Multi-Facet Internet Resource Management System*

H. Liu
Dept. of Elec. and Comp. Engineering
University of Massachusetts Dartmouth
285 Old Westport Road
North Dartmouth, MA
USA
{hliu@umassd.edu}

H. H. Dempsey
Network Management Department
GTE Internet (formerly BBN)
40 Sylvan Road
Waltham, MA
USA
{hdempsey@bbn.com}

Abstract

The Internet, like most packet switched networks, is vulnerable to aggressive behaviors of emerging multimedia applications. Such deficiency can be overcome through the intervention of a network manager. The manager makes a management strategy aimed at improving the entire network's performance and guides self-optimizing users to achieve their individual performance objectives in harmony.

This paper presents an Internet resource management system with multiple facets that provides tailored functions and associated user interfaces to a hierarchy of users. Users assure their Quality of Service guarantees from the Internet by resource reservation through the system, whereby a higher authorized user acts as a manager to guide other users in their resource reservations. The work focuses on functionality that the system can have. Our real experiments show that the system works effectively with the dynamics of the Internet's distributed resource reservation.

Keywords

Management of the Internet resources, Quality of Service (QoS) management, case studies and experiences

1. Introduction

The Internet offers a single class of "best-effort" service. Best-effort service requires no admission control and involves no resource reservation, therefore, the Internet has no assurance about when or whether packets would be delivered. This basic architecture fulfilled the fundamental goal of effectively interconnecting data networks and supported a wide variety of data applications [1].

Now, the Internet is facing new challenges as it is commercialized and become globally ubiquitous [2, 3]. Multimedia applications like interactive voice and video

*This work was co-sponsored with BBN by NSF/GOALI Award NCR-9612852. The views expressed are of the authors and do not represent the position of NSF or BBN. The paper is approved for public release-distribution unlimited.

require real-time and high bandwidth, they also involve multicast of disseminating data to multiple recipients at once. When running over the current Internet, multimedia applications not only do not always perform adequately due to delay variation and packet loss, but they also interfere with the traditional data applications. The technical foundation of the Internet needs change, extending its best-effort service model to include integrated services with quality of service (QoS) assurance [4]. The first augment is a technology to set aside needed bandwidth for multimedia applications, called Resource Reservation Protocol (RSVP).

Traditional data applications tolerate packet delays and packet losses. Moreover, these data applications, with the underlying transport protocol TCP, guard the Internet from congestion collapse by decreasing their own transmission rates in the presence of network congestion. Multimedia applications, on the other hand, optimize their individual performance objectives, aggressively without taking the overall health of the network in consideration [5].

One approach for keeping the overall network health under noncooperative environment is to employ a network manager. The manager acts as a leader, imposing its strategy on the self-optimizing followers in a game. It is proved that, in certain conditions, such a maximally efficient strategy exists uniquely and can be specified explicitly for the manager to drive the network into its global optimum [6]. This approach is superior to other methods, such as incentive mechanisms that need a new pricing component to the networking structure [7], isolation and regulation that require a priori design decision on the network service disciplines [8], or capacity allocation that demands a priori design decision on the resource configuration [9]. This manager approach architects the network operating point at run time, i.e., during the actual operation of the network. Thus, it is more flexible to the reservation dynamics in the integrated services of the future Internet.

This paper presents an Internet resource management system with multiple facets. Multiple facets are user interfaces to a system with different regions of geographical scope and different levels of access right. To manage large-scale and distributive-administrated networks like the Internet, a system should be able to present focus points of a partial network for the regional manager, in addition to a bird-eye view of the entire network for the top manager. To confine end-users/applications to the management strategy in competing the Internet resources, a system should be capable of restricting access rights on fly to its users by higher authorized users. This work concentrates on the functionality of the system. The effectiveness of the system for managers to pilot self-optimizing users achieving the entire network stability will be our further study.

The paper first shows RSVP’s vulnerability to end-user/application misbehaviors. Then, it explains how a multi-facet resource management system on top of RSVP confines end-user/application behaviors. Each facet of the management system is presented next. Finally, it demonstrates, through several real experiments, the versatility of the system functionality for the Internet resource management.

2. RSVP – Internet Resource Reservation Mechanism

The resource ReSerVation Protocol (RSVP) [10] is a signaling protocol that enables senders, receivers, and routers to exchange information for setting up resource reservations in order to support the services with QoS requirements. RSVP identifies a communication session (either multicast or unicast) by the combination of destination address, transport-layer protocol type, and destination port number.

Figure 1 depicts an RSVP-capable device, host or router, which has two parts. One is data forwarding with Classifier and Scheduler. Classifier determines route as well as selects data packets to use the reserved resource. Scheduler allocates the resources with which to transmit the data packets. The other part is background control with RSVP Daemon and Admission Control. RSVP Daemon handles reservation messages. Admission Control accepts or denies the reservation request, i.e., Resv message. If accepted, RSVP sets the parameters to Classifier and Scheduler for resource reservation; if denied, RSVP returns an error indication, i.e., ResvErr message, to the receiver's application. It is important to note from the figure that the Resv messages are receiver initiated and follow the reverse data-forwarding path. The reverse data-forwarding path is previously set up by the sender-issued Path messages that may contain information for receivers to figure out their end-to-end QoS requirements.

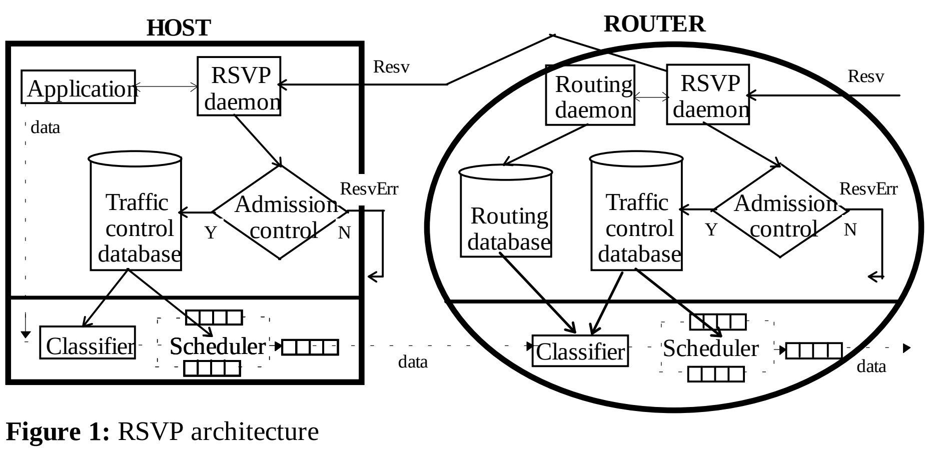

Figure 1: RSVP architecture

An end-user/application can use RSVP to assure from an integrated services Internet the end-to-end QoS guarantee by expressing its QoS requirements to the Internet and having the Internet promise to support such [11]. The information about QoS requirements is conveyed in the Resv message mentioned above. RSVP targets to operational issues regarding resource reservation, through which individuals can strive to meet their performance objectives. However, RSVP does not address management issues, leaving the Internet unprotected from abuses.

For example, reducing retry interval lowers delays for the end-user/application to establish a session. RSVP, for the sake of controlling protocol overhead, prefers larger timer values (or teardown delays) in refreshing soft state. When teardown delay greater than retry interval, a retry for reservation setup can build upon the earlier partial reservation still in place, which would be superior to that for the end-user/application to contend for new resources at every link. The scenario left some resources reserved but not used, and hence wasted. Such wasted resources can be accumulated, when the network load is sufficiently high, that degrades the entire network performance by blocking other users' reservation requests.

Incremental retry policy in multicast, of which only blocked receivers issue new reservation requests while successful receivers keep their reservations until a session being established, is another example. The incremental addition of session receivers always results in fewer retries than that of all receivers retry upon failures by some of them, thus it is beneficial to the end-user/application. However, the scenario could cause deadlock – some hold resources forever in waiting for others to obtain resources that they are holding. The network throughput drops to zero when the entire network deadlocks.

Thus, to keep the future Internet healthy, work must be done regarding the management aspects of RSVP.

3. An Internet Resource Management System

The ProVision Group at the Network Management Department of BBN Planet has developed an Internet resource management system on top of RSVP [12]. BBN Planet, now part of GTE Internet, is one of the largest commercial Internet service providers in USA. The ProVision project is to package a value-added service to subscribers of BBN Planet, which provides reserved and multicast bandwidth-on-demand through RSVP management. The ProVision Group is well equipped with the future Internet testbeds and the project goes through a rigid cycle of software development and deployment.

Figure 2 shows where the management system fits in the RSVP architecture to manage the Internet resources. There are two ways of their interactions. In one way, the RSVP management system presents to the manager a high-level view of the Internet resource state in the form of dynamic visual abstractions. The abstraction helps the manager to understand the situation and provides the basis for the manager to make his decision. The manager can then tune resource reservation for network performance without concerning specific mechanisms and their implementations used in resource reservation. In the other way of their interactions, the manager, through the RSVP management system, manipulates a set of management parameters that are translated into control parameters. Management parameters trade-off between entire network bandwidth utilization and individual users' QoS guarantees. Samples of management parameters are resource utilization, QoS constraints, and network properties. Control parameters influence the behavior of the RSVP components, and therefore, change the way the Internet resources are reserved. Control parameters could be admission bounds, traffic priorities, and teardown delays. This split between management parameters and control parameters came from the idea of separating computational task from operational task, which makes it effective to manage the Internet resources supporting integrated services [13].

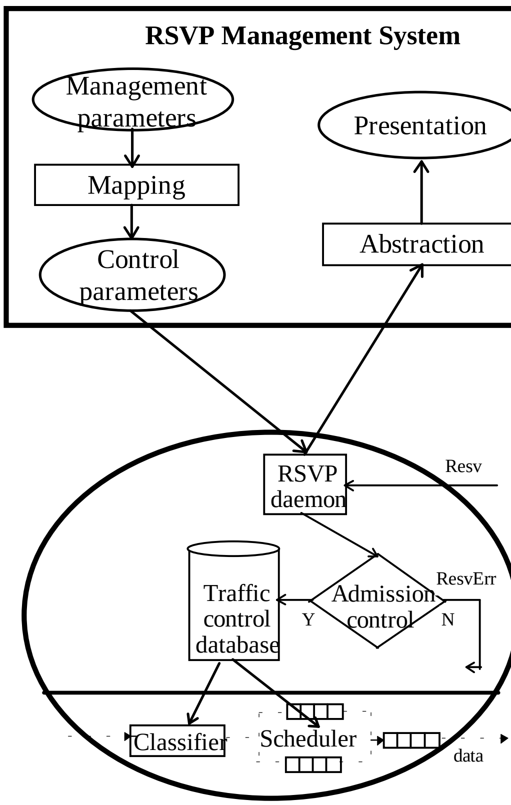

Figure 2: System interaction with RSVP

The RSVP management system contains three major parts as shown in Figure 3. Modeler is the central part, which holds the abstractions of the Internet resource state and maps users' management parameters to devices' control parameters. Abstractions model network entities with status reported from network components and/or controlled access by network operators. The RSVP MIB [14] provides a simple and quick operational framework for defining some primitive network entities. Modeler also provides what-if scenarios for users to analyze their decisions before actual control parameters are issued. The underlying network technology is ever changing and application requirements are unclear at the early stage of the integrated services Internet, therefore, object-oriented (OO) approach is suitable to develop modelers for its design- and code-reusability.

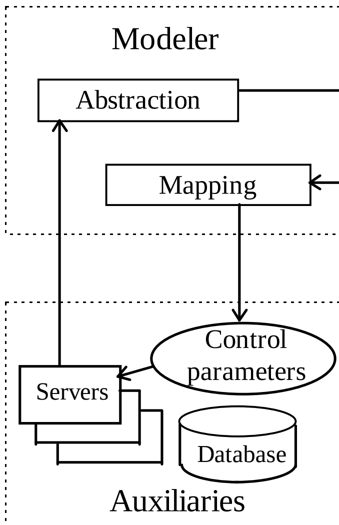

Figure 3: System architecture

The next part is Auxiliaries that implement control functions of the Internet resource management. It contains a database and several servers:

- User Database that keeps user realm, records management strategies and user-defined functions, and tracks user's reservations and usage.
- Configuration Server that retrieves and/or changes the RSVP configurations.
- Event Server that records all the events occurred, such as failures of reservation requests or reaching thresholds of resource usage.
- Security Server that serves as a web server with security features of controlling access to the information on the Internet.

The Graphical User Interface (GUI) provides functional accessibility tailored according to the authorization level of the user classes. It presents multi-facets of the Internet resources and allows different levels of resource management.

The three parts of the RSVP management system and their components reside on different processes running on dispersed hosts. These processes are organized in a distributed-hierarchical paradigm. Notice that the distributed-hierarchical paradigm only applies to resource management for services with QoS requirements, the

Internet remains its distributed-cooperative paradigm to best-effort service. As of hierarchical (i.e., centralized), a manager is in charge of its own domain, which could be divided into sub-domains maintained by middle managers who have dual roles of manager and agent. As of distributed, when some global knowledge is needed in the process or a problem comes from outside his own domain, a manager can delegate his peer managers to retrieve information or to fix the problem [15]. This paradigm has the advantage of simplicity and efficiency featured by centralized paradigm and scalability and reliability rendered by distributed paradigm.

For details on the implementation of the ProVision's Internet Resource Management System, refer to our NOM'98 paper [12].

4. Multi-Facets of Management

The multi-facets of the Internet resource management work in the following way. The users are classified by authorization levels. For example, the ProVision's System has four classes: Network-Analyst, NOC-Operator (where NOC stands for Network Operation Center), Site-Administrator, and End-User. Users of the four classes try to obtain their respective performance objectives by reserving the Internet resources through the System while they follow the rules set by the higher authorized users to keep the Internet stable. A NOC-Operator acts as a domain manager who is aware of the noncooperative behavior of the users and makes its management strategy to improve its overall network performance. Site-Administrators as middle managers, on one hand, supply their sites' capacity coincident with the global state to complete end-to-end QoS delivery, on the other hand, they pilot the End-Users of respective sites to optimize their sites' local performance. End-Users converge their individual performance objectives to the global optimum under the guidance of the management strategies. Network-Analysts work on background to support NOC-Operators in diagnosis and trouble-shooting. The functions of the ProVision's System are categorized into the five areas based on the ISO (International Organization for Standardization) Network Management Forum [16].

Each facet has a unique set of functions and an associated BUI (Browser-based User Interface) to access them. A facet inherits all the functions from its lower authorized facet; for example, a site-administrator has all the functions of an end-user besides his additional functions.

4.1 End-user facet

After a user logins into the System, he is prompted with a BUI tailored to the class that he belongs to. The BUI has access to a RSVP setup tool offered by the Configuration Server and customized to the user's class. For an End-User, the RSVP setup tool is a Speed-Dial panel shown in Figure 4. The Speed-Dial panel contains a list of Speed-Dial sessions, each of which maps a set of RSVP directives and parameters pre-configured so that the end-user can skip the cumbersome computations involved in the RSVP management and tedious process of RSVP setup. The end-user simply chooses the session that he is interested in and the System will establish the reservation for him to tune in the session. The end-user is allowed to make some changes on the parameters within a pre-specified range.

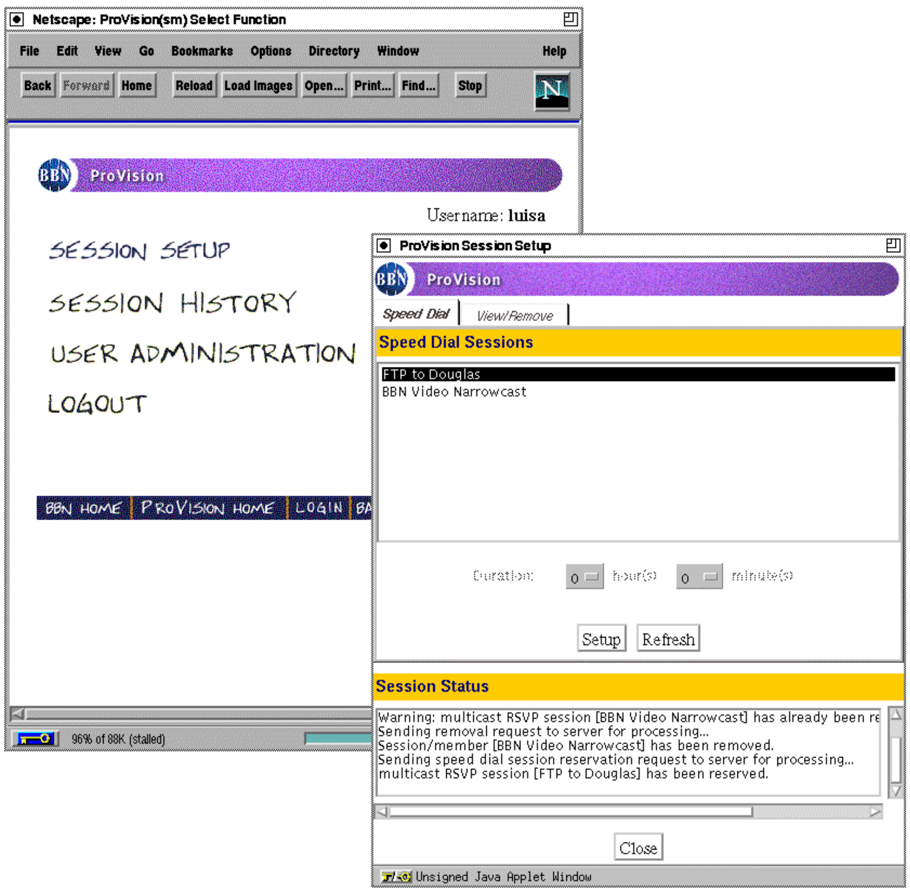

Figure 4: Sample BUI for end-user

The ProVision's System monitors reservation activities and their results. At low level, router status is polled periodically and any significant changes will result in an event, which will be recorded in the Event Server. At high level, monitoring tools will notify state changes as events reflecting particular tasks and filter events to appropriate users based on their classes and access rights. When a setup tool issues a reservation request, it will initiate a monitoring process that checks the affected resources for expected behavior; if there is a mismatch, an event is generated to notify the NOC-Operator and the appropriate Site- Administrator.

Once a reservation is established, the monitoring process continuously checks the health of the reservation – any changes, such as joining/leaving the set of receivers, will generate an event. Significant events that might indicate a problem will raise an alarm to the NOC-Operator for proper handling. For an End-user, the fault/performance management is an Event Window that lists all the events relevant to him.

User identities are authenticated with a login and password mechanism. Users are unable to access the System's URLs until they have logged in. Once logged in, a user's state is preserved so that the user's ID is known for subsequent accesses. All transactions between the web server (i.e., the Security Server) and browsers are encrypted and therefore, no one snooping on the Internet can steal the information.

4.2 Site-administrator facet

A site-administrator has full and sole control of RSVP in his site routers. As the consequence, a RSVP setup session may define a subset of a reservation, which needs to be completed by a NOC-operator and possibly other site-administrators. In a case of crossing regions, a site-administrator may only have a setup session with partial reservation in place along the path while the rest has to be done by others before an end-user can access it. In another case of distributing contents, a sending site-administrator may simply have a setup session with Path messages installed, no reservation will be made until receiving site-administrators are ready to join in.

If a site-administrator permits, a NOC-operator or another site-administrator can set up a reservation remotely by entering this site upon authorized. It gives the convenience for a sending site-administrator, in the case of distributing contents, to establish a reserved channel for all his receivers.

A site-administrator has authority to define Speed-Dial sessions and to define which users under his supervision are allowed to use particular sessions. Figure 5 shows a typical BUI for Site-Administrator, which includes tools for Speed-Dial editing and testing, adding/removing users, and authorizing remote administrator.

There are a number of monitoring tools for a Site-Administrator besides the Event Window:

- Distribution Tree – It presents a map of the tree from the sender to the receivers of a specific multicast address that obtained the reservation successfully;
- Receiving Tree – It shows a map of the reverse tree from the receiver back to the senders of his interests, that succeeded in the reservations;
- RSVP Trace – It is a map of the path from the sender to a specific receiver who obtained his reservation successfully;

- Busy Nodes – It shows those routers of which CPU/memory utilization, or routing capacity (in packets per second), reached thresholds defined by the site administrator;
- Busy Links – It shows those links of which reservation utilization reached user-defined thresholds.

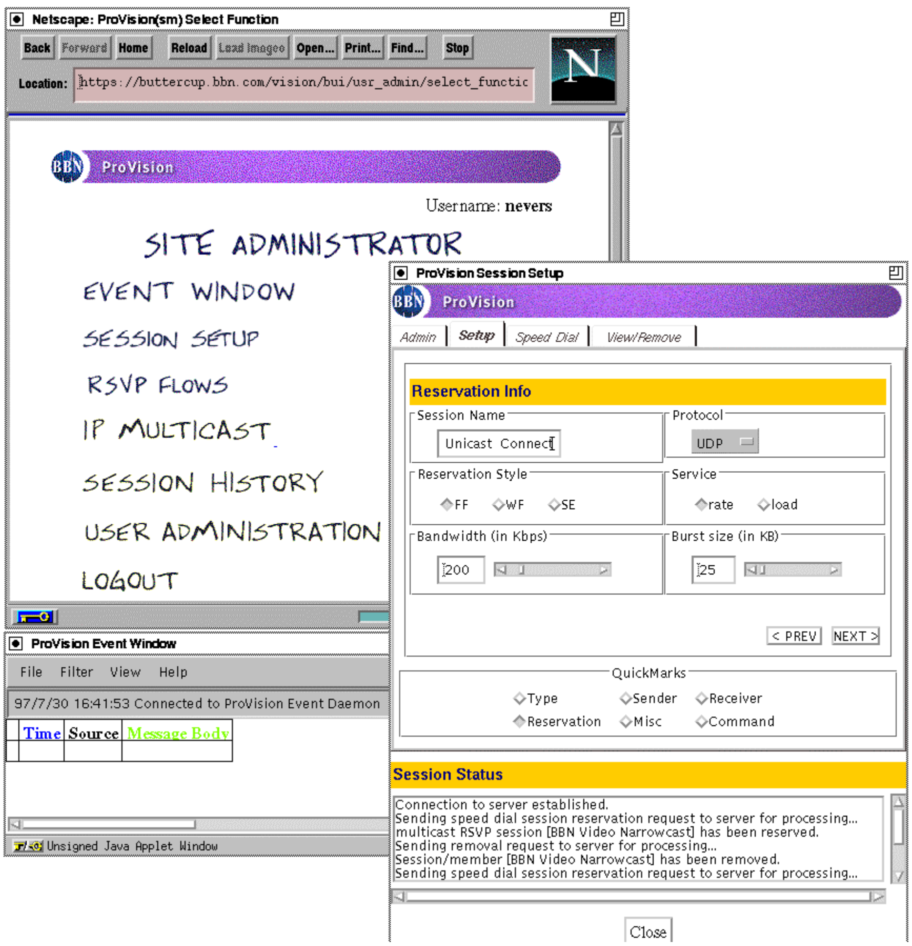

Figure 5: Sample BUI for site-administrator

4.3 NOC-operator facet

A NOC-Operator is supported with tools to facilitate retrieving and presenting the fielded configuration values in routers, comparing them with what expected, and updating the values if needed to assure correct router configuration. In addition, a NOC-Operator can enable/disable RSVP for some routers in overall or at individual interfaces to assure that a site is able/unable to establish reservations to certain regions.

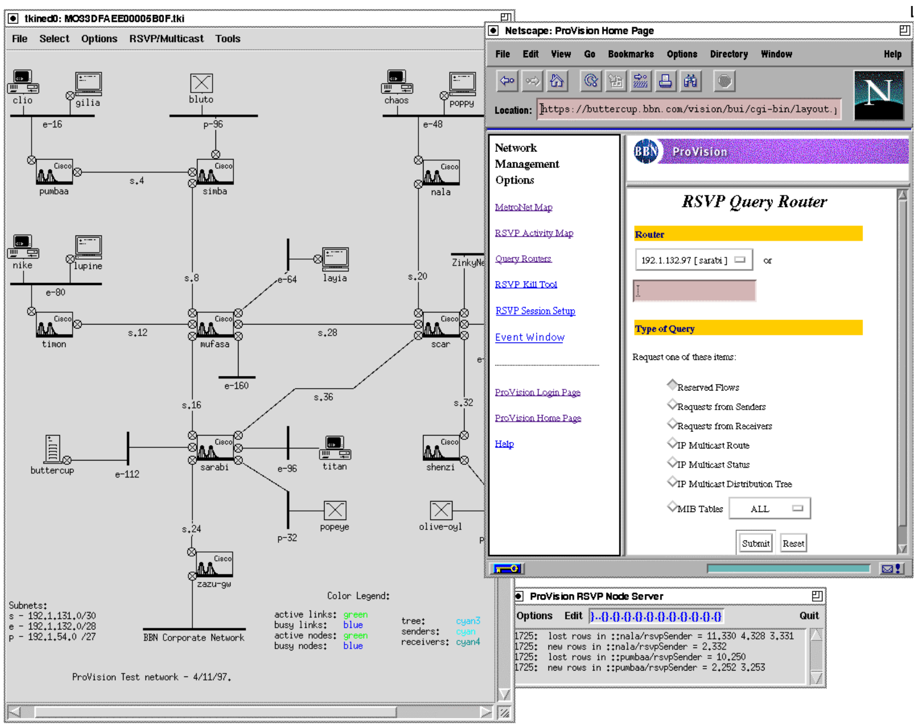

Figure 6: Sample BUI for NOC-operator

Figure 6 shows a typical BUI for NOC-Operator. A NOC-Operator has a SOS tool through which one can manually set up or remove (i.e., kill) a reservation. A NOC-Operator can view his entire network using the monitoring tools. For security, a NOC-Operator can disable sites or patrol for a site to identify unusual activities.

5. Field Test Runs

The ProVision's Internet Resource Management System has undergone several field runs on BBN Planet by configuring facets on fly [17]. Each experiment is done by comparative observations on two identical applications running side-by-side, one with and the other without RSVP setup. Our observed results of RSVP performance are encouraging. We found that translation of QoS parameters into RSVP control parameters could be tricky, which we will study further [18].

5.1 Predictable file transfer

Predictable File Transfer (P-FTP) is a value-added FTP service. It provides predictable performance in terms of elapsed time for point-to-multipoint file transfer in the face of unpredictable Internet load. Files are gathered together and are transferred in batch over the Internet at a pre-scheduled time from one site to a set of other sites and the job is expected to finish by certain time. All the sites are mutually trusting. The source site only needs to send one copy out for all to receive.

The ProVision's System supports this application by its Speed-Dial feature. The Site-Administrator at the source site defines a P-FTP session, of which the RSVP parameters are calculated according to the size and duration of the job. The source Site-Administrator, as a remote site-administrator to all the destination sites, programs those sites to start the P-FTP session at the pre-scheduled time requesting for resource reservations. A NOC-Operator may have to join the process of completing the resource reservation requests over the Internet that connecting the sites. Upon the prompt of the successful reservation at the pre-scheduled time, files can be transferred using a multicast ftp such as Starburst.

5.2 Narrowcast

Narrowcast is a distribution of real-time video/audio programs to a wide Internet audience, which inspires the name from television "broadcast." Some selected sites are distributors of video/audio programs in channels that are the Internet class D addresses for multicast. Customer sites subscribe to some channels by joining the multicast groups. For security, distribution across the Internet is multicast only to those sites that subscribe to the channels and authentication is done by only allowing the legitimate program packets entering a site as a result of that site subscribing to that channel.

A NOC-Operator creates a narrowcast channel for a distribution site by defining a Speed-Dial session for an Internet class D address to be used by the channel and establishing a partial setup session with only Path messages installed. A customer site subscribes to a channel by requesting a NOC-Operator to enable his site to the multicast group with reservation ability. When the customer is ready to tune in a live channel that he subscribed, he uses the Speed-Dial to set up a reservation and then receives the programs. A customer site can unsubscribe to a channel at any time.

5.3 Distance learning application

Distance learning application uses commercial multipoint room/desktop video conferencing systems over the Internet. It is similar to narrowcast in the way of video/audio distribution for teacher's presentation. It may be simpler because all participants trust each other. It may also be more complex because reverse paths are needed for receivers to raise questions.

A NOC-Operator creates a class session by defining a Speed-Dial session for a course. The course session has reserved resources from the teacher's site to all students' sites for quality delivery of the lecture contents. The multicasts from any students' sites are simply non-RSVP for economics.

6. Summary

This paper presents a multi-facet Internet resource management system developed by BBN ProVision. The system offers tailored functions and associated user interfaces, called "multiple facets," to four classes of users, namely End-User, Site-Administrator, NOC-Operator, and Network-Analyst. Through the system, a higher authorized user enforces his management strategy to users under his supervision competing for the Internet resources to meet their individual QoS objectives without straining the Internet. The system is field-tested over BBN Planet, part of the real Internet, for several specific applications such as predictable file transfer, video/audio distribution, distance learning, and ad hoc multimedia applications. Observations of these experiments show the effectiveness that the system guards the Internet from being drained by multimedia applications.

The technology of the system is based on the leader-followers model of game theory, where a leader makes a rule that confines followers in playing a competitive game. The architectural design of the system separates the computational task from the operational task in managing RSVP to support integrated services over the Internet. This idea coincides with the IETF organizations.

Our further work will be on the effectiveness proof by systematically collecting data and thoughtfully analyzing them. Another further study is to show how the system can compromise the scalability problem of RSVP when the Internet upgrades to very high speed.

Acknowledgment

The authors greatly appreciate all the members of ProVision group at BBN for their good work, which led to this paper. Claudio Topolcic, Tim O'Neil, and many others have offered their insights on RSVP and Internet QoS issues during our numerous discussions, which the authors took for granted. The authors are impressed by the dedication of the IM reviewers and thank the reviewers for their advice and detailed comments, which brought up the quality of this paper

References

[1]. D.D. Clark, "The design philosophy of the DARPA Internet protocols," Proc. of ACM SIGCOMM '88, (Stanford, CA), pp. 106-114, August 1988.

[2]. Feature Topic: The Global Internet, IEEE Communications Magazine, vol. 11, no. 3, May/June 1997.

[3]. http://www.ucop.edu/irc/projects/Internet-II_Arch.html

[4]. S. Shenker, “Fundamental design issues for the future Internet,” IEEE Journal on Selected Areas in Communications, vol. 13, no. 7, pp. 1176-1188, September 1995.

[5]. D.J. Mitzel, D. Estrin, S. Shenker, and L. Zhang, “A study of reservation dynamics in integrated services packet networks,” Proc. of IEEE INFOCOM ’96, (San Francisco, CA), pp. 871-879, March 1996.

[6]. Y.A. Korilis, A.A. Lazar, and A. Orda, “The role of the manager in a noncooperative network,” Proc. of IEEE INFOCOM ’96, (San Francisco, CA), pp. 1285-1293, March 1996.

[7]. R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, “Pricing in computer networks: motivation, formulation and example,” IEEE/ACM Transactions on Networking, vol. 1, no. 4, pp. 614-627, December 1993.

[8]. S. Shenker, “Making greed work in networks: a game-theoretic analysis of switch service disciplines,” IEEE/ACM Transactions on Networking, vol. 3, no. 4, pp. 819-831, December 1995.

[9]. Y.A. Korilis, A.A. Lazar, and A. Orda, “The designer’s perspective to noncooperative networks,” Proc. of IEEE INFOCOM ’95, (Boston, MA), pp. 562-570, April 1995.

[10]. R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource ReSerVation protocol (RSVP) – version 1 functional specification (RFC 2205)”, ftp://ftp.isi.edu/in-notes/rfc2205.txt.

[11]. P.P. White, “RSVP and integrated services in the Internet: a tutorial,” IEEE Communications Magazine, vol. 11, no. 3, pp. 100-106, May/June 1997.

[12]. H. Liu and H.H. Dempsey, “WWW-based Internet resource management system,” Proc of the Tenth IEEE/IFIP Network Operations and Management Symposium, (New Orleans, LA), pp.11-20, February 1998.

[13]. G. Pacifici and R. Stadler, “An architecture for performance management of multimedia networks,” Proc. of the Forth IFIP/IEEE International Symposium on Integrated Network Management, (New York, NY), pp.174-185, May 1995.

[14]. F. Baker and J. Krawczyk, “RSVP management information base using SMIv2 (RFC 2206),” ftp://ftp.isi.edu/in-notes/rfc2206.txt.

[15]. Y. Yemini, G. Goldszmidt, and S. Yemini, “Network management by delegation,” Proc. of the Second IFIP/IEEE International Symposium on Integrated Network Management, (Washington, DC), pp.95-107, April 1991.

[16]. A. Leinwand and K.F. Conroy, Network Management: A Practical Perspective (2nd Ed.), MA: Addison-Wesley, 1996.

[17]. L. Feldkhun, M. Marini, and S. Borioni, “Integrated customer-focused network management: architectural perspectives,” Fifth IFIP/IEEE International Symposium on Integrated Network Management, (San Diego, CA), http://www.comsoc.org/confs/im/97/F.html, May 1997.

[18]. P. P. Pan and H. Schulzrinne, "YESSIR: a simple reservation mechanism for the Internet," Proc. International Workshop on Network and Operating

System Support for Digital Audio and Video (NOSSDAV), (Cambridge, England), July 1998. also IBM Research Technical Report TC20967.