Security and Quality of Service Interactions

Panel Chair

Susan Hinrichs, Cisco Systems, Inc.

Panelists

Klara Nahrstedt, University of Illinois

John McHugh, CERT Coordination Center

Partha Bhattacharya, Cisco Systems, Inc.

Abstract

Security and quality of service (QoS) are two critical network services in today's inter-networked world. Security mechanisms are used to provide proof of identity, preserve protected information, and ensure that information received has not been tampered with. Quality of service enables multi-media and other real-time services to use public data networks instead of a more expensive dedicated networks.

Security and quality of service mechanisms are not independent. Choices of security mechanisms impact the effectiveness of quality of service and visa versa. Quality of service requires security mechanisms to ensure appropriate service assignment and billing. Poor security mechanism selection and placement can reduce the performance of a carefully queued network. Inappropriate service level selection can leak extra information about the importance of packets in the traffic stream, but clever manipulation of quality of service parameters might even help to reduce leaking of information through covert channels.

Without a good understanding of these interactions, poor network design choices may result in weaker than expected security and/or less effective quality of service guarantees. Therefore, both services must be considered together when designing and implementing a network infrastructure to achieve the best possible security and quality of service levels.

This panel session will be geared for attendees interested in network management and design. In particular, this session will be of interest to attendees responsible for the security and/or quality of service aspects of network design and management.

The panelists' backgrounds span the areas of industry research and development, government research, and university research. This variety of perspectives should provide an interesting range of insights. The panelists will describe their experiences in addressing the congruence of security and quality of service enforcement. Each panelist will give a brief presentation and the audience will be encouraged to interact with questions and their own observations and experiences.

Can Network QoS and Security Live in Symbiosis?
(Position Statement)

Klara Narhstedt
Department of Computer Science
University of Illinois, Urbana-Champaign
klara@cs.uiuc.edu

Network quality of service research and deployment has focused for several years on problems such as bandwidth guarantees, loss rate, jitter, end-to-end delay and other performance-related quality guarantees when transmitting data over the Internet. Results in network research, such as weighted fair queuing, integrated service framework with its Resource Reservation (RSVP) protocol, differentiated service framework with its service classes for premium, assured and best effort class differentiation, and their standardization through IETF and ATM From, demonstrate the strong emphasis on these issues. Rarely would security be mentioned in this area, and it is a silent assumption that, if one wants network performance and QoS for the data traffic (e.g., multimedia), security can't be part of the equation.

On the other hand, various events in the past, such as serious security holes in the operating systems on routers, denial of access attacks on web servers, intrusion into the routers changing RSVP parameters, and others, shook the networking community. The question is out there if network QoS and security are still orthogonal to each other or should one consider security as another QoS parameter and integrate it with the performance-related QoS results. So the main question is "Can network QoS and security live in symbiosis or not"?

Our belief is that network QoS and security can live in symbiosis if security is put in the right places and at the right time. Problems such as protection of crucial QoS parameters during connection setup, protection of data packets during their transmission in a timely manner, protection against intrusion and denial of service attacks are only some issues which security and QoS need to consider when marrying each other. If security mechanisms, such as authentication, access control, encryption, denial-of-access-sensitive admission control, are enforced during the QoS connection setup, this should be sufficient to distribute the QoS requirements and provide proper resource reservation/allocation/access in a secure fashion. If security mechanisms and policies at routers, gateways and fire-walls, such as intrusion detection, digital signature and encryption with variable key lengths, scalable key management, water-marking, security policy management are available, this could provide for a secure transmission path, content protection and end-to-end QoS provision.

We will show two different examples of integration between QoS and security: (1) authentication security approach, placed into gateways and performed during multimedia transmission phase, which violated end-to-end QoS, and (2) QoS parameter protection security approach, placed into routers and performed during setup phase, which protects end-to-end QoS setup for multimedia transmission. Both examples will argue for symbiosis of security and QoS, however a careful selection of QoS and security mechanisms and policies at the right place and right time will be emphasized.

Security and Quality of Service Interactions
(A Position Statement)

John McHugh
CERT Coordination Center
jmchugh@cert.org

In its original conception, the internet was egalitarian with respect to service guarantees. Delivery on a best effort basis generally meant that the likelihood that a given user's packets would be dropped due to an overload on a segment was proportional to the user's contribution to the load. Since all users were assumed to be friendly, denial of service due to deliberate overloading was not considered to be a real threat. At the same time, it was generally agreed that the protocols were not suitable for applications that required hard service guarantees. The limited bandwidth of the original ARPA Net (56Kbps as recently as the mid 1980s) precluded most real time communications applications such as voice and video which were mostly carried on dedicated circuits.

As the bandwidth of data channels increased and transmission latencies were reduced, it became feasible to consider adding services with strict latency and jitter requirements to the internet traffic mix. One- and two-way audio and video are good examples. For these services to be considered usable, both the time between transmission and delivery (delay) and the regularity with which delivery occurs (jitter) must be carefully controlled. This is often done by reserving the resources necessary to ensure that the delivery goals are met. This can interact with security in a number of interesting ways:

- Services requiring assured delivery can deny service to services that are security (but not QoS) critical by reserving excessive resources.
- The protocols used for negotiating QoS agreements may be subject to attack or interference by non-participating parties.
- Security services such as encryption can prevent delay requirements from being met by introducing additional latencies. Algorithms whose timing is data dependent may introduce additional jitter, as well. Irregular operations such as re-keying may do this also.
- Security services can benefit from QoS measures, as well. To the extent that QOS measures limit delay and jitter, control of such features as a covert signaling measure is depreciated.
- To the extent that QoS operates under a business model that requires assurance of network management services for provisioning, auditing, and billing, the QoS mechanisms may well take advantage of existing network security services.

Both QoS and security are resource management problems and conflicting demands for limited resources are to be expected. Prior experience with similar problems indicates that the treating the conflicts as a risk management problem and applying the risk driven process model (the Spiral Model) developed by Boehm at TRW is a useful way to design and build systems that have conflicting requirements. Under this approach, risk factors, such as the resource conflict between QoS and security services, are identified at each stage of the development from requirements gathering to deployment and maintenance. Development does not proceed until an adequate risk mitigation has been worked out. Risk mitigation techniques that are applicable to resource allocation and performance conflicts include analytical models, simulations, and prototyping. Although the model has not been applied to network QoS problems as far as we know, it has been successfully applied to other security related resource allocation problems including the ABM battle management problem and a high assurance windowing system.

Security and Quality of Service Management
(Position Statement)

Partha Bhattacharya
Cisco Systems, Inc.
pbhattac@cisco.com

Quality of service (QoS) and security services are both vital and affect the entire network infrastructure. While both services are necessary for safe and adequate network operations, in many organizations separate groups are responsible for security and QoS. However, security and QoS implementations will have an impact on each other. Without information about QoS requirements, a poor choice of encryption endpoints may reduce the effectiveness of QoS performance queuing. Without information on security requirements, a poor assignment of QoS performance levels may lead to denial of service for vital but low bandwidth data.

Therefore, QoS and security requirements must be considered together, but it is quite difficult to find people who are expert in both areas. A network policy framework can fill this expertise gap and identify conflicts in security and QoS requirements. Security and QoS requirements can be entered in to the policy framework through a single organization policy, or the security policy and QoS policy can be entered separately. In either case, if the policy framework has sufficient information about the network system, security requirements, and QoS requirements, the framework can resolve or at least identify conflicting requirements.

Enforcing both security requirements and QoS requirements can be viewed as resource allocation problems. When the policy framework is the single point that is solving both resource allocation problems, conflicts can be found or allocations can be altered to deal with the global set of requirements. When security or QoS requirements are considered separately some resource allocation decisions can be arbitrary. For example, when considering encryption requirements, two routers in the network may satisfy the security requirements equally well, but when QoS requirements are also considered, the choice may not be so arbitrary.

Current policy framework systems can adequately deal with static resolution of requirements for security or QoS. It is not a big leap to deal with security and QoS together. The policy framework systems will have to continue to evolve to deal with interactions between administrative domains, more dynamic network requirements, and new network services.