



IP Premium Agenda

- Services specification and implementation discussion
- Qos Parameters



Géant QoS Services Specifications

Mauro Campanella
Tiziana Ferrari

Mauro.Campanella@mi.infn.it
Tiziana.Ferrari@cnaa.infn.it



Requested Services

From the Géant proposal and TEN-155 operation, the following services are recommended:

- 1 Managed Bandwidth Service (MBS), a replacement of what is implemented today through the use of dedicated ATM CBR permanent virtual circuits;
- 2 Guaranteed Bandwidth Service, to provide minimum guaranteed bandwidth between two network nodes,
- 3 IP Premium service “ to provide priority over the basic IP service in order to meet defined Quality of Service parameters”.
- 4 Last, but not least, Best Effort !



Proposed framework

- Leverage overprovisioning where possible
- Use the Diffserv model
- Collapse MBS and IP Premium and map it to the EF PHB
- Use the AF PHB for the GCS service



IP Premium/MBS

MBS provides:

- 1 isolation between flows at the **link layer** at two levels:
 - packet level;
 - protocol level: native transport of e.g. MPLS, IPv6 without any need for tunnelling.
- 2 bandwidth guarantees;
- 3 bounded one-way delay and one-way delay-variation;
- 4 null packet loss in case of well-shaped input traffic complying with the peak rate associated to the pvc.

IP Premium/MBS (cont.)



Marking: EF PHB DSCP to build a low loss, low latency, low jitter, assured bandwidth, end-to-end service through DS domains".

Policing: strict policing at the diffserv domain ingress, so that packets exceeding the contract are dropped.

Strict policing ensures that the average amount of IP premium traffic injected in a priority queue is limited

Queuing: priority queuing is the most suitable queuing mechanism for minimizing packet loss, one-way delay and delay variation. WFQ and DRR can be used alternatively, provided that the service rate of the premium queue is appropriately over-estimated. The premium queue is exclusively used by IP premium packets.

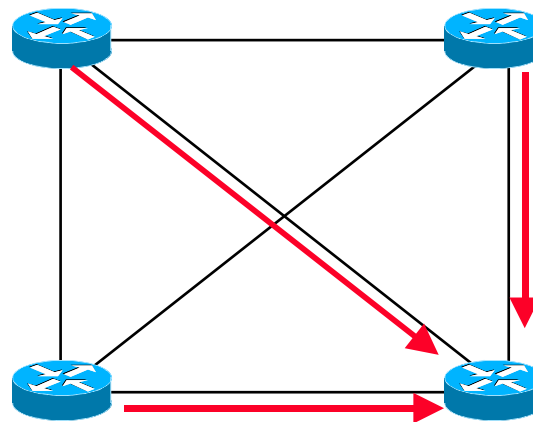


IP Premium/MBS (cont.)

Open issues

- shaping. May be needed at the ingress and in the core to avoid packet loss at aggregation points, even if BW is available;
- service can be
 - 1- point to point (allows a precise load estimate)
 - 2- independent of the destination
- need to police at the egress in case 2 ?
- asymmetric SLA
- need for different ACLs in and out in various places based on IP static address only
- transport of protocols other than IPv4. Can be done using MPLS or again at the link layer (SDH,) ?

If the SLA is for a “generic” IP Premium service and, for example at each node the SLA (towards the core) allows 10Mb



each node can receive up to 30 Mb of IP Premium traffic

GCS



GCS should provide:

- transport of at least a nominal BW
- loose or none guarantees on delay and ipdv
- allow packet loss in case of congestion of GCS traffic or burstiness
- allow transport of traffic exceeding nominal BW, up to a maximum BW, if capacity is available

GCS (cont.)



Marking: AF PHB DSCP using a separate (hw) queue from BE.

Policing: traffic is metered according to a token bucket algorithm, whose parameters are defined by the per domain behaviour. No dropping of excess traffic is used, it is rather down marked as a different color (or BE) DSCP.

Queuing: WFQ where weight is directly dependent on the minimum bandwidth provided to the queue.



IP Premium/MBS (cont.)

Open issues

- shaping is not probably needed
- number of colors in the AF PHB
- possible packet reordering when SLA is exceeded (if excess packets are transferred to the BE queue)
- asymmetric SLA, egress policing
- TCP-friendly vs TCP-unfriendly traffic (fairness to microflows)
- instantaneous loss rate can be high, but not on average
- RIO or WRED for congestion avoidance

- performance if on the same queue of BE



Working together

- “static” provisioning (with BGP policy propagation)
- mapping to real hardware features, i.e. choose the features common to most platforms
- need for at least 3 hardware queues (IP Premium, GCS and Best Effort)
- compute maximum available IP Premium and GCS bandwidth, given the core link capacities (no cheating allowed)

- Measure the baseline of the values on an empty but working network
- Extension to multiple domains
- Monitoring



Dimensioning

If the SLS allows, for example for each node in a mesh of 25, towards the core 10Mb/s of “generic” IP Premium service and 10 Mb/s of GCS, the total reserved bandwidth on a link can be up to $25 \times (10 + 10) = 500 \text{ Mb}$.

Proposal: overprovision IP Premium at one third of the sum of SLS BW, monitor links for performance

Ask for destination of traffic



QoS parameters

Agreed set:

- Delay, ipdv (moving target as definition, but agreed as concept)
- packet loss, bandwidth
- *provisioning, BER, availability*
- *MTU size (fragmentation)*
- *duplicate and out-of-order packets*

QoS value ranges



	Single value (SV)	Short range (class 0)	Medium (class 1)	Wide range (class 2)
One-way Delay	Measured value at empty network (baseline)	less than SV + 50 ms (150 ms)	less than SV + 250 ms (400 ms)	less than SV + 10 s (1 s)
ipdv	Between 0 and the time needed to transmit one full MTU at line speed	25 ms (50 ms)	50 ms (50 ms)	none (1 s)
Packet loss (Probability)	null	$< 10^{-4}$ (10^{-3})	$< 10^{-3}$ (10^{-3})	< 0.1 (10^{-3})
Bandwidth (speed 64Kb/s)	Fixed value Allow for a burst of one MTU packet	n.a.	n.a.	a minimum of one full MTU size packet per second

Between parenthesis are ITU-T Y.1541 draft values



Next

Map proposed model to actual hardware implementation

Start measuring values on an unloaded/overprovisioned network

Map QoS on WAN to QoS in LAN (802.1p)