Technical Framework for Introduction of Network QoS into JANET

JANET QoS Development Project Report

Version 1.1

Authors:
Victor Olifer, Rina Samani

Table of Contents

1. Introduction

This document complements the QoS Policy Framework document [1] and provides additional details on implementing the prototype QoS model on JANET. This document provides guidance to organisations wanting to deploy QoS across their network, in conjunction with knowledge of their own routing equipment to develop the configurations required.

The JANET community uses a wide range of equipment from many suppliers, so the approach taken here has been to describe a set of generic techniques required to implement the QoS model. These techniques are applicable to unicast IP traffic, which is the predominant form of traffic. Implementing QoS for multicast traffic, and configuring specific networking devices such as firewalls and VPN gateways to support QoS, are beyond of the scope of this document. Additional details about router sample configurations can be found in the Test Results document [3] and the final deliverables of the project partners published on the UKERNA web site.

The JANET QoS model will be implemented within the framework of the differentiated services (DiffServ) model using a small number of traffic classes which aim to provide the different service characteristics required by broad groups of applications. Details of these classes can be found in the Policy Framework document [1] and the Description of QoS classes [4].

2. Generic techniques required for implementing QoS

Implementing the QoS model will require the routing equipment to be configured to perform the functions required, using a combination of techniques for processing network traffic. These generic techniques are described briefly below.

Equipment from different vendors may not provide identical functions, but it is expected that an equivalent functionality can be found which is adequate to perform the tasks required. The generic functions are:

- Classification – Information in an IP packet header is examined in order to decide which traffic class to place it in. At the edges of the network (site) this may be protocol type, port number, IP address etc. Within the networks, classification will normally be restricted to examination of the existing DiffServ Code Point (DSCP) marks so that traffic can be scheduled appropriately.

- Marking – once a classification decision has been made, packets are marked with the relevant DSCP information. This identifies them for subsequent handling within the networks.

- Rate measurement – this is typically performed by token-bucket type algorithms, though details of these and their associated implementation and configuration parameters may vary between suppliers and also between equipment from a single supplier.

  Rate measurement is important in two contexts, traffic shaping and traffic policing, which are described below.

- Traffic shaping – this is performed on egress from a domain to ensure that traffic rates conform to those required by a Service Level Specification (SLS) (or other constraint). In order to verify that rates are not exceeded, the rate for each class must be measured. If short-term rates are too high then traffic may be buffered briefly to keep mean rates at or below the required levels. Chronic over-subscription to a particular class may require re-marking traffic to a lower priority class or even dropping the excess. Traffic shaping has consequences such as increased latency and delay variation between packets.

- Traffic policing – this is performed on ingress to a domain to check that traffic rates conform to SLS requirements. Out-of-contract or illegally marked traffic may be either re-marked to a lower priority traffic class or dropped altogether, according to the policy in operation between the domains.

- Congestion avoidance – when congestion occurs, the routing equipment may run out of buffer space to store packets and consequently some will be dropped. To avoid the rapid onset of large packet loss, algorithms such as WRED (Weighted Random Early Drop/Detection) may be applied – these avoid sudden buffer exhaustion for a particular queue by randomly dropping a proportion of packets once the buffer fills past a certain threshold. As a congestion control mechanism, this relies on the TCP feedback mechanism at the source where recognition that a packet has been lost causes a reduction in the transmission rate. Note that this technique does not work for UDP transport which has no such feedback mechanism.

- Queuing and Scheduling – all of the routing equipment needs to be configured to handle the different traffic classes with the correct priority. Once classified, traffic will be placed into one of several queues for forwarding, and a scheduling algorithm will service the queues taking into account their respective priorities.

3. The JANET QoS model

3.1 Overview

It is expected that traffic from organisations connected to JANET will be marked by the time it is forwarded to the Regional Network hosting the connection. The basis on which traffic is marked is a matter of policy for the site concerned, though there may be guidelines that could be used as an initial position. Traffic may be marked at source by the applications, or the site access router may classify the traffic according to site policy and mark it as belonging to a particular class, or a mixture of the two approaches may be used.

Figure 1 shows a view of end-to-end communication between systems located at JANET sites across the various management domains involved. The QoS model is implemented by combining groups of the generic QoS functions described above and deploying them at the points marked on the diagram.

It should be noted that although the diagram focuses on functions at the edges of the management domains, the routing equipment within the domain must also be configured to classify traffic based on the marks in the packet headers, and to schedule traffic handling with queues which are suitably defined and serviced.

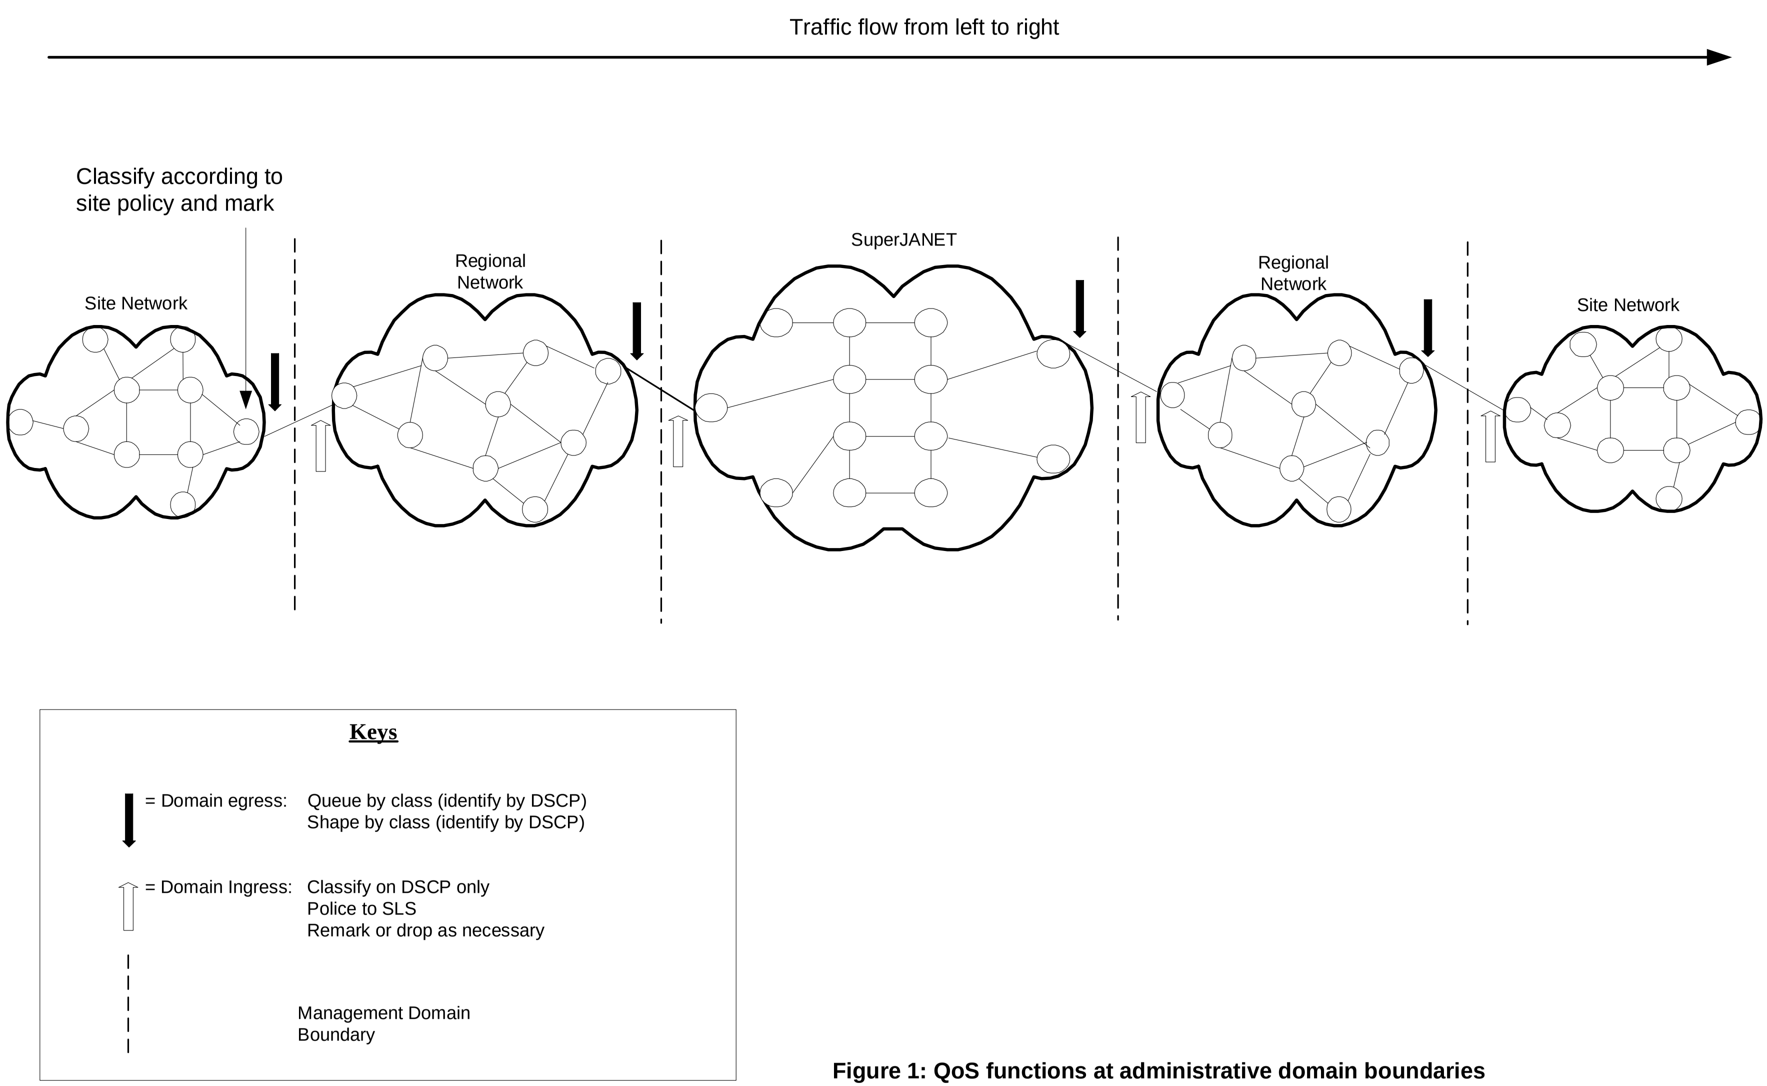

Figure 1: QoS functions at administrative domain boundaries

3.2 QoS classes

There are several well-known QoS classes, which are designed for typical classes of network applications. The most important among them are IP Premium, IP Plus, Best Effort (BE) and Less than Best Effort (LBE).

The IP Premium class of transport service is designed to provide the best QoS service, with strict and low limits for QoS metrics such as packets delays, delay variation (jitter), loss and disordering. These desired values are guaranteed only when the rate of IP Premium traffic does not exceed an agreed upper limit. All time-sensitive traffic, for example Voice over IP (VoIP) and Videoconference traffic, should be served by the IP Premium service.

The Best Effort class is the regular and only Internet-wide service, based on normal processing of IP packets by ISP routers. Routers try to deliver IP packets but without any guarantees, dropping excess traffic during periods of congestion.

The IP Plus (IP+) class is an intermediate service between BE and IP Premium, and could be useful for a number of applications that do not have such strict demands for latency parameters as have applications such as VoIP, but which still need bandwidth guarantees. Typical representatives of this class of applications are Enterprise Resource Planning (ERP) and content-streaming applications.

In contrast to the IP Premium and other QoS services that are privileged relative to the BE regular Internet service, the LBE service has lower than BE priority. Basically it conforms to the DiffServ idea of providing services with *different* QoS levels.

The LBE class is intended for traffic of sufficiently low value (where "value" may be interpreted in any useful way by the network operator), where all other traffic takes precedence over the LBE class of traffic and consumes the bandwidth on the network link. One possible interpretation of "low value" traffic is its low priority in time, which does not necessarily imply that it is of low importance. For the JANET community, the LBE service could be used for bulky GRID traffic that could starve regular Best Effort service of bandwidth because of its high amount of data, i.e. terabytes or even petabytes. When such traffic doesn’t have strict timing requirements it can be transferred during network idle times (which could be quite short, perhaps lasting only a few milliseconds). LBE allows network users to define which traffic could be transferred with the lowest priority.

The participants of the JANET QoS Development Project have not deployed the IP+ class, so the remaining sections of this document only describe the IP Premium, Best Effort and LBE class of services.

3.3 Behaviour of the QoS model

In broad terms, the overall effect desired is for higher priority traffic classes to displace the lower priority classes up to the level at which they are provisioned, but to permit lower priority traffic to make use of bandwidth available in the absence of higher priority traffic. In implementing this behaviour, the following features need to be maintained:

- Protect the characteristics of IP Premium traffic up to the level provisioned.
- Prevent the starvation of BE and LBE traffic by an excess of IP Premium traffic.
- Prevent the starvation of BE traffic by an excess of LBE traffic.
- Give LBE a minimum amount of bandwidth to protect TCP sessions from breaking during periods of congestion.

- Process out-of-contract traffic by either dropping it or re-marking it to place it in a lower priority traffic class.

Implementing this behaviour requires the creation of queues with appropriately tuned packet drop characteristics which are then serviced using a suitable scheduling algorithm.

3.4 Congestion and packet loss

Congestion control will be implemented by dropping packets. To retain the relative performance characteristics for each class, a scheme such as WRED for TCP-based traffic (usually served as BE and LBE) will be used to drop lower priority packets preferentially as congestion develops. It is expected that tuning the parameters for the algorithms will be an iterative process, with refinements made throughout the project.

The use of Explicit Congestion Notification (ECN) is not foreseen at present.

4. Implementing the QoS model

The following sections outline the main tasks that need to be performed in each domain to configure routing equipment to implement the QoS model. These functions are also summarised in figure 1.

4.1 All routers in all participating domains

- Configure the traffic classes. Most modern routers and switches support policing and scheduling on the basis of user-defined traffic classes. The simplest method of creating QoS classes is based on DSCP values, e.g. IP Premium packets have a DSCP mark equal to 46 in the DS field; the LBE class consists of packets with a DSCP mark equal to 8.
- Configure queue scheduling.
    - For the IP Premium class the best way of scheduling is to direct the traffic into the highest priority queue. Most vendors of networking equipment provide this kind of scheduling. It can have different names but the crucial point here is that this queue will be emptied by schedule before any other queue of a router or switch. Cisco strict priority (low latency) queuing is an example of this. In the absence of priority queuing support, weighted scheduling for IP Premium traffic could be used instead. In this case the weight allocated for IP Premium should be twice as big as its maximum agreed policed rate to accommodate a possible burstiness of traffic without delays.
    - LBE should usually be processed in a weighted queue with a weight of 1% - 3%. This reserved amount of bandwidth is usually enough to prevent the breaking of LBE sessions during congestion created by IP Premium and BE traffic.
    - BE will receive all unreserved bandwidth during congestion.[1]

- [1] . Please note that some routers/switches can have a default reservation for the network control traffic; in this case BE traffic will receive an unreserved bandwidth minus this default level.

Whenever policing or shaping functions are required, appropriate rate-measurement features must be configured and enabled.

4.2 JANET Backbone

The core and Backbone Access Routers (BARs) on JANET are configured by the JANET Network Operations Service Centre (NOSC); the approach taken to implement QoS is identical to that described for the Regional Networks below.

4.3 Regional Networks

Edge routers

Ingress traffic:

- Classify packets by the DSCP mark (or mark them if the site networks do not carry out this function).
- Police according to the SLS, dropping or re-marking as out-of-contract or illegal traffic.
- Allocate packets to a queue appropriate to the class.

Egress traffic:

- Classify packets by the DSCP mark.
- Allocate packets to a queue appropriate to the class.
- Apply shaping to each class to conform to the appropriate SLS (site or JANET)

Internal routers

- Classify packets by the DSCP mark.
- Allocate packets to a queue appropriate to the class.

4.4 JANET Sites

This section assumes that a JANET site will only implement QoS functions on the router which connects it to a Regional Network (its access router). It may implement a more complex internal model based on similar principles to those described above for the Regional Networks and the JANET backbone, but details will be a matter of local policy and are beyond the scope of this document.

Site traffic outbound towards the Regional Network or JANET:

- Classify traffic according to local policy and allocate to a traffic class;
- Mark packets according to the class chosen;
- Allocate packets to a queue appropriate to the class;
- Apply shaping to each class to conform to the SLS with the RN.

Inbound traffic from the Regional Network or JANET:

- Classify packets by the DSCP mark;
- Police according to SLS, dropping or re-marking out-of-contract or illegal traffic;
- Allocate packets to a queue appropriate to the class.

5. Testing and Measurement

Testing allows the user to evaluate the correctness and efficiency of the deployed QoS model. The testing should include:

- Subjective assessment of the quality of the applications
- Objective measurement of the application traffic against the QoS metrics

For voice applications, the subjective parameters of quality degradation include the distortion of speech, e.g. echo, cracking, word clipping. The quality of video applications can be assessed by the slowing down or freezing of motion, image distortion and so on.

Objective measurement allows users to obtain the quantitative values of the QoS metrics: one-way delay (OWD), inter-packet delay variation (jitter) and packet loss. These metrics describe the negative effects of queuing on packet transmission by the network devices during congestion. OWD characteristics involve the packet delays between a sender and a receiver that have occurred because of the waiting in router queues. Jitter characteristics involve the variation in the arrival time between packets, which is a result of the stochastic nature of the queuing processes. Packet loss reflects a ratio of packets dropped by the routers because of a buffer overflow.

To conduct objective measurements, a test infrastructure would need to be deployed. This infrastructure should include the following elements:

- a set of real applications that can test the QoS functions
- traffic-generating tools that can generate both TCP and UDP test traffic
- a monitoring infrastructure with probes and collecting agents.

The monitoring infrastructure should consist of tools for both active and passive measurement, which are installed in the same subnets as the end nodes of the applications that are to be tested.

Traffic-generating tools are required to create artificial congestion when natural load on the network links is quite low. Traffic generating tools are only required when verifying the QoS configuration on the routers within the network.

Details of the test applications and measuring techniques used for the JANET QoS Development Project are described in the Test Plan [2].

6. References

1. Policy Framework for Introduction of Network QoS into JANET.
2. JANET QoS Development Project – Test Plan
3. JANET QoS Development Project – Test Results
4. JANET QoS Development Project – Description of QoS classes

The above documents can be found on the JANET QoS website:
http://www.ja.net/development/qos/