Premium IP on GÉANT

Early Experiences with an Inter-Domain Diffserv Deployment

Simon Leinen

November 17, 2002

Context: Research Networking in Europe

Hierarchy of research networks

- Trans-European Backbone
  - current generation: *GÉANT* (before: Europanet, TEN-34/155)
  - one PoP per country (exception: NORDUnet access sharing)
- National Research & Education Networks
  - wide variety of sizes, backbone speeds, technology...
- (In some places: regional research networks)
- Campus networks

GÉANT http://www.geant.net/geant/

- Current European research backbone
- Link speeds: 9.6 (9), 2.4 Gb/s (17), 622 (6) and 34-155 Mb/s (10)
- Routers: mostly Juniper M-160
- Designed and built by DANTE
- NOC operated by contractor (commercial)

SEQUIN Project Overview

Background

Replacement for TEN-155's (GÉANT's predecessor) Managed Bandwidth Service. MBS was based on ATM PVPCs/PVCCs that were extended through NRENs and campus networks.

Advantages:

- Operational experience with a QoS offering spanning domains.
- MBS usage hints at the potential for new service.

Disadvantages:

- Biases design towards connection-oriented services

Requirements Analysis

Interviews with 10 user groups from GÉANT community

Extensive questionnaire about target applications and coverage,
qualitative and quantitative QoS expectations, current connectivity.

QoS-Enhanced Services Considered

“Premium IP” service: low-jitter and strictly rate-limited.
“IP+” (assured rate) left for further study.

Implementation Concepts

Within a domain: classical EF-based implementation.

Specify different policing granularities at core/edge domains.

- core (GÉANT) should police on source/destination AS combination, but this cannot be easily implemented right now.

For easier implementation, be lenient with respect to bursts.

Early Tests and Measurements

Laboratory testing at high data rates

Tested policing and scheduling (WRR/MDRR, WRED) behavior of core router platforms (Juniper M160 and Cisco 12400 with “Engine 3” line cards) at STM-16 (2.4 Gb/s) and STM-64 (9.6 Gb/s) line rates.

- Commercial packet generators and loggers (Smartbits)
- QoS features found to work as advertized and without noticeable impact on performance

Early Tests and Measurements

International tests with H.323 and synthetic traffic (1)

Given that videoconferencing, and H.323 in particular:

- was frequently mentioned as an application in the interviews,
- is actively being used in the research community,
- has traffic patterns that are fairly well understood,
- has been noted as problematic over wide-area networks,
- and should make for a nice multi-party test case,

We decided to use this as a basis for more extensive wide-area tests.

Early Tests and Measurements

International tests with H.323 and synthetic traffic (2)

1. Perform actual H.323 videoconferences with and without Premium IP, and have experts assess perceived quality.
2. Controlled measurements with traffic generation/capturing software, using traffic patterns similar to those seen in H.323 usage.

The goal was to learn more about:

- Performance of Premium IP under realistic conditions.
- Adding Premium IP capability to various types of networks.
- Provisioning Premium IP service instances.

Early Tests and Measurements

International tests with H.323 and synthetic traffic (3)

Other results:

- Packet reordering (even where IPDV was very low).
- It's useful to have a modified traceroute[a] that reports DSCP changes.

[a] ftp://ftp.login.com/pub/software/traceroute/beta/

Initial Deployment

Lessons from beta testing with a few research groups

- Implementing Premium IP is still hard for NREN/campus networks.
- If application is an overlay network (as in AQUILA and MOICANE), configuration and troubleshooting/monitoring become hairier.
- User satisfaction and the performance of the underlying mechanisms are not that directly related.
- The “provisioning” process needs to be streamlined
  $\Rightarrow$ work on SLA/SLSs; decoupling from EF implementation

Next Steps

1. Turn Premium IP into a “production service” (DANTE)
2. Study other diffserv applications for GÉANT (IP^+, LBE, ABE?)

Outlook

Focus shifting from QoS mechanisms and SLAs to a more holistic view of the “end-to-end performance problem”.

- Performance monitoring and problem diagnosis
  - Continuous monitoring (QoS “beacon” matrix)
  - End-to-end QoS trace similar to traceroute?
- PERT (“Performance Emergency and Response Team”)—build expertise across several areas that are traditionally separate, but whose interaction determines end-to-end performance: network performance proper (core/campus), computer/OS architecture, transport and higher-layer protocols, distributed algorithm design…

Somewhat orthogonal to whether networks provide differentiated QoS.