White Paper:

Quality of Service Assurance over IP Networks

Editor :	Vincent Alexandre (BRI)
Contributors:	R.McMahon (BRI)

Release:	1.0
Date :	8/3/02
Distribution :	Public

ABSTRACT

The Internet is, since the invention of the World Wide Web in 1994, the preferred method for information exchange and human communications. It provides a wonderful encyclopaedia with efficient table of contents, the search engines, and an easy mechanism for human communications, the handy and free e-mail service. To reduce costs and delays, enterprises will probably carry out the next revolution of the Internet. Businesses will adopt the Internet to trade with their customers, providers and peers, and replace actual phone, fax or e-mail-based negotiation, invoice and payment mechanisms by automatic, Internet-based ones. Quality and security assurance will be two major issues.

This document describes the requirements for quality assurance over the Internet, in a business-to-business framework, and in relationship with the FORM project managed and co-funded by the European Commission. It proposes to use the Internet2 bandwidth brokerage concept as a protocol for service activation, and to introduce another protocol for quality of service negotiation. It provides also some implementation issues.

KEYWORDS

Table of Contents

1 Introduction

The success of the Internet is surely mainly due to a well-engineered protocol. The Internet Protocol (IP) [1] has been designed in such a way that the stream of information is split into a multitude of packets, each packet being carried independently of the other ones. The main advantage is that a physical line is never reserved for a communication; the IP protocol is very scalable and perfectly suitable for burst traffic, such as web data.

Now that it has proved to be a strong and reliable protocol, the telecommunication community is widely applying it to other fields. The Universal Mobile Telecommunication System (UMTS), the next generation for mobile telephony from the Third Generation Partnership Project (3GPP) [2], is based on an “all IP” core network, where the actual circuit-switched architecture, which reserves a physical line during the whole communication process, is replaced by an IP packet-switched model. And even the fixed phone world could be based in the future on the Internet protocol, using models such as voice over IP.

2 Quality of Service Assurance Standards

Internet traffic is today based on a best effort scheme: when congesting (i.e. approximately from 10 a.m. to 6 p.m. each working day of the week), the network discards packets that overload its charge, in order to come back to a manageable situation. Discarding packets could appear catastrophic: how coherent can a data stream be if some of its bytes are deleted? Fortunately some upper protocols (e.g. TCP) automatically deals with congestion problems by requesting again a packet if it is discarded: as a result the data stream is not corrupted, but simply delayed.

Over the current Internet all traffic endures the same quality, whoever they belong to. Quality of Service (QoS) [3] solves the issue differentiating traffic and prioritising it, and then discarding first, during congestion, packets belonging to a traffic that does not require high quality. In order to achieve this, several models have been defined and tested in a more or less large scale. The models we will focus on are those of the Internet Engineering Task Force (IETF).

2.1 QoS models

A first example of QoS models from the IETF is the Multi-Protocol Label Switching (MPLS) architecture. It proposes to insert in each packet a label which makes an IP routed network behave such as a connection-oriented one. In MPLS, a router determines the path to reach a destination using a label as an index in a table that provides the next router identifier, and a new label.

The IETF has also designed the Integrated Service (IntServ) model, intended to create on an IP packet-switched network the closest model to circuit emulation. During a signalling phase, the ReSerVation Protocol (RSVP) checks the routers onto the route to the destination and their available resources. If all the routers can assure the requested QoS, then their resources are effectively reserved and will remain reserved during the whole communication. Bandwidth is not over-sold, service is guaranteed: it is, in effect, a virtual circuit.

The QoS model that is the most promising is the IETF Differentiated Service (DiffServ) [4] model, because it fully respects the Internet philosophy in that it does not reserve any resources but tries instead to maximise their usage. DiffServ makes use of the IP Type of Service (ToS) field, which codes the behaviour of the network elements, per-hop, onto six-bits. Two classes of service have been defined, corresponding to well-defined ToS values:

- the Assured Forwarding (AF) class proposes to allocate a specific amount of network resources for each level of service (AF1→4), prioritising the different levels and discarding or downgrading traffic that exceed the agreements when there is congestion
- alternatively the Expected Forwarding (EF) class creates a virtual leased line, and only services whose QoS can really be fulfilled are accepted.

When entering a domain, the IP packets are identified and classified by the edge routers into the class of service they should belong to, according to the agreements that have been negotiated. This classification can be done according the date and time of the day, the source and destination IP addresses, or other IP fields. Edge routers mark the ToS field of each packet, which allows core routers to apply the good behaviour to the flow. We speak about aggregated flows, since, for a single ToS value, the same behaviour will be applied whomever the traffic packets belong to.

These models have been existing already for a while, and it is possible to find components on the shelves that implement them, e.g. routers that speak a QoS language or another one. As a matter of fact it is today possible to assure QoS onto a same domain and so to give a better quality to some users. But it is impossible to assure QoS through different domains; a protocol, agreed and spoken by all the network operators, is urgently needed.

2.2 Internet2 QBone initiative

The QoS architecture adopted and proposed in this paper is based on the IETF DiffServ model, but uses also results from the Internet2 (I2) QBone initiative [5]. This initiative is part of the Internet2 consortium, led by 180 North America universities to "develop and deploy advanced network applications and technologies". The QBone initiative aims to deploy on a real network, this of the I2 universities, new IP QoS technologies. It is based on the Bandwidth Broker (BBr) concept, agent responsible for:

- receiving Resource Allocation Requests (RAR) from end-systems, and answering them
- managing the resources of the domain it is responsible for
- configuring edge devices
- communicating with peer BBrs to ensure QoS can be fulfilled across domains.

The I2 QBone BBr architecture has defined the Simple Inter Bandwidth Broker Signalling (SIBBS) protocol, a protocol over TCP used to receive and answer RARs. Because it is likely to be extended to the international community once it would have proven its strengths on the American education network, the model we present here uses and completes the I2 QBone protocol. This will be presented in part 5.1.

3 Business to Business

It is not unreasonable to think, considering the striking and promising development of the Internet, that more and more telecommunication fields will use the IP protocol to carry their data. And Business-to-Business (B2B) [6] communications, because they are forecasted to be incredibly important in the next decade, are an interesting field. B2B is defined as “dealing with the electronic commerce conducted between businesses”. Considering two companies trading the one with the other, communications are performed: one or the two of them provides the other one with products or services and the characteristics of the trade have to be negotiated, controlled, and the service has finally to be paid:

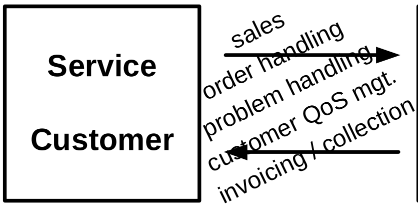

Figure 3-1: Classical B2B model.

In order to organise the services to be carried out between the two businesses, Business Processes (BP) have been identified by the Telecommunication Management Forum (TM Forum) [7]:

- service definition, sales, and order handling, part of the Fulfilment BP
- invoicing and collections, part of the Billing BP
- problem handling and customer QoS management, part of the Assurance BP.

Some activities resulting from these BPs can be outsourced, and services resulting from this outsourcing will have to be managed. This is what the FORM European project proposes to deal with.

4 FORM

The FORM project is part of the research and development program of the European Commission: the Fifth Framework’s Information Society Technology (IST) program (project IST-1999-10357). It aims to “Engineer a Co-operative Inter-Enterprise Management Framework Supporting Dynamic Federated Organisations Management”. This framework introduces the concept of an Inter-Enterprise Service Provider (IESP), responsible for the management of services whose characteristics are represented and negotiated upon a Service Level Agreement (SLA):

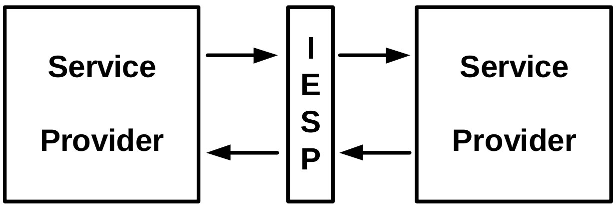

Figure 4-1: FORM B2B model.

The consortium has set up three working groups, one for each of the BPs designed by the TM Forum, and proposes to identify their interactions and the data flows that are transmitted between them. For example, if the quality that should be assured for a service is not fulfilled then financial compensations can be expected. We will describe now the QoS assurance BP.

5 FORM QoS Model

5.1 Negotiation vs. activation

The I2 QBone BBr architecture supposes as one of its preliminary requirements that Service Level Specifications (SLS) have been agreed between domains, during a SLA negotiation step; we integrate SLS to the set of (parameter, value) pairs which compose the SLA.

SLA negotiation is analogous with a flight ticket reservation. In the flight model, parameters (e.g. dates, price, class) are negotiated between a service provider (a travel agency) and a customer (a traveller). In the FORM model, they are negotiated between the IESP and a service customer, as well as between the IESP and different Internet Service Providers (ISP), and between the different ISPs themselves. The part of the SLAs defining QoS parameters are represented in FORM through an XML [8] based protocol allowing an inter-domain multi-service negotiation, and the XML schema used for the request is also used for the answer. Let's take as an example an educational service provided by a university on a first domain, to a student on another domain, each Friday from 8 a.m. to 10 a.m., and composed of two services: a videoconference and an internal newsgroup:

```
<?xml version="1.0"?>
<bbs:bandwidthBrokeredSLA […] >
   <version>20010101120000</version>
   <sLAId/>
   <sLAName>Educational System</sLAName>
   <state>REQUESTED</state>
   <rARs>
      <rAR>
         <rARId/>
         <rARName>Video-conference</rARName>
         <sourceIPAddress>10.10.20.2</sourceIPAddress>
         <sourcePort>10000</sourcePort>
         <destinationIPAddress>10.10.20.3</destinationIPAddress>
         <destinationPort>8080</destinationPort>
         <ingressRouterIPAddress>10.10.20.1</ingressRouterIPAddress>
         <activationMode>AUTO</activationMode>
         <activationDelay>0.0</activationDelay>
         <startTime>969541200000</startTime>
         <stopTime>969548400000</stopTime>
         <dSCP>1</dSCP>
         <bandwidth>128000</bandwidth>
         <delay>11</delay>
         <jitter>10</jitter>
         <lossProbability>13</lossProbability>
         <state>REQUESTED</state>
      </rAR>
      <rAR>
         <rARId/>
         <rARName>Internal newsgroup</rARName>
[…]
   </rARs>
</bbs:bandwidthBrokeredSLA>
```

Figure 5-1: Bandwidth brokered (BBred) SLA example.

From this compound service definition, SLAs will be defined and set between the IESP and all the ISP domains to be crossed. Resources will so be early planned, the educational service's traffic flows aggregated with others flows and resource usage efficiently optimised.

Nevertheless on some Friday the service will not be consumed (e.g. the student is in holiday). The FORM QoS model proposes so to confirm a resource reservation through an activation request, in the same way as a flight reservation is confirmed during the check□in operation; this double request mechanism ensuring resources, critical onto the Internet, will not be wasted. For this activation request, the I2 QBone BBr protocol is used. When the service application is run, it sends automatically an activation request as a RAR message, over the TCP protocol as defined by the I2 QBone initiative, whose characteristics comply with the SLA that has been previously negotiated. The BBr first checks if a SLA has previously been negotiated or not for the time the application tries to activate the service. If it has not been, then a dynamic SLA negotiation will be performed, ensuring that even at the last minute available resources can be still used, and so sold.

But the BBred SLA XML schema defines also an "activation mode" parameter, set to "auto" on Figure 3. In the "auto" mode, the BBr will automatically perform the activation: the customer has so not to handle the two requests, but whatever it uses or not the service it will be fully charged. In a "non-auto" mode, a value for the "activation delay" parameter has to be provided. If the activation request does not occur during this delay, then the network resources can be reused for another service, and the customer will be partially refunded. In the "QoS-flight" reservation comparison, this "activation delay" can correspond to the end of the check-in time, and an airline company is potentially allowed to resell a place if the check-in is not performed in time.

5.2 Tariff mechanism

The comparison we did between QoS reservation and flight reservation processes can be extended to charging issues. For instance, since the network operators are then able to efficiently aggregate traffic flows, and potentially update or complete their network, it is expected that the sooner a customer will request a service, the cheaper the service will be:

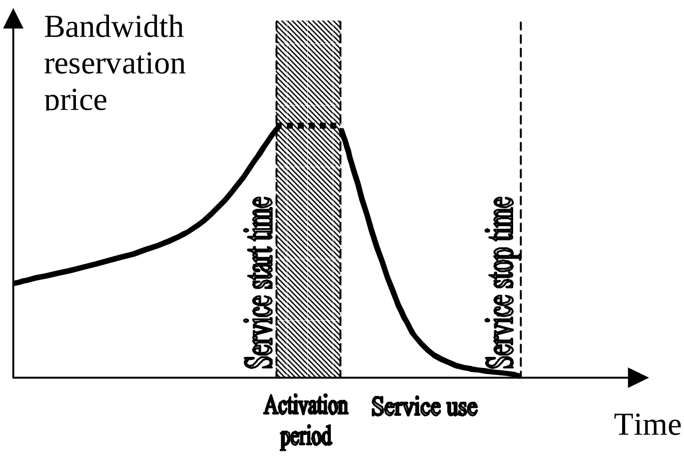

Figure 5-2: Tariff scheme example.

In the same way, resource reservations are frozen during the activation period. And during the service usage period, the price decreases in order to encourage potential late customers to buy the service. The policy for the provider is then to fully use its network capacity whatever, almost, the price is, since resources will be lost otherwise.

6 FORM Network Provisioning

Let's go back to the QoS Assurance model in FORM. There is a need to divide it into two levels:

- an Information Service Assurance (ISA) level, responsible for providing end-to-end QoS at the application level
- a Network Provisioning Assurance (NPA) level, concerned with network connection and QoS provision.

The QoS parameters (i.e. bandwidth, delay, jitter and loss probability) at the NPA level are estimated from the quality that has been negotiated for an application at the ISA level, and resource reservation reports are regularly, or on demand, transmitted from the lower level to the upper one. We will present now and in details the architecture of the NPA model.

6.1 Components architecture

The Java 2 [9] Enterprise Edition (J2EE) [10] architecture, from Sun Microsystems, has been chosen to implement the FORM QoS assurance engine. Enterprise JavaBeans (EJB) are used to represent the BBr, Service Negotiation Requests (SNR) and Service Activation Requests (SAR), and Network Elements (NE):

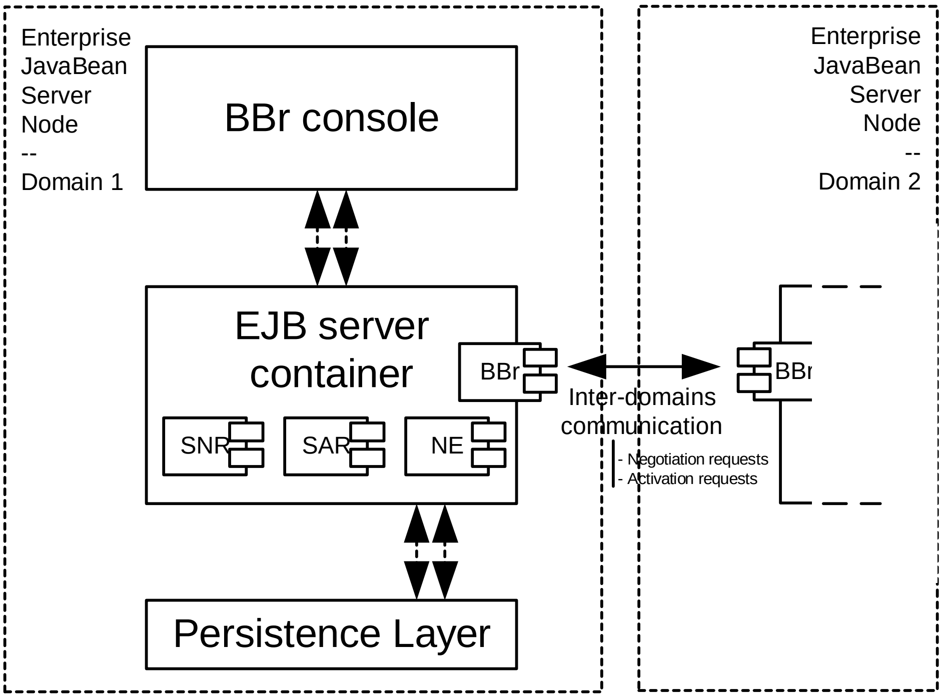

Figure 6-1: Components architecture for a bandwidth brokerage in FORM.

The use of SIBBS for the activation mechanism has determined the protocol to use for this communication. SIBBS is an over-TCP protocol that defines, bit after bit, information that is negotiated/activated between a QoS provider (BBr) and a QoS customer (end-system, or another BBr).

The negotiation protocol is performed through an EJB invocation call, where the SLA is provided as an XML document. The QoS customer accesses a Java Naming Directory Interface (JNDI) service to find a reference of the BBr responsible for its domain. It then requests the service specifying the SLA as a parameter. And the answer of the BBr is finally returned by the EJB as a document fulfilling the same XML-schema as the request: the BBred SLA XML-schema described on Figure 3.

The BBr is a Container Managed Persistency (CMP) entity bean. This model of components indeed fully fulfils the requirements of a BBr:

- need for access and storage of intra- and inter- domains topology information
- need for serving client requests.

One instance of a BBr needs to be deployed on each of the domains; and so an EJB server/container (s/c) is needed. The JBoss and the Orion EJB s/c have been assessed: both are free of charge for a developer use, and the first one is even available as open source. Although it was expected that the two EJB s/c would behave in the same way since the EJB architecture is a set of specifications that should normally be implemented similarly, some problems of compatibility appeared, fortunately solved today.

In fact, the object model design proposes a three-layer architecture to fulfil the specificity of this double request mechanism. The first layer offers an object model representation for the messages of the SIBBS protocol: it is mainly responsible for formatting and parsing the TCP stream bits to/from the architecture components. The second layer proposes a general Application Programming Interface (API) for a bandwidth brokerage, Internet2 QBone compliant: this layer offers for instance interfaces to represent network links, and it will be the responsibility of a network administrator to implement these interfaces to map this general API with his network. And the third layer finally proposes an implementation of this general layer, based on the FORM specifications and original ideas, and called B3 (for Broadcom BBr):

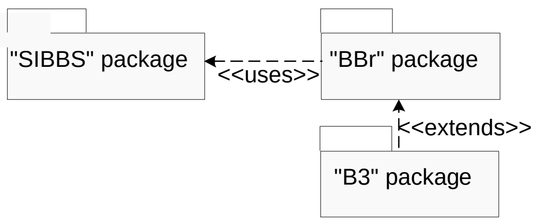

Figure 6-2: Package dependencies.

6.2 Admission control

The BBr, designed onto the second layer of this architecture, is responsible for the admission control; in other words it has to receive RARs, compute and finally send an answer. Performing this can be very complex and introduce statistics and even artificial intelligence to use as efficiently as possible all the network resources. It is also a big challenge to implement the whole DiffServ architecture, and in particular to assure bandwidth with the AF class since traffic can be discarded.

The BBr we implemented answers to some of these issues. For each network resource (e.g. a link), a special amount of resources (e.g. bandwidth, buffer size) is allocated for the different DiffServ classes of service. The BBr responsible for a domain manages all the resource reservations that have been allocated, and it is so able to accept or reject all new RARs, or to suggest an alternative according the scheme:

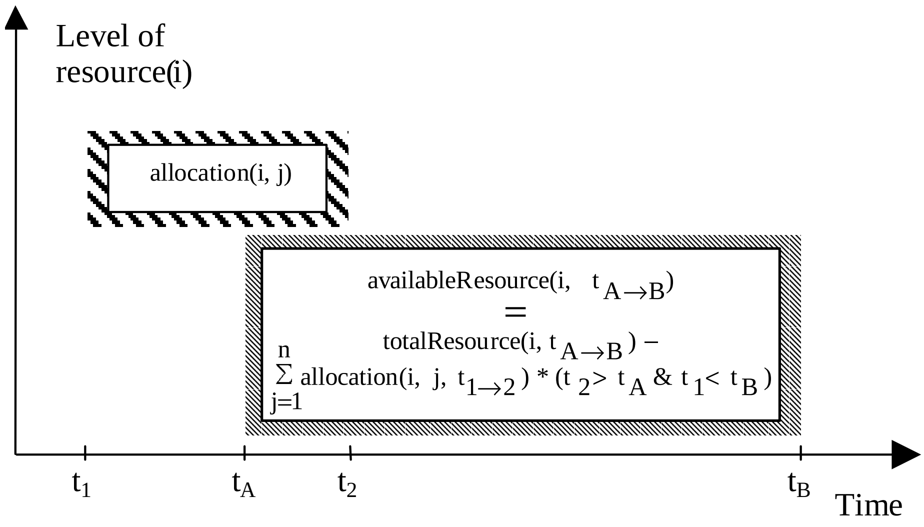

Figure 6-3: Admission control policy.

For each of the network resources, the level of available resource is computed from the total amount of resource allocated to the class of service, minus the sum of the amount of resources that have already been allocated to services. And in order to ease the logic, when considering the period [tA, tB] and an allocation from t1 to t2 with t2>tA and t1>tB, (i.e. the two intervals "have" an intersection), the period the resource is allocated on is "extended" to the period the available resources are computed.

The admission control ends, in our prototype, with the enforcement of the QoS into the network elements. Free BSD Linux boxes implementing the Class Based Queuing (CBQ) QoS algorithm compose the two domains of our test bed. Onto the reception of an activation request, if it is accepted, the BBr connects to the Free BSD ingress router and telnets a CBQ configuration file, providing network resources to be allocated and traffic characteristics information.

7 Conclusion

This paper presented how Quality of Service can be assured over the Internet protocol, in an inter-domains architecture, and how it has been implemented in the FORM European project. It presented briefly the Internet protocol and the main standards dealing with Quality of Service assurance. The business-to-business model, and how the FORM project fits into this model, was also introduced. A strong emphasis was made on the FORM Quality of Service Network Provisioning Assurance and both its business model and its technological architectures were outlined. We showed how the Internet2 Simple Inter Bandwidth Broker Signalling protocol has been used to handle service activation, and how it has been extended with a negotiation protocol in order to optimise network utilisation according to customers demands. We introduced also some implementation issues.

In the two-year lifetime ending on April 2002, FORM has reached its halfway milestone. A first trial has been performed, showing conclusive results for each of the Fulfilment, Billing and Assurance processes, and some first interactions between their different components. Technologies and tools have been assessed, for the application field and the market FORM fits in. As FORM Quality of Service Network Provisioning Assurance is concerned, a complete end-to-end negotiation has been performed onto a virtual network model; and it remains now to implement the activation protocol and the enforcement into the physical network elements.

8 References

[1] Defense Advanced Research Projects Agency, Information Processing Techniques Office: "INTERNET PROTOCOL DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION", September 1981, http://www.ietf.org/rfc/rfc791.txt

[2] Third Generation Partnership Project web site, http://www.3gpp.org/

[3] QoS Forum web site, http://www.qosforum.com

[4] S. Blake, D. Black, M. Carlson, E. Davies, Wang Bell, W. Weiss: "An Architecture for Differentiated Services", December 1998, http://www.ietf.org/rfc/rfc2475.txt

[5] QBone Forum web site, http://qbone.internet2.edu

[6] Business to business directory web site, http://www.geobiz2biz.com/

[7] Telecommunication Management Forum web site, http://www.tmforum.org

[8] W3C: Extensible Markup Language, http://www.w3.org/XML

[9] Bruce Eckel: "Thinking in Java 2nd Edition", June 2000, http://www.mindview.net/Books/TIJ/

[10] Ed Roman: "Mastering Enterprise JavaBeans™ and the Java™ 2 Platform Enterprise Edition", 1999, http://www2.theserverside.com/resources/book.jsp