Project Number: IST-1999-20841
Project Title: SEQUIN

Deliverable D4.2

QoS Implementation Plan

Deliverable Type:	PU-Public
Contractual Date:	31 May 2002
Actual Date:	31 May 2002
Work Package:	4
Nature of Deliverable:	RE - Report

Authors: Mauro Campanella, INFN-GARR,
Roberto Sabatino, DANTE

Abstract:

The suggested solution for QoS is Premium IP. This deliverable addresses the technological, operational and management aspects of implementing Premium IP on a wide scale on operational networks such as GEANT and NRENs.

Keywords: QoS, GÉANT, Premium IP

1 Executive Summary

An overview of the essential aspects of implementing QoS in a multidomain and heterogeneous environment are presented. These aspects are technological, operational and management aspects. The solution for QoS as recommended by SEQUIN is a Premium IP service which is based on Differentiated Services and can interoperate with ATM and over provisioned networks. From a technological point of view there are a number of different tasks that need to be performed by the network nodes. These tasks are policing, admission control, classification, marking and scheduling. Each network node in each network domain has a specific task to undertake. Border routers perform policing and admission control, core routers perform scheduling whilst marking is done as close as possible to the source.

The operational aspects for the implementation of Premium IP are considered. Premium IP packets which fail admission control are to be treated as normal Best Effort packets whilst Premium IP packets that exceed the contracted bandwidth must be discarded. A provisioning process based on identifying individuals with clear responsibilities in each participating domain, combined with the use of discussion lists proved successful in ensuring that Premium IP configuration requests can be dealt with efficiently. Finally, from an operational perspective, it is fundamental to define an interdomain monitoring infrastructure for Premium. SEQUIN recommends a monitoring infrastructure based on a mixture of passive and active measurement systems. The passive systems monitor bandwidth and packet loss whilst the active measurement systems monitor delay and ipdv. The active measurements are taken on a per domain and between domain basis, and where possible end to end as well.

Finally management and policy aspects are crucial for the wider implementation of QoS. The results of the work of SEQUIN must be widely disseminated, and the expertise widely distributed. Of crucial importance is the strategy adopted by networks (NRENs and GEANT) for the support of QoS: first of all it is important that suitable router hardware is in place and secondly it is important that networks commit resources (both human and infrastructure for monitoring) for the support of QoS.

2 Introduction

The work carried out in WP2 resulted in the definition of a Premium IP service which can operate on top of the IP protocol, on an end to end basis across multiple management domains. Premium IP is the recommended solution for providing a high level of QoS in a multi-domain and heterogeneous environment and is based on the Differentiated Services Expedited Forwarding Per Hop Behaviour (EF-PHB). It can be implemented by combining different data link technologies (ATM, POS) and techniques (network overprovisioning, prioritised queuing) in each domain or even on each network node. A domain can be a connected subset of the whole network. One of the most important aspects of Premium IP is the fact that each IP packet that belongs to a Premium IP flow must be tagged in its header with a defined value called DSCP (Differentiated Service Code Point). The value 46 (bit pattern 101110) is chosen for Premium IP. This is to ensure the correct identification of a Premium IP packet so that appropriate packet treatment techniques can be applied on domain and network node in order to ensure guarantees in terms of the four network parameters identified in D2.1: capacity, packet loss, one way delay and ipdv (IP packet delay variation)

This deliverable will provide a technological overview of how the required functions can be performed on each network node, giving concrete configuration examples. The deliverable addresses the operational requirements in terms of provisioning process, monitoring and recommended hardware. Finally, management aspects such as availability of expertise, investments in hardware already committed to by network domains and levels of support offered by different network domains are discussed.

3 Technological overview

A description of functions to be performed on the various network nodes is available in deliverables D2.1 [D2.1] and D3.1 [D3.1]. Results of testing these functionalities on two high speed router platforms (Juniper and Cisco 12K) are available in D5.1 [D5.1]. Addendum 1 [A1] to D2.1 gives also a detailed description of which specific functions should be performed by which specific network nodes. A brief summary of Addendum 1 to D2.1 is provided in the following sub-sections and is complemented by configuration examples.

Figure 1 outlines a user in network *L1* which sends Premium IP traffic in network *L2*, which are interconnected via intermediate networks *N1* and *N2*, and a *core* network.

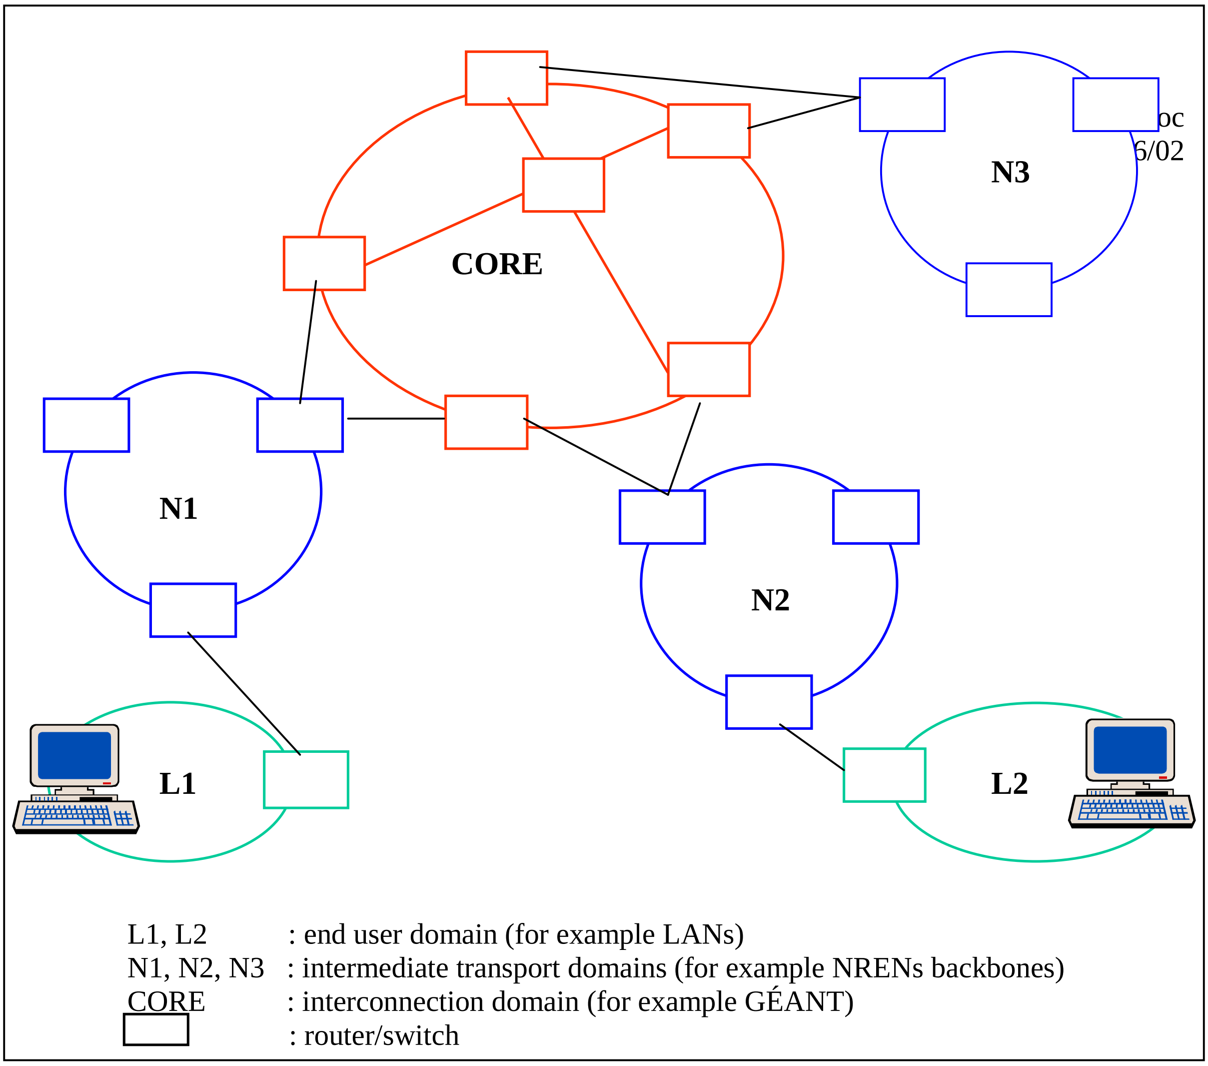

Figure 1: simplified multidomain network

3.1 Policing

Each microflow must be policed as close as possible to the source of the traffic flow. The parameters for the policing function should be at least three: source and destination IP addresses and contracted sending rate. In figure 1, assuming that the source is the workstation in L1, then the network node in L1 must implement the first policing function. Policing is to be performed also at the boundaries between different administrative domains, but only on ingress traffic. Therefore at the network boundary between L1 and N1 only the router in N1 needs to perform the policing action. The same applies between the core and N1, whereby only the core will perform policing on ingress from N1. Once the Premium IP traffic reaches the core, it is suggested to simplify the task of the networks by not performing any other policing functions when forwarding traffic to the final destination.

The Premium IP model is based on DiffServ whose basic concept is to provide guarantees on aggregate traffic rather than single flows, to ensure scalability and simplicity. It is therefore recommended that at the core (and correspondingly at the intermediate levels) network level the policing parameters should be based on the aggregate contracted Premium IP traffic between the peer network domains. Taking the figure as an example, the core will police all traffic received from N1 destined to N2 with the aggregate Premium IP traffic contracted to be sent between these two networks.

Detailed configuration examples are available at the SEQUIN technical web site
http://www.dante.net/sequin/technical

3.2 Admission control and classification

Admission to the Premium IP service is done as close as possible to the source host and is based on the IP source and destination addresses or prefixes.

Between domains packet will be classified according to the QoS tag (DSCP or an appropriate IP Precedence value), "trusting" the ingress domain.

The admission control can be based also on other parameters, as defined case by case. A particular case is that the source is capable of tagging the packets and admission is then granted when the tag is present, which is discouraged on a LAN due to security concerns.

Admission control and classification must be enabled on all border routers of the DiffServ domain in the form of a general deny unless explicitly allowed by a rule. The general “deny” rule must be active before the service is started.

Detailed configuration examples are available at the SEQUIN technical web site
http://www.dante.net/sequin/technical

3.3 Marking

Although a single DSCP value for all domains is not mandatory according to the Diffserv specification, an identical value of 46 (bit pattern 101110) on all domains is strongly suggested to avoid additional complexity.

In the LAN environment 802.1p may be used instead of Diffserv to provide QoS assurances to selected flows. In this case there is the need that LAN edge routers translate between DSCP and 802.1p markings. A VLAN, coupled to 802.1p, in which only Premium IP allowed nodes are members, can provide an effective access control to the service itself.

Detailed configuration examples are available at the SEQUIN technical web site
http://www.dante.net/sequin/technical

3.4 Scheduling

For the Premium IP service the scheduling must use the highest priority queuing algorithm available, for example Priority Queuing or Weighted Round Robin with the maximum weight on Premium IP queue.
High priority queuing must be enabled at all nodes of involved Diffserv domains, or at least along all the relevant paths.

Detailed configuration examples are available at the SEQUIN technical web site
http://www.dante.net/sequin/technical

3.5 Interaction with ATM based networks

ATM based network can offer virtual circuits with different types of quality of service. The Constant Bit Rate Permanent Virtual Circuit (CBR PVC) in particular offers the equivalent of a leased line of constant capacity at the data link layer. For the Premium IP service an ATM CBR

PVC can be considered equivalent to a single hop exhibiting the EF behaviour. Only EF traffic must flow through the ATM CBR PVC and the CBR PVC's configured capacity determines its rate.

The ATM CBR PVC can also be a concatenation through ATM switches of many PVCs of appropriate, matching capacity. Two ATM CBR PVCs connected through a router DO NOT constitute a single EF hop.

For the sake of avoiding packet losses, unless the CBR PVC is used also as a policerto drop excess EF aggregate traffic, it is suggested to configure the CBR PVC capacity to a value 15% larger than the computed maximum aggregate EF traffic. The increase in capacity will allow absorbing burstiness and it is preferred to configuring a large the ATM Maximum Burst Size on the PVC.

As the PVC is dedicated to Premium IP traffic, there is no need of a priority queuing inside the PVC itself.

In case of multiple PVC on the same router interface, the Premium IP traffic flows must have precedence in respect to other flows on different circuits. Care must be exercised also when switching packets from an ATM PVC to a different data link interface, to ensure appropriate servicing priorities. The techniques are a function of hardware type.

3.6 Interaction with over provisioned networks

A link is called "overprovisioned" here when it possesses enough capacity so that, at all instants, there is never congestion nor packet drop. In addition to overprovisioning the link itself, also the hardware at its ends should be able to cope with the link speed and perform the requested functions on each packet at wire speed. The mentioned characteristics should be valid for a substantial part of time, compatible with SLA/SLS agreements.

It can be shown that overprovisioning, as outlined above, does provide reasonable assurances to QoS parameters, although there is no commitment that overprovisioning will hold true over a long period of time.

The Premium IP service can thus benefit of overprovisioned network even if the service is not implemented technically as described in the model and the domain is not compliant.In particular in the LAN environment and the core backbone, when overprovisioned, can provide paths to packets that avoid packet losses and minimise delay variation.

4 Operational aspects

The previous section has summarised which are the functions to be performed where they should be performed and has given concrete configuration examples. There remain a number of important operational aspects that need to be addressed. These include the standardisation across network domains of packet marking, the treatment of packets which fail admission control, provisioning processes, monitoring and troubleshooting.

4.1 Packet marking

The Diffserv specification does not require that Premium IP packets are tagged with the same DSCP in each network domain, but in principle every network domain may choose its own DSCP for Premium IP. Nevertheless, to simply the multi-domain implementation of Premium IP it is strongly recommended to apply the same DSCP value for Premium in each network domain. The recommended DSCP value is 46 (binary 101110). This value is the value recommended by the IETF in the specification of the EF-PHB.

4.2 Treatment of packets that fail rate limitation

When a packet exceeds the contractual Premium IP capacity it must be dropped. The policing action should be performed as close as possible to the sending source, preferably on a single flow or on an aggregate of few flows.

4.3 Treatment of packets that fail admission control

The issue of how to treat packets that fail admission control for Premium IP has been topic for extensive discussions within the SEQUIN project and with the wider community represented by the GEANT [GEA] Access Port managers (APMs). A study has been conducted on GEANT to assess the amount of IP traffic with DSCP values other than IP Best Effort, whose DSCP value is 0. This study showed that 94.6% of packets transiting GEANT had a DSCP value of 0, whilst 0.05% of packets were already tagged with DSCP equal to 46. The remaining packets were tagged with different DSCP values.

SEQUIN recommends that for the wider implementation of Premium IP on networks such as GEANT, Premium IP packets which fail admission control are re-tagged with DSCP value equal to 0 (IP Best Effort) and forwarded accordingly. This is considered as the least disruptive or intrusive measure taken to enable a smooth introduction of Premium IP. The possibility of discarding the packets is considered too intrusive and could cause significant disruption to end-users.

4.4 Interaction with other Services

The Premium IP service has no interaction or collision with other services, like multicast. Capacity contention is solved by definition through the explicit agreements on its deployment.
Premium IP can as well be deployed in parallel with other QoS services, like Less than Best Effort (LBE), provided that the hardware supports both at link speed. Premium IP uses only one of the 64 possible DSCP values, the others are left untouched by the service.

4.5 Router Hardware

As stated previously, D3.1 and D5.1 discuss the functions required by routers in order to implement Premium IP. Especially when dealing with high capacity networks (Gbps), it is important that these functions are supported at line rate. The work carried out in WP5 has outlined that when operating at Gbps capacities, Juniper routers support all the functions needed at line rate. Cisco routers are also capable of handling the functions required at line rate, but it is important to install the correct version of line cards. More specifically, the engine-3 line cards (which operate at 2.5Gbps) have this ability, whilst at the time of carrying out WP5 the engine-4 line cards (which operate at 10Gbps) were able to perform only a subset of the functions required.
Cisco engine-2 line cards, which operate up to 2.5Gbps, also support only a subset of the functions required.
When operating at lower capacities, for example up to 155Mbps, the Cisco 7500 or 7200 series routers are able to support all the functions required.

4.6 End to end provisioning process

Provisioning Premium IP in a multi-domain environment is a complex task which requires extremely careful co-ordination and operational discipline to ensure that the correct configuration is put in place in each domain in a timely fashion.
For each group of users one technical person assumes the role of Single Point of Contact (SPC) who co-ordinates the configuration process in the various domains involved. The SPC will avail of the assistance of a Technical Contact (TC) from other network domains involved and will

ensure that all available documentation on configuration examples are forwarded to the TCs. The SPC is also responsible for ensuring that the SLA/SLS documents as described in Addendum 2 to D2.1 are filled with the requested information.

D6.1 [D6.1] provides a detailed description of the recommended processes that would apply to the research environment composed by GEANT and the NRENs. However, the work carried out in WP6 (testing in a user environment) has shown that a different approach to inter-domain provisioning proved more effective than the one set out in D6.1. This different approach is based on detailed collaboration and exchange of information between all parties, at times in a somewhat un-coordinated fashion.
What is set out in D6.1 remains the recommended way of working in a multi-domain environment and has been specified based on the experience gained with the TEN-155 Managed Bandwidth Service (MBS) which was based on ATM. The TEN-155 MBS operational structure [MBS] proved successful mainly because the technology being used (ATM) was already being used for several years and its parameters and configuration were well understood by all parties involved. This is not yet the case for Premium IP, which is why initially a more collaborative approach is needed with the aim to finally be able to adopt the process set out in D6.1.

As a practical comment, to provide the user with the requested QoS at the application level, an additional debugging step is needed, which links network QoS to the end to end application performance. This quite complex task has to be performed in strong collaboration by the users and a person or a team with knowledge of applications and networking.
Although it is outside the scope of network management, it is a key element for the success of the service. It is suggested that NRENs evaluate the need of such a professional figure.

4.7 Monitoring

As with any type of service, monitoring is a fundamental complementary component for the operation of a Premium IP service. We distinguish between metrics to be measured, interdomain measurement strategy and measurement technology used. A detailed description of these issues is available in Addendum 3 [A3] to D2.1, whilst a summary is provided here.

The measurement metrics correspond to the four parameters associated with network QoS: bandwidth, packet loss, one way delay and ipdv. Packet loss can be measured on the aggregate traffic, whilst delay and delay variation can be measured using real traffic flows end to end or injecting a small capacity flow in the aggregate inside a domain or between different domains.
The current state of the art of monitoring performance parameters relies on passive and active measurement techniques. Passive measurement techniques are based on software modules that query the network devices for information regarding performance parameters. Routers for example contain a Management Information Base (MIB) from which it is possible to retrieve information using SNMP. This technique is useful for collecting information on traffic volumes (capacity) and packet loss. QoS specific MIBs have been standardised by IETF, but are not yet implemented.
Active measurement techniques rely on external devices, or probes, which inject test traffic into the network. This test traffic normally contains an absolute or relative time stamp and it is received by another external measurement device which can calculate performance parameters based on the time of arrival and the time stamp set in the IP packet. This technique is useful for measuring parameters such as one-way-delay and ipdv.
The recommended monitoring infrastructure uses both passive and active measurement techniques. The passive measurements are used for collecting information in respect of aggregate capacity and packet loss associated with Premium IP ideally on each network node involved in

transporting Premium IP traffic. Polling agents (workstations with appropriate SW) located at network nodes collect the information and make it available for presentation via standard web pages using modules such as MRTG or Cricket. A tool which performs these tasks has been developed and deployed within the GEANT network (taksometro)

Active measurements are proposed to collect information in respect of one way delay and ipdv. It is proposed that these measurements are taken on a per domain and between domain basis and where possible on an end-to-end basis if the appropriate equipment and software is available at the end user sites. It must be noted that end users are encouraged to run the appropriate SW to enable them to verify the SLA/SLS.
Figure 2 outlines the proposed interdomain active measurement infrastructure set out in Addendum3 to D2.1

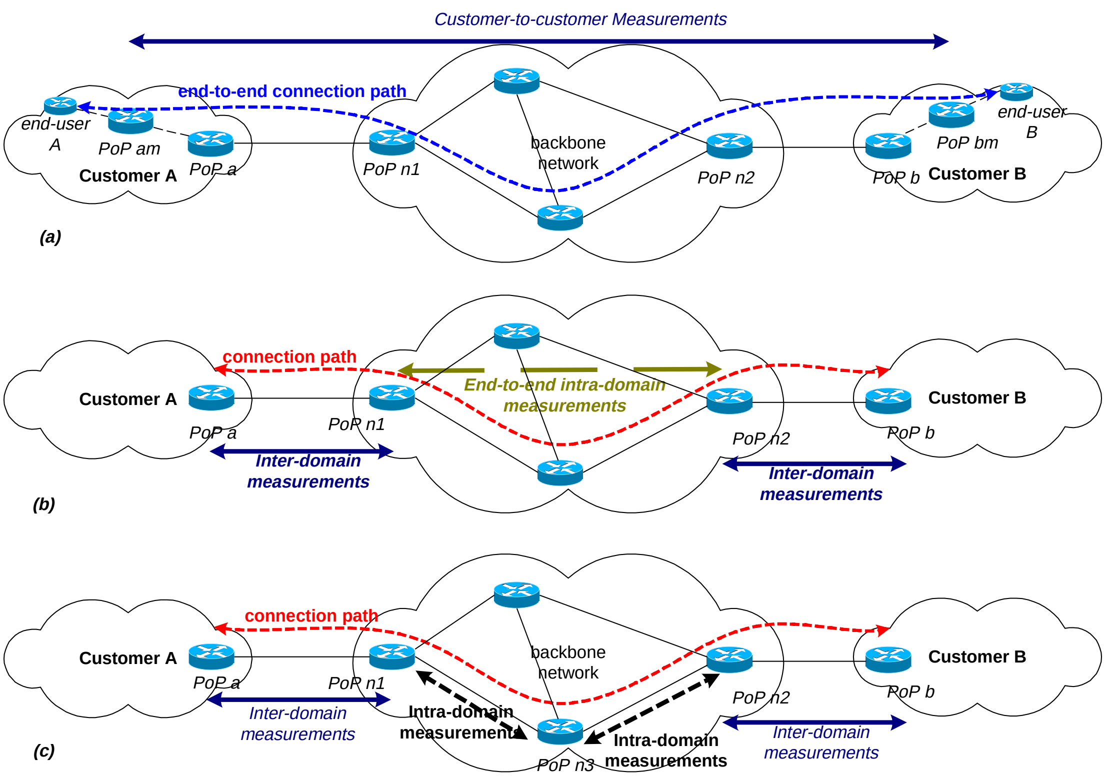

Figure 2: Interdomain active measurements infrastructures

Addendum 3 to D2.1 contains a comprehensive analysis of two toolsets to measure one-way delay and ipdv (rude/crude SW vs. RIPE TTM systems) and two ways of providing accurate timing/synchronisation information (NTP vs. GPS). It is obvious that from a technical perspective RIPE TTM or dedicated measurement boxes like Smartbits from Spirent combined with GPS is the most accurate way to provide the information. However there are possible high costs associated , logistical difficulties in ensuring that GPS receivers can be installed in network PoPs. It also appears that the accuracy obtainable with NTP is perhaps sufficient for the pan-European performance requirements. For these reasons, it is proposed to implement a monitoring infrastructure based on rude/crude software initially, with the view to migrate to a more costly and accurate setup once Premium IP is adopted on a truly wide scale.

5 Management aspects of implementing Premium IP

It is the view of the SEQUIN partners that the largest obstacles in the way of the wider implementation of Premium IP are the management aspects, which include:

- Dissemination of knowledge
- Expertise available in different administrative domains
- Investments in hardware already committed to or required
- Level of commitment and support to Premium IP by different administrative domains
- combined application-network debugging

Despite these obstacles, the SEQUIN partners firmly believe that a Premium IP pilot phase with a selected number of users in a limited number of countries is possible. This pilot phase, which should be carried out in the context of GEANT, is key to providing the necessary critical mass that is needed to facilitate the wider adoption of Premium IP.

5.1 Dissemination

SEQUIN has made considerable efforts to disseminate the results and findings of its work. There are regular meetings with the joint DANTE/TERENA task force TF-NGN which is attended by approximately 40 experts from across Europe and there have been two discussion sessions with the GEANT APMs. This enabled detailed discussion with the wider community, which ultimately lead to the detailed definition of Premium IP. A workshop with the user community was held and was attended by a number of international user-groups, all of which have expressed enthusiasm with the work carried out by SEQUIN. Four of the user groups which attended the workshop have committed to take part in Premium IP trials with GEANT and the NRENs. The results of SEQUIN have also been presented at a variety of national and international conferences.
From this point of view, SEQUIN has done all what was in its possibilities and scope to address the wider community.
Nevertheless this still represents a minority of the wider community and generating greater awareness is a task for all interested NRENs to reach their own end-users and international user groups.

5.2 Expertise

At least 5 administrative domains are involved when setting up Premium IP between two users in a pan-European environment. These are the two local area user networks, the two national research networks and the pan-European network GEANT. Often, at a national level, there are even more network domains involved, especially when the national infrastructure is based for example on regional networks, MANS and so on. It is reasonable to expect up to eight network domains to connect two users in two different countries. GEANT currently connects 27 NRENs, and has interconnections with other regional research networks such as ABILENE, CANARIE, SINET. Considering that Premium IP can be requested by any user in any country gives an idea of how many network entities and network administrators are potentially involved and need the right expertise.
It is to be expected that not all of these individuals have the required expertise, given that each of them deals with different realities on a daily basis. Many network domains are understaffed, and the available staff is already under pressure to perform even the simplest networking tasks.
SEQUIN has made available via the web technical information and configuration examples. This certainly helps in conveying the essential technical information, but it must be taken into account that especially at the end user sites a wide variety of networking equipment is used. For this wide

selection of equipment, SEQUIN has obviously not had resources, nor was it within its scope, to perform performance analysis and provide configuration examples.

Although it is outside the scope of network management a person or a team with knowledge of both applications and networking.it is a key element for the success of the service. It is suggested that NRENs evaluate the need of such a professional figure.

5.3 Investment in hardware

The section on operational aspects has outlined that, especially for networks with Gbps capacities, the correct hardware be in place to support Premium IP. Some NRENs have recently or are in the process of upgrading their networks to Gbps capacities and in this case they have the opportunity to procure the recommended hardware. It becomes a national, conscious decision at this point whether the NREN decides to procure the recommended hardware.

Other NRENs completed networks upgrades before the results of SEQUIN were clear and disseminated or even before the recommended router hardware was actually available on the market (this is particularly true for the Cisco 12K engine-3 line cards). In this case, NRENs may have chosen HW which was state of the art at the time but not fully supportive of Premium IP at Gbps capacities (such as the Cisco engine-2 line cards). These are the problematic cases in that NRENs have recently invested considerable resources in equipment that is not fully supportive of Premium IP.

5.4 Levels of Support offered by NRENs

For an NREN we can distinguish 4 positions with regard to Premium IP:

- *incompatible*: if a NREN resets or discards packets with Premium IP DSCP values, it will be hard if not impossible to connect users via such an NREN;
- *indifferent*: NREN preserves Premium IP DSCP value but applies BE treatment;
- *supportive*: NREN preserves Premium IP DSCP value and offers environment where in general Premium IP characteristics hold, (e.g. by overprovisioning);
- *compliant*: NREN offers a realisation of Premium IP that is compliant with the specification (the implementation techniques may vary (ATM, PQ, WRR, POS).

These positions can be dictated by a number of parameters, which include the level of expertise available, human resources available to support Premium IP, strategic choice within the network domain to support Premium or only Best Effort, availability of suitable HW.

If an NREN (or any other network domain) is incompatible or indifferent to Premium IP then it is not feasible to enable Premium IP for users whose connectivity crosses those network domains, as no guarantees can be offered and furthermore no monitoring or support for troubleshooting can be offered. The optimal situation is whereby a network domain is fully compliant. Network domains which are not fully compliant but which are at least supportive can be used to offer Premium IP, **without explicit end to end guarantees,** to end users. This greatly simplifies the initial end to end deployment of Premium IP as a significant number of NRENs have implemented high capacity networks in which the general Premium IP characteristics hold. In these cases though it is recommended that the NREN or network domain allocate resources to assist in troubleshooting and making available the recommended monitoring infrastructure.

6 Conclusions - Introducing Premium IP

The difficulties outlined in the previous sections for the wider deployment of Premium IP can be minimised, and practical actions can be taken be each network domain to assist in its implementation. As many network domains are now highly provisioned, these network domains can with relatively little effort become *supportive* of Premium IP. The main efforts required in this case are the human resources to assist in troubleshooting and setting up the monitoring infrastructure. A considerable amount of expertise from the SEQUIN partners is available to assist in this process and is available within the context of GEANT beyond the project lifetime of SEQUIN.

In terms of router configuration, Premium IP does not need to be configured on each router within a network domain but only on those along the path towards an end user. Given that the number of end users requesting Premium IP is expected to be limited for each domain, the effort required to do the recommended configurations is viewed to be manageable.

Where investments in out of date hardware have been made (such as Cisco engine-2 line cards instead of engine-3) it is still possible to perform some of the main functions such as scheduling. Other functions such as policing (for which the engine-2 line card has performance limitations) can be performed by different domains and complemented by network monitoring. It is feasible that a network domain may trust a downstream network domain to always perform policing closer to the users therefore all Premium IP packets received can be considered legitimate and within the contracted bandwidth limits.

Addendums 1,2 and 3 to D2.1 provide descriptions of the functions to be performed, the SLA/SLS and the recommended monitoring infrastructure to be used for the implementation of Premium IP.

This deliverable, drawing from the experience of WP6 (reported in D6.1) has also addressed the operational and management aspects for its wider deployment and implementation, and despite there being several obstacles it is viewed that Premium IP can be supported by a large number of NRENs with a manageable effort.

Dissemination of skills and knowledge to the wider network administrators community remains an issue which can be solved only by dedication to the cause by the experts available, with the support of the NRENs, and the extended use of technical information being maintained on web sites.

7 Acronyms

APM	Access Port Managers
ATM	Asynchronous Transfer Mode
CBR	Constant Bit Rate
DSCP	Differentiated Services Code Point
EF-PHB	Expedited Forwarding Per Hop Behaviour
GPS	Global Positioning System
IETF	Internet Engineering Task Force
IP	Internet Protocol
LAN	Local Area Network
LBE	Less than Best Efforts
MAN	Metropolitan Area Network
MBS	Managed Bandwidth Service
MIB	Management Information Base
MRTG	Multi Router Traffic Grapher
NREN	National Research and Education Network
NTP	Network Time Protocol
POP	Point of Presence
POS	Packet over Sonet
PVC	Permanent Virtual Connection
QoS	Quality of Service
SLA	Service Level Agreement
SLS	Service Level Specification
SNMP	Simple Network Management Protocol
SPC	Single Point of Contact
TC	Technical Contact
VLAN	Virtual Local Area Network
WRR	Weighted Round Robin

8 References

[A1] Addendum 1 to SEQUIN D2.1,
Implementation architecture specification for the Premium IP service
M. Campanella, October 2001

[A2] Addendum 2 to SEQUIN D2.1,
SLA spec for IP Premium Service
A. Sevasti, M. Campanella, October 2001

[A3] Addendum 3 to SEQUIN D2.1
Monitoring and Verifying Premium IP SLAs
Athanassios Liakopoulos, April 2002

[D2.1] Definition of Quality of Service
SEQUIN deliverable D2.1
M. Campanella, P.Chivalier, A. Sevasti, N. Simar
May 2001

[D3.1] Definition of Testbed
SEQUIN deliverable D3.1
M.Campanella, M. Carboni, P.Chivalier, S. Leinen, J. Rauschenbach,
R. Sabatino, N.Simar – April 2002

[D5.1] Proof of Concept testbed
SEQUIN deliverable D5.1
M. Przyblyzki, R. Sabatino, N.Simar, J. Tiemann, S. Trocha – December 2001

[D6.1] Testing in a User Environment
SEQUIN deliverable D6.1
R. Roth – May 2001

[EF-PHB] RFC-2598
An Expedited Forwarding Per Hop Behaviour
V. Jacobson, K. Nichols, K. Poduri – June 1999

[GEA] GEANT, The Pan-European Network for Research
http://www.dante.net/geant

[MBS] The TEN-155 Managed Bandwidth Service procedures
http://www.dante.net/ten-155/brochure_MBS.html