Framework for Marked Traffic allocation:

The goal of differentiated services is controlled sharing of bandwidth. It can be done by individuals(users) who will set bits for their packets to represent services , or by Agents also called Bandwidth Brokers (BBs) who know about organizations policies and prorities. A BB is associated to a paricular region or Organizations . It establishes limited trust with peer BB.

BB has two responsibilities:

1. Allocating different profiles to Marked traffic allocations and setting Leaf routers appropriately

with that Profiles.

2. It has to manage the messages that are sent across the boundaries to adjacent BBs.

Only a BB can configure a Leaf router (which is necessary for a secured system)

Traffic allocation procedure:

When a service (allocation) is required for a paricular flow , user send this request to BB. Request should include service type , a target rate , a maximum burst , and time period when service is required.

1. BB first authenticates the message to know whether it is a valid request from valid user.

2. Allocates the required bandwidth if it is , by reducing avialable bandwidth by the amount allocated .

3. If requesters destination is out this trust region , the request must fall within the class allocation through the next hop trust region estabilshed by bilateral agreement between two regions.
a) The requester BB informs the adjecent BB that it needs some bandwidth from its allocation
b) Other BB authenticates the request and sets the Border router with appropriate profile to allow that particular traffic.

4. Now BB configures the appropriate leaf router with the information about the packet flow to be given a service at the time the service is to commence.

5. This configuration is called **Soft state** that BB will periodically refresh.

The RSVP signalling procedure can be used in conjunction with BBs.