

Test Results of the Commercial Internet Multimedia Trials

Mark Baugher[†] and Saib Jarrar[‡]

This paper reports the test results of applications and services from the Commercial Internet Multimedia Trials. The Trials were a twelve-month effort by three companies to evaluate the product readiness of multimedia applications and services in business environments by supporting multimedia services on production IP networks. The test beds were enabled for IP multicast routing; one of the test beds was enabled for RSVP. The results of our RSVP tests and user surveys are reported, and some of issues related to the deployment of multimedia services and applications on corporate IP networks are considered.

1 Introduction

Large-scale experiments such as Internet Mbone have shown that IP-based multimedia services can successfully support large conferences that integrate audio, video, graphics, text and other media on IP networks and even on the Public Internet [1]. The Mbone experience demonstrates that multicast service can scale to thousands of users for individual events, and that a multicast IP network can scale to thousands of subnetworks. The real-time and synchronized media used in these conferences, however, are run as best-effort flows without guarantees or mechanisms for service quality. This problem has been addressed by the Internet-standard resource reservation protocol RSVP [2].

RSVP and packet-scheduling services in routers and hosts provide bandwidth management for multimedia application flows. To date, experience with RSVP has not approached the scale of the IP multicast experience. At the September, 1995 Network+Interop Conference, six vendors demonstrated router and host implementations of RSVP which provided internetwork quality of service (QoS) for multimedia applications. Thus, an end-to-end solution of IP-based multimedia services appeared for the first time that included product prototypes from major vendors of internetworking hardware and software systems for RSVP, IP multicast, and multimedia applications.

Employees of Cisco Systems, Intel Corporation and MCI Corporation came together shortly after the Interop RSVP Technology Demonstration to plan, build and operate IP network test beds that supported multimedia services. The "Commercial Internet Multimedia Trials," sought to address many of the challenges of taking an end-to-end multimedia solution from a technology demonstration to a prototype service on production networks. The "Trials" consisted of a "Technology Trial" of protocols and services and a "Customer Trial" of multimedia applications, content and tools.

A successful customer trial of new IP services should help answer the many questions of feasibility, customer demand, and whether a positive return on investment (ROI) can be expected. Although there are many factors that contribute to the success of new technology products

[†] Intel Corp., 2111 NE 25th Ave., Hillsboro, OR 97124, mbaugher@intel.com

[‡] MCI Corp., 901 International Pkwy., Richardson, TX 75082, saib.jarrar@mci.com

beyond costs and expected financial returns, it is prudent to assume that positive answers to ROI questions are prerequisite requirements for a commercial product offering of new network services. Thus the approach taken in the Trials was to form a small, focused alliance of vendors devoted to overcoming problems in the deployment and operation of IP-based multimedia applications with the aim of showing compelling content and a positive ROI. The tests were undertaken with the goal of showing how well the technologies worked.

The test results of the Technology Trial and the customer experiences of the Customer Trial are reported in the pages that follow. The next section, *Network Testing*, describes two of the test beds that were used and some of the tests that were run. The *Applications and Customers* section discusses the three types of applications that were run on the test beds along with the survey results of customer participants. The *Results* section summarizes the results of the Trials, and the *Conclusion* considers issues related to the commercial deployment of the services and applications that were evaluated.

2 Network Testing

The search for customers, applications and application tools began concurrently with setting up the test beds. There were three test beds. First, MCI provided use of its commercial HyperStream Switched Multimegabit Data Service (SMDS) public network, and connected four sites in the Western U.S. into an IP multicast/RSVP WAN test bed. Two MCI sites, one Intel site and one Cisco site were on the SMDS test bed. Second, Intel's Information Technology Division agreed to link four Intel sites to an IP multicast/RSVP test bed that overlays its DS3 ATM WAN at Intel. And third, MCI Infolink, the group that maintains MCI's internal IP network, agreed to evaluate IP multicast and multimedia service on a portion of their network for two events.

2.1 The Test Beds

MCI's SMDS network supported a multi-site test bed for multicast and RSVP evaluation. Each of the sites was connected with a DS1 circuit to the SMDS network. Each site was equipped with a Cisco router, Ethernet and several Pentium* PCs. As shown in Figure 1, the use of SMDS as the underlying network to interconnect the test routers provided the flexibility of testing two different network topologies, multi-access and point-to-point topologies. First, the four test routers formed a single IP subnetwork using the inherent capability of SMDS as a multi-access network at the data link layer where each router is a hop away from any other router. The routers used the SMDS multicast capability to transport IP multicast traffic. A second subnetwork topology was later established and tested between pairs of routers in a point-to-point topology. Each router was configured with two sub-interfaces to connect to two other routers. Each sub-interface was assigned its own SMDS and IP addresses and established a virtual point-to-point connection with a peer sub-interface on a different router. This configuration formed a network of four nodes with a square topology as shown on the left side of Figure 1 and it was used to test recovery from route failure. The multi-point configuration is shown on the right side of the figure.




Figure 1: Point-to-Point and Multipoint SMDS Topologies

For the purpose of IP multicast testing, the test routers were configured with the Protocol Independent Multicast (PIM) routing protocol. Both PIM Dense Mode (PIM-DM) and PIM Sparse Mode (PIM-SM) were tested successfully. Sparse mode is recommended in a WAN environment when many sites are interconnected, but only few are participating in a multicast session. Limiting distribution to these sites may conserve WAN bandwidth. For the purpose of QoS testing, the Cisco routers and the PCs ran RSVP to signal QoS requirements. The routers attempt to enforce these requirements by implementing the Weighted Fair Queuing (WFQ) mechanism for outgoing traffic, especially on the serial interfaces connected to the SMDS network.

The RSVP functionality was tested end-to-end between workstations at different sites across the routers and the SMDS network. The QoS functionality was tested on the routers with the SMDS network playing a passive role. There was no intent to involve the SMDS switches in the network QoS testing. SMDS does not support any QoS capability in terms of guaranteed bandwidth. The SMDS network delivers packets on a best-effort basis. In addition, the testing assumed that the SMDS network is not congested and hence is not considered as a bottleneck. In fact, the SMDS network utilizes a DS3 backbone; the testing verified that the backbone is not congested and ample bandwidth was available to carry the test traffic. This was done using Ping and throughput tests at different intervals during the trail. The round-trip delays were consistently below 100 milliseconds and with insignificant variations. The routers were able to achieve close to line speed (DS1) throughput. The DS1 ports/circuits were considered to be the bottlenecks, and congestion was induced on these ports to test the QoS functionality.

However, the fact that SMDS is a multi-access network and does not support QoS combined with the fact that RSVP applies to outgoing but not incoming interfaces presented a problem. This problem is most prevalent in the multipoint configuration in Figure 1. In this configuration, multiple sites may stream simultaneously into a single site resulting in a total traffic load that exceeds that site's link speed (DS1). Consequently, the outgoing interface on the SMDS switch serving that site will become congested. The SMDS switch will drop packets based on a FIFO strategy, irrespective of the QoS requirements associated with different packet flows. Thus QoS could not be enforced end-to-end. This behavior was taken into account in designing the test cases. In the point-to-point configuration, this scenario was avoided by implementing traffic shaping by the routers on their SMDS interfaces. This limited the output bit rate on each interface and ensured that no two sites can overload a single destination site. Obviously, this solution does not scale, but it was sufficient for completing the RSVP tests on the small SMDS test bed.

Intel's ATM test bed, however, featured QoS but not end-to-end RSVP. The RSVP product feature in Cisco's IOS* 11.2 did not support interfaces at ATM speeds at the time of our testing (a more recent product release does support ATM). It was necessary to proceed with the first training classes of Intel factory workers without RSVP service in the Fall of 1996. It was possible, in principle, to over-provision across the ATM backbone. In practice, however, we were limited to running at DS1-speeds by agreement with the backbone service provider, Intel Information Technology Division. As shown in Figure 2, an isolated network made up of ATM point-to-point Variable Bit Rate (VBR) Permanent Virtual Circuits (PVC's) were used for the test bed. A DS1-rate circuit was guaranteed through the ATM network, but this bandwidth could be used for other flows when not used for Trials events. Without RSVP and some packet scheduling mechanism on the VBR PVC, however, real-time and synchronized media competed with other best-effort traffic such as FTP and HTTP on the test bed. FTP, TFTP and HTTP applications were used between the sites for a variety of purposes including software updates. The TFTP server for router IOS updates was located on the Jones Farm site test bed LAN. Unicast routes

from one test bed router to the other naturally went through the DS1-rate VBR PVC's that were used for the multimedia and multicast service.

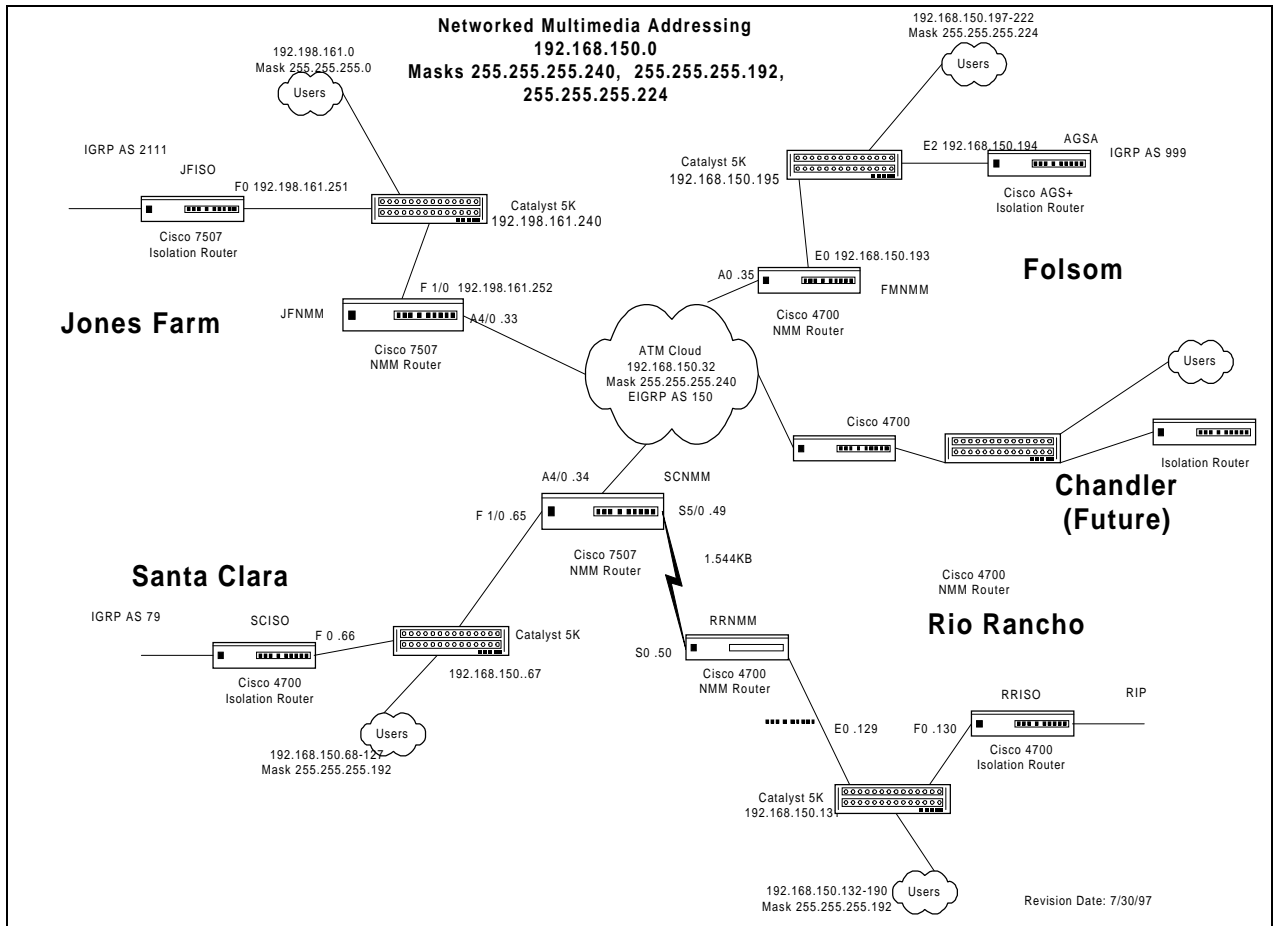


Figure 2: ATM WAN Test Bed

It would have been possible to simply add a second router at each site with a serial line between it and the router connected to the ATM network. This solution would have doubled the cost for each site, however, and negatively affect the ROI – one of the criteria for evaluating the costs and benefits of the service. The chosen solution was to control the traffic using packet filters (Cisco IOS access lists) over the DS1-rate VBR PVC's between the sites as a means control traffic. This solution was messy, inefficient and prone to error. Over-provisioning was a time-consuming practice on this test bed. RSVP would have been a much better solution. As discussed in the Results section, subnetwork bandwidth management is a potential alternative for subnetworks such as our two test beds that do not feature resource reservation.

2.2 RSVP Test Plan and Experiences

An RSVP test plan was successfully completed on the SMDS test bed. A total of fifteen RSVP tests were executed. The tests of Table 1 were selected based upon the network, applications and tools that were being used. RSVP Wildcard and Shared Filters were not tested. In most cases, the Controlled Load service model was used. The Guaranteed Service Model was tested in case 2.3. Some use was made of the router user interface commands for installing and removing

reservations, but all signaling tests also were run from a host. Table 1 lists the RSVP tests that were run on the SMDS network.

Case	Description
2.1	Verify RSVP session establishment, maintenance, tear down with FF reservations for a 1:1 session
2.2	2.1 for a 1:N session where $N \leq 6$ and verify merging
2.3	Verify QoS parameters change for 2.1 and 2.2
2.4	Verify reservation independence for each direction of 2.1
2.5	Verify that hosts and routers can maintain multiple reservations
2.6	Verify cleanup of reservation state when cleanup interval expires
2.7	Verify routers correctly provide reservation confirmation when requested
2.8	Verify admission control rejects request based on bandwidth availability
2.9	Verify policy control function of the router can reject requests based upon sender or receiver
2.10	Verify QoS enforcement by WFQ for single session
2.11	Verify QoS enforcement by WFQ for multiple sessions
2.12	Verify RSVP routers treat excessive RSVP session traffic on a best-effort basis
2.13	Verify traffic shaping reduces ingress port load
2.14	Verify reservation reestablishment after route failure and recovery
2.15	Verify operation when RSVP routers are connected via non-RSVP routers

Table 1: RSVP Functional Tests Performed on SMDS Testbed

There was considerable discussion during the planning of Trials as to whether the testing should be done on a production WAN or in a room. Experience has shown that most of the tests of Table 1 needed to be run on the production network since many of the problems that we encountered were unique to the production network configuration. Many of the tests, however, were first run on a small, isolated test bed to catch major problems as we moved from one version of host or router software to another.

How test cases were evaluated varied according to the test. In many cases, subjective evaluation of audio and video on the hosts served to identify problems in the end-to-end delivery. Loss and delay variation on the hosts were also measured. For most test cases, router interfaces were monitored for loss and utilization through the router's user interface commands.

A few RSVP functions were problematic across multiple versions of host or router software.

- Merging
- Reservation recovery following route change
- Guaranteed Service Model delay-specification tests

The RSVP merging test was ultimately completed successfully. The recovery of a reservation following a route change (induced by a simulated failure on the test bed) suffered from problems in both the host and router implementations. A race condition would occur when a PATH TEAR message propagated back to the host faster than reservation re-establishment occurred. The host RSVP implementation sometimes exposed this condition to the user as a lost reservation, or it failed to reestablish the reservation. Ultimately, the route-repair test was the only test case that was not completed successfully since it was among the last test cases run, and time had run out on RSVP functional testing. Our delay-specification tests suffered from a lack of appropriate test tools for measuring delay.

Some testing problems were not related to the RSVP implementation.

- DVMRP tunnels terminating in PIM networks
- Finding appropriate test tools that scale to large numbers of session and which can be used to measure end-to-end performance
- RSVP user interface specification

The third point, RSVP user interface specification, is discussed in the next section.

2.3 RSVP Tools

RSVP-enabled tools were needed for testing purposes as well as for actual use during Trials activities. Synthetic workload generators were used for testing routers and hosts, but most of the end-to-end testing relied on subjective evaluation of audio and video quality to determine how well multimedia services were being provided on the test bed. A total of seven RSVP test tools were used from three vendors or sources.

Tool	OS used	Source
RSVP ProShare® Presenter	Windows 95*	Intel
RSVP ProShare® Conferencing	Windows 95	Intel
RSVP Announce (sdr tool)	Windows 95	Intel
RSVP Test Kit*	Windows 95	Precept
PC-RSVP Test Tool	Windows 95	Intel
RSVP RTPTest Tool	Windows 95	Intel
RSVP-enabled mgen/drec	FreeBSD	NRL/Intel

Table 2: RSVP tools used for testing and by customers

Most of the test tools used were run under Windows 95 even though several of the tools from Table 2 worked under Windows NT* as well. Windows 95 was the lowest common denominator for a wide variety of multimedia application products that were evaluated on the test beds. Intel modified two ProShare® products and an RTP test tool to use RSVP signaling. In general it would take about one staff month to design, implement and test the RSVP-enabled prototype. There are a number of alternative designs that were followed across many of the RSVP prototype tools shown in Table 2. First, there are options for exposing or not exposing RSVP state to the user of the tool: RSVP state can be presented in a detailed manner, such as “PATH Received” or “RESV Sent”, it can be presented in a simplified manner such as reservation “On” or “Off”, or RSVP state can be hidden altogether. In fact, all three approaches were used in the set of RSVP tools shown in Table 2. One Intel tool and the Precept tool provide a level of detail suitable for debugging RSVP messages or API calls. Apart from debugging tools, however, the question of how to present RSVP state to the user is far from obvious.

A second issue is how many RSVP parameters should be presented to the user or how to parameterize an RSVP application product. The Intel tools that were RSVP versions of commercial products merely assumed a single sender or unicast service and Fixed Filter reservations only. It is not clear that the application tool will always be capable of automatically determining which filter parameter should be used. The tool also needs to determine the media rate, but many of the RSVP-enabled tools that were evaluated consistently failed to make the correct reservation for their RSVP flows at some points during operation. It was observed in the Trials that products from a variety of vendors failed to limit video rates to preset parameters. Those RSVP-enabled tools that merely attempted to predict a reasonable upper bound for the reservation often failed to do so correctly. Those tools that tried to adapt to changes in

transmission rate similarly failed to correctly adjust FlowSpecs or FilterSpecs to media (e.g., video) rate.

Different application tools used different RSVP Application Programming Interfaces (API) depending on when they were enabled to use RSVP. Since the start of the Trials effort, the API for network resource reservation has evolved on Windows* platforms. Braden and Hoffman specified the original RSVP API, called *RAPI*, for the reference implementation of RSVP done by the USC Information Sciences Institute. An abstraction of an RSVP programming interface is presented in the RSVP specification [2]. The Windows 95 applications that used RSVP in the Trials, however, originally used a specification from the Winsock2 RSVP Annex [3]. Currently, the Windows API for QoS is the Generic QoS (GQOS) Mapping that is available in Draft form [4]. GQOS is a generalized API that is designed for a variety of resource reservation protocols with RSVP being the first protocol having a GQOS mapping. Microsoft has announced RSVP and GQOS for the next release of Windows NT. As products become available that use these services, it will be interesting to see how the issues of the presentation of reservation state, reservation parameters, and the adjustment of reservation requests to match varying CODEC rates are handled.

In addition to application tools that sent and received audio and video, synthetic workload generators such as mgen from NRL were used for router and host testing. Mgen was modified to generate hundreds of constant-bit-rate flows to evaluate router-forwarding capacity. The Cisco router RSVP user interfaces were used to generate large numbers of reservations for measuring RSVP signalling capacity. These tests were not intended to serve as product evaluations and so are not reported in this paper.

3 Applications and Customers

The goal of the testing was to evaluate and understand the network technologies to support multimedia applications on corporate IP networks. Three classes of multimedia applications were evaluated through customer use. As shown in Table 3, over 2000 employees participated in over fifty events during the Trials. Each row of Table 3 shows the total number of people (#PARTICIPANTS) in all events (#EVENTS) by company and application type. The events consisted of training classes, network-based briefings and small workgroup collaboration sessions.

APPLICATION	COMPANY	#EVENTS	#PARTICIPANTS
Collaboration	Cisco/Intel/MCI	24	>150
Briefing	MCI	2	>1000
Briefing	Intel	12	>750
Training	Intel	14	123

Table 3: Trials Applications and Participants

The "Collaboration" application used IP-based videoconferencing for small workgroups. This application was used on a daily basis from desktop to desktop on two of the three network test beds, and it was used from room-to-room and room-to-desktop at several Intel sites. On the SMDS and Intel test beds, videoconferencing was used to coordinate testing and configuration. Two other groups at Intel used IP-based videoconferencing for their weekly meetings that involved cross-site groups. In general, desktop-to-desktop videoconferencing was used mostly to

coordinate testing. The testing team found this useful though there was not much demand for desktop-to-desktop videoconferencing on the IP network during the Customer Trials. Some employees experimented with use of videoconferencing on the test bed and an H.323 firewall was installed at Intel to make such use more secure. Use of videoconferencing in the Trials was essential for some training classes where the application-sharing feature of the ProShare® conferencing product was requested by the instructors.

"Briefings" are a remote presentation application that had the largest number of Trials participants. The Briefings application involved both rooms and desktop participants where use of IP multicast service was essential to the event. IP multicast is probably the best alternative when the goal is to involve large numbers of remote participants in a conference. In this, the Trials reproduced the experiences of the Mbone on the smaller scale of a single enterprise. MCI involved more than a thousand employee participants over two events. The participants were selected to be those employees most likely to be comfortable with the new technology: One event was the MCI Webmaster's Forum that ran IP multicast over more than 100 routers operated by MCI's network information services group, InfoLink. MCI used a variety of tools for distribution and presentation of audio, video and meeting presentations across the MCI corporate IP network and dial-up lines. The SMDS test bed was used to connect a site in England for one of the events. RSVP was not used for this event since a decision was made not to upgrade production-router software to experimentally run RSVP on MCI's production network.

Intel held many small remote presentations that had fewer than 200 participants. Prior to the Trials project, Intel Architecture Labs engineers organized large remote presentations of up to 2000 participants on Oregon campuses. The Intel Trials effort during 1996-97 was intended to link sites across Intel's domestic WAN and to install an infrastructure suitable for a future production service, at least among four sites in the Western U.S. As shown in Table 3, Intel ran twelve events with over 750 total participants (in fact, the count of events and participants on Intel's test bed exceed this number since events are still being held even though the formal evaluation procedures stopped in June, 1997). Only one campus was enabled for IP multicast service, it is the Jones Farm campus that has thousands of employees with desktop computers and a total of about 50,000 ports on the campus IP network. At the three other sites, service was to rooms and small numbers of desktop PC's. A few large events attracted in excess of 100 participants. For some events, however, there would be only one half-dozen or so employees in a room, though these people avoided travel or avoided missing a meeting by using the remote presentation service. Intel's experience is that presentations by corporate or technical leaders or presentations about compelling issues (such as changes to the stock plan) can attract hundreds of participants. Routine briefings and colloquia can attract dozens. Whether this level of participation will change as IP multicast becomes more widely deployed on more campuses remains to be seen.

STUDENT RESPONSE	PERCENTAGE
Positive	23%
Positive with Reservations	59%
Negative with Reservations	12%
Negative	6%

Table 4: Student Evaluation Breakdown by Percentage

The third application shown in Table 3, "Training," used both videoconferencing and remote presentation tools between rooms and desktop computers that were located in training rooms.

There is considerable interest in training at Intel and many other large corporations. U.S. corporations spend over \$50B U.S. annually on training, and this figure increased by 5% per year in the period between 1986 and 1995 [5]. Training costs can be considerable, and network-based training is being widely deployed in an effort to reduce the costs of this activity by eliminating student or instructor travel, to provide "just in time" training for equipment operation, and to integrate computer-based education into the corporate curriculum. Of the 123 students who participated in training classes on Intel's Trials network, most were employees who worked in "Fabs", Intel's semiconductor manufacturing plants. Most of the Fab classes were Occupational Health and Safety Administration (OSHA) safety classes.

An educational psychologist who conducted surveys and interviewed participants formally evaluated the Trials activities. As Table 4 indicates, 82% of the participants rated the class positively - as "Positive" or "Positive with Reservations." 12% of the participants evaluated a class as "negative with reservations." The reservations had to do with audio quality (loss of data packets on the network, fidelity of the audio CODEC, microphone/speaker ergonomics) or feature deficiencies in the training product that was used (such as the lack of a video back channel from student to instructor). Reliability of the training product was an issue in at least one class. Although improvements in products and service delivery should increase the overall positive rating, not all people will like or even be willing to use distance-learning services - about 6% of the participants rated the class as being a negative experience.

4 Results

RSVP was successfully tested on the SMDS test bed. No significant functional problems were found in the RSVP router or host implementations that were evaluated. Network testing was done on a part-time basis owing to staffing constraints and customer use. For this reason, the elapsed time spanned months. The presence of mature RSVP and IP-based multimedia products today should greatly reduce the time needed for testing in future efforts. Testing time would have been reduced if more extensive IP multicast testing had been done since IP multicast routing issues consumed much engineering time. As discussed above, it was necessary to perform most tests on the production network that supported the service. Automated test tools are necessary to efficiently test new services. It is desirable to completely run multi-site tests from a single location rather than rely on having test personnel at each site. Thus test tools need to be remotely manageable and automated. More time spent in acquiring the appropriate test tools prior to the start of testing is time well spent. Host tools that can automatically report missed deadlines when running real workloads were not used in Trials testing, but these are essential for efficient testing audio and video delivery. Synthetic workload generators that can transmit hundreds of flows at varying rates and hundreds of RSVP reservations are also needed prior to the commencement of RSVP testing.

RSVP was not used by customers in the Trials since RSVP was not supported on the ATM network routers at the time of customer use though it is today. In both the ATM and SMDS test beds, the problems of supporting RSVP and packet scheduling on the periphery of a subnetwork loomed large in the Trials experiences. Incremental deployment of RSVP is an important issue: Corporations that deploy multimedia services on their IP networks are likely to have high standards for presentation quality, but they are not likely to introduce QoS on corporate routers on some "Flag Day." Managing traffic on an over-provisioned WAN subnetwork proved to be difficult and prone to error on our test beds. Many of the problems may be solved by deploying a

subnetwork bandwidth management [6] system to manage reserved flows across the over-provisioned subnetwork though such a solution remains to be shown.

Despite the problems and complexities of engineering a backbone WAN to support briefings, training, and workgroup collaboration, practically all of the events were deemed to be successful by the participants. Customer satisfaction was evaluated and an 82% rating of “positive” or “positive with reservations” was achieved. Multimedia products were found that successfully supported each application. Product performance results are not contained in this paper since the goal of the effort was not to painstakingly evaluate products relative to one another but rather to apply suitable products for particular applications. Costs and benefit analyses were also done, but ROI information is specific to the business conditions and network infrastructure of particular companies and so is not presented in this paper. ROI estimates are not only specific to an enterprise, but are also subjective, especially when attempts are made to estimate savings in employee travel or the qualitative improvements to the ability of employees to do their jobs.

Corporate briefings, a remote presentation application similar to many Mbone conferences, had the most participants overall. This application is being introduced at a number of corporations that feature multicast business content on their corporate IP networks. Commercial products are available which are easy to install and use for 1:N presentations with limited back-channel capability. The Trials experience suggests that briefings can be offered on many corporate IP networks at low cost. Typically, a remote presentation will have three media flows for video (about 300 kbps for good-quality, real-time video), voice (up to 80 kbps) and foils of text and graphics (approximately 80 kbps). Many businesses are likely to find that briefings using IP networks is an effective means to disseminate information to employees at their desktops.

Corporate training is an interesting area for application of multimedia technology since so much money is spent on it just in the U.S. The fourteen classes held at Intel saved at least fourteen classes worth of instructor time by having at least two sites per class, and instructor travel was eliminated for most classes. Savings in student-employee time through network-based training are an even larger source of savings for some companies. There are tradeoffs in doing corporate training over IP networks versus ISDN, but IP-based solutions can generally run over ISDN though the converse is not true. Trials training classes involved fewer sites and a smaller number of participants relative to corporate briefings. Trials customers, however, had more stringent requirements for interactivity for training than for briefings. Some instructors wanted n-way video and audio so that all members of the class could see and hear the students as well as the instructor. The application tools that were used for briefings and which scaled up to hundreds or thousands of users did not in general scale down to feature n-way audio and video, floor control, and shared control of presentation materials. Although remote presentation tools were used for many classes, even room-to-room training sessions, these products were designed to work well for 1:N presentations and did not feature the floor control protocols or level of interactivity required by many instructors. Both room and desktop services were evaluated in the Trials, though the desktop-based training used PC's configured in classrooms rather than on employee's desks. All of the Trials training classes were held in rooms. Workplaces that feature cubicles are often considered to be unsuitable for distance learning. Although the training application may have the greatest potential ROI, the experience of the Customer Trials was that most of the time spent on acquiring, configuring, deploying and supporting application tools was spent on the training tools.

5 Conclusion

The Commercial Internet Multimedia Trials project described in this paper was one attempt to answer some of the questions regarding the readiness of the IP-based multimedia services and the applications they support. Similar efforts have been undertaken by public and private service providers such as BBN [7] and some Fortune 100 companies over the past year [8]. Other efforts such as the IP Multicast Initiative [9] and the Networked Multimedia Connection [10] are bringing applications and service providers together to focus business and engineering attention on the rollout of new IP services for multimedia and multicast applications. Whether done through a public service offering, privately by an individual corporation, through an "alliance," or by a consortium, these efforts are probably more useful to the extent that the results are publicly shared. The overall experience of the Trials was that there were a several good application tools that were indeed "product ready," and we achieved a high level of customer satisfaction over fifty events that involved a few thousand employee-customers.

Another objective of the Trials was to show that the costs of deployment and operation of multimedia services on corporate IP networks could be low. These costs include both the investments in new equipment (as low as a software change for relatively late-model routers and hosts) and network resource usage (new CODEC's such as H.263 video and G.723 audio offer significant reductions in bit rates). Even with use of a separate multimedia backbone network, the costs of the service were very low for remote presentation services. Access to corporate IP networks for remote presentation will probably be more closely controlled than access to the Mbone; content providers in most cases will need to be sanctioned by the corporation. For these reasons, it may be generally true that the use of remote presentations in even large businesses will not result in more than a few concurrent sessions though the standards of quality and reliability for those sessions probably must be much higher than that found on the Mbone. Such a service should be inexpensive to implement on the IP networks of many Fortune 500 companies today. How low the needed investment will be depends on the particular company, its infrastructure and its policies. How high the return will be will also vary by company, and it may be very hard to estimate with any precision. It has been the experience of the Commercial Internet Multimedia Trials, however, that technology readiness in terms of services and application tools is probably not an impediment to deployment.

6 Acknowledgements

Many people contributed to the effort described in this paper including Chedley Aouriri, Jim Archuleta, Seenu Banda, Fred Baker, Alan Batie, Richard Beckwith, Lisa Birch, Guy Blair, Jackson He, Russell Kool, Vicki Lee, Stacey O'Rourke, Henry Sinnreich, Lou Poehlitz, Carlos Tapang, Abel Weinrib, and Raj Yavatkar. The authors also thank Christian Maciocco, John Richardson and the anonymous reviewer for their comments on previous drafts of this paper.

7 References

[1] Eriksson, H., MBONE: the multicast backbone, Communications of the ACM, Vol. 37, No. 8 (Aug. 1994), Pages 54-60.

[2] Braden R., Zhang, L., Berson, S., Herzog S., Jamin, S., Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification, Internet Engineering Task Force RFC 2205, September, 1997.

[3] Windows Sockets 2 Application Programming Interface, Revision 2.2.0, Microsoft Corporation, May 10, 1996.

[4] Bernet, Y., Stewart, J., Yavatkar, R., Andersen, D., Tai, C., Winsock Generic QoS Mapping (DRAFT), Microsoft Corporation, September 11, 1997.

[5] The Economist, Learning organisations. Those who can, teach, October 28, 1995.

[6] Yavatkar, R., Baker, F., Hoffman, D., Bernet, Y., SBM (Subnet Bandwidth Manager): Protocol for RSVP-based Admission Control over IEEE 802-style networks, Work in Progress, <http://ds.internic.net/inte>