Bay Area Gigabit Testbed Network (BAGNet): Experience with a High-Speed, Metropolitan Area, IP over ATM Network[1]

William Johnston

Imaging and Distributed Computing Group
Ernest Orlando Lawrence Berkeley National Laboratory
University of California
Berkeley, CA, 94720

Abstract

For two years, fifteen organizations in the San Francisco Bay Area participated in an ATM testbed to develop and deploy the computer multimedia network infrastructure needed to support a diverse set of distributed applications in a large scale, metropolitan ATM network environment. We present some of the lessons learned in the testbed, and suggestions for structuring and operating a successful computer network testbed.

1.0 Introduction

BAGNet was an IP over ATM, metropolitan area testbed that operated in the San Francisco Bay Area (California) for two years starting in early 1994. The participants included the computer science and telecommunications groups from fifteen Bay Area organizations. The goal was to develop and deploy the infrastructure needed to support a diverse set of distributed applications in a large-scale, IP over ATM network environment. The organizations that were involved were: Apple Computer, DEC – Palo Alto Systems Research Center, Hewlett-Packard Laboratories, International Computer Science Institute, Lawrence Berkeley National Laboratory (LBNL), Lawrence Livermore National Laboratory (LLNL), NASA Ames Research Center, Pacific Bell – Broadband Development Group, Sandia National Laboratories, Silicon Graphics, Inc., SRI International, Stanford University, Sun Microsystems, Inc., University of California, Berkeley, and Xerox Palo Alto Research Center (PARC).

The testbed was an IP over ATM network consisting of a full mesh unicast ATM/PVC, link level structure supporting four end nodes at each of the fifteen sites, and a full mesh ATM point-to-multipoint (multicast) link structure for each of the 15 sites. The unicast mesh provided an ATM "best-effort" quality-of-service over a 155 Mb/sec SONET infrastructure between the (approximately) 60 connected systems. This infrastructure was used for an IP network supporting a variety of distributed applications – see, for example, [Johnston96a] and [Wiltzius]. The ATM point-to-multipoint mesh was used to provide IP multicast. This capability supported high quality multimedia teleseminars using the MBone tools: vic, vat, and wb. (See [MBone1] and [MBone2].)

This paper describes some of the lessons learned in BAGnet, together with experiences in MAGIC, and several other ATM testbeds, in a form intended to assist others in organizing successful ATM computer network testbeds. We also suggest (drawing on experiences in four such testbeds) some criteria for a successful testbed and the supporting factors. General descriptions for BAGNet and MAGIC are available ([BAGNet1], [MAGIC]), as is an overview of the evolution of the MAGIC testbed [Fuller].

---

1. This work is supported by the U. S. Dept. of Energy, Energy Research Division, Mathematical, Information, and Computational Sciences office (http://www.er.doe.gov/production/octr/mics), under contract DE-AC03-76SF00098 with the University of California, and by ARPA ITO (http://ftp.arpa.mil/ResearchAreas.html). Author's address: Mail Stop: 50B-2239. Tel: +1-510-486-5014, fax: +1-510-486-6363, wejohnston@lbl.gov, http://www-itg.lbl.gov/~johnston. This document is report LBNL-39037.

2.0 The BAGNet Environment

2.1 Internet Architecture and Asynchronous Transport Mode Networks

The architecture of the traditional Internet is that of collections of “hot potato” packet routers that communicate with each other regarding the probable best way to get packets to their destination. These routers forward IP packets over a heterogeneous collection of link interconnects, and typically over LANs between routers and end node hosts.

While there are several models for data transport using ATM networks, the principle focus of the BAGNet testbed (and most other testbeds) was to use ATM networks as a link-level infrastructure for IP (Internet Protocol) datagrams. Once this is accomplished, then all Internet services operate over the ATM network, and interoperate with the many other link-level networks that comprise the global Internet.

To first order, there are arguably two basic Internet transport services: unicast and multicast. In the former, packets are labeled with source and destination addresses of communicating end systems (hosts) and transported unreliably from a single source to a single destination. TCP (transmission control protocol) builds on IP to provide reliable data streams between processes (computer programs), and UDP (user datagram protocol) and RTP (realtime transport protocol – used mostly for audio and video) provide connectionless, unreliable delivery, user-level messaging. IP multicast is an mechanism whereby many hosts can receive the same datagram, and the key issue is to determine the minimum number of routers that have to replicate datagrams in order that all interested parties receive them. IP multicast is yet not supported by many of the commercial routers used in the Internet, so a multicast backbone (“Mbone”) overlays the larger IP network by having multicast-capable routers that are separated by incapable routers, communicate with each other via TCP. (That is, TCP streams that go through conventional routers are used as the interconnect between multicast routers.) See [MBone2].

Asynchronous Transfer Mode (as opposed to the synchronous operation of SONET, which is the typical physical-layer protocol for ATM) is a circuit-oriented, micro-packet (“cell”) protocol. ATM is circuit oriented in the sense that ATM cells follow the same path for the life-time of a connection. Circuits are statically configured (permanent virtual circuit – PVCs) or, like a telephone call, an elaborate signaling process precedes data transmission in order to identify the entire path that data carrying cells will follow (switched virtual circuits – SVCs). ATM defines a number of transport capabilities called adaptation layers (AALs), one of which – AAL-5 – was tailored for IP.

Since it is fairly clear that the telecommunications industry will be deploying SONET and ATM infrastructure for the bulk of future high-speed networks, the question is how to map Internet protocols to ATM. The importance of this mapping, rather than scrapping Internet protocols in favor of data transport directly over ATM is two fold: In the Internet there is a great deal of experience in designing and building general purpose distributed systems. These systems, and the semantics of the IP-based transport, have grown up together, and it would not be a trivial (or even ultimately useful) task to adapt distributed applications to the semantics of an ATM-based transport. Secondly, the global Internet will be a very heterogeneous link-level environment for a long time to come, and IP is very successful in hiding this heterogeneity from the transport and applications layer.

While it is perhaps too early to say that there is consensus on the numerous issues of integrating ATM networks into the Global Internet, there is now enough shared experience in both the telecommunications and Internet communities that various “middle grounds” are emerging. See [RFC1932] for a good discussion of the issues.

2.2 BAGNet Architecture

Pacific Bell's CalREN program awarded to each participant a grant covering the cost of two years of OC-3 (155 Mbits/sec), ATM service. The service was defined as best-effort virtual circuits to the other BAGNet sites, and was provided by single-mode SONET with up to a 15 db loss from the customer site to the Pacific Bell equipment.

2.2.1 The Physical-level

The network backbone consisted of several Pacific Bell operated, Central Office ATM switches connected via a SONET network. SONET is a topic of its own, but for the purposes of this discussion a few brief comments are in order.

SONET is the optical fiber implementation of the "synchronous digital hierarchy." SDH specifies a synchronous communication mechanism and a set of interrelated communication speeds. It provides framing conventions that allow many different "flows" of different bandwidths to be carried on a single physical circuit. These flows can migrate in and out of streams through the use of multiplexers and other SONET-specific equipment. SONET "networks" are relatively static, with the flows typically corresponding into physical end points of equipment. Reconfiguring a SONET network is something that typically takes minutes to hours. SONET is designed for high- speed, long haul communication, with OC-48 (≈ 2.5 gigabits/sec) being the typical bandwidth between SONET equipment operated by the telecommunication industry in the wide area. SONET streams are typically subdivided into several OC-12 (622 mbits/sec) and OC-3 (155 mbits/sec) flows operating between ATM (or other communications) equipment. SONET equipment is relatively expensive, and in BAGNet, SONET was used primary to interconnect the Pacific Bell central office ATM switches, and to get the user switch port connections (as opposed to the inter-switch trunking ports) to within 15 db fiber loss of the customer site. That is, SONET from the Pacific Bell ATM switches was used to reach the central office that is within about 10 km (15 db fiber loss) of the customer site. The customer had to provide equipment capable of driving 10 km of single-mode fiber from their site to a Pacific Bell SONET terminal.

The overall Pacific Bell, San Francisco Bay Area ATM network is used by organizations other than those identified as part of BAGNet (some of these were, and still are, partnered with BAGNet organizations, and some worked independently). During the BAGNet period there were about 30 sites connecting to the metropolitan ATM network at the 155 mbits/sec rate. (The two access rates for the Pacific Bell ATM network were OC-3 (155 mbits/sec) and DS-3 (45 mbits/sec).) There were also connections to the Monterey Bay area (80 miles south), with inter-LATA ATM service being provided by Sprint, and there was a similar Pacific Bell network in the Los Angles area. (Some of the other sites and testbeds are still operating.) See [Johnston95a] for illustrations of the configurations.

2.2.2 The Link-level

As mentioned, the Pacific Bell, Bay Area ATM network was PVC based and fully meshed.This means that there were a fixed set of virtual circuits between every communicating pair of end-nodes in the network. The reason for using PVCs instead of SVCs was that, at the time, the much more complex mechanisms needed to support SVCs had not been implemented by all of the ATM switch vendors, and for SVCs to work all of the switches have to talk to each other during the circuit setup phase. (Even today – mid 1996 – heterogeneous SVC networks are just being established.) The several disadvantages of PVC meshes are first that they grow as the square of the number of sites, and they are typically configured by hand (an error-prone process). The BAGNet mesh – about 1800 ($60^2/2$)

PVCs – was the largest that could be accommodated by the central office ATM switches used by Pacific Bell.

A virtual circuit is a path through a collection of ATM switches. The mechanism of providing this path is that there is a table in every switch with entries corresponding to every path through the switch. Note that neither PVC or SVC directly addresses routing. (That is, how do you determine by what physical path – sequence of switches – you get from source to destination?) PVCs are statically routed by “hand”. The routing issues for IP and SVC networks are turning out to be essentially the same, and the same (or merged) protocols may end up being used for both (separately or in combination). See [RFC1932].

2.2.3 The Network-level

To provide IP service BAGNet used the approaches given in [RFC1483] and [RFC1577]: ATM “Adaptation Layer - 5” (AAL-5) is used to carry IP packets. AAL-5 defines structure that organizes collections of cells to carry larger unites of data. AAL-5 has no multiplexing and no cell sequence numbers. The basic idea is to provide an efficient and reliable way to transport a protocol data unit (PDU). A PDU of 1 to 65535 bytes is structured in a cell sequence marked by a trailer. Cells are not sequence numbered within the PDU under the assumption that cells may be dropped, but not reordered by the ATM network. The entire PDU is protected by a 32 bit CRC.

IP packets with IEEE 802.2 Logical Link Control (LLC) headers are placed in the AAL-5 PDU. (See [RFC1483].) The use of LLC potentially allows identification of the protocol (i.e. IP is encapsulated via 802.2). Other protocols could be – but in BAGNet were not – transported using this mechanism.

Address resolution (RFC-1577) is done by ATM ARP (“address resolution protocol”) servers whose address tables are statically configured. (ATM is a “non-broadcast multiple access” medium and so Ethernet style ARP will not work.)

The current trend (e.g. in the MAGIC testbed) is toward an addressing and routing architecture where NHRP (“next hop resolution protocol”) is used to generate the ARP tables automatically within autonomous system / routing domains, and a routing protocol like BGP is used to provide NHRP servers with information for inter-domain connections. See [RFC1932] and [MAGIC].

3.0 Application use of IP Over ATM

Much of our interest and experience with ATM testbeds is in experimentation with approaches to problem that are enabled by a high bandwidth network infrastructure. Network performance and characteristics are obviously an important factor in this circumstance, and we have developed several performance monitoring strategies for characterizing network performance when they are used for high data rate applications.

3.1 Performance Issues Related to ATM

For high data-rate applications there are virtually no behavioral aspects of an ATM “network” that can be taken for granted, even in an end-to-end ATM network. By “network” we mean the end-to-end data path from the transport API through the host network protocol (TCP/IP) software, the host network adaptors and their device drivers, the many different kinds of ATM switches and physical link bandwidths, and then up through the corresponding software stack on the receiver. Further, the behavior of different elements at similar places in the network architecture can be quite different because they are implemented in different ways. The combination of these aspects can lead to complex and unpredictable network behavior for high performance applications. (See [Tierney].)

The next two sections provide some information and observations on these issues. We only address issues in the flow of data after the host-to-host connection is established. (That is, we do not address routing and signaling.)

3.2 ATM Network to Workstation Performance

The architecture and resulting environment for high performance applications is typically a “data movement intensive” environment. That is, in the normal course of dealing with large data-objects in a fully distributed environment, large amounts of data are constantly on the move. This circumstance generates many the R&D issues, and we discuss a few of them here.

The following synopsis describes the progress in the area of host ATM interface performance, and the use of multiple network interfaces to provide striped (parallel) data transport over the network (which is very important in high-speed distributed applications). Some of the progress noted below is due to the maturing of ATM network code stacks, and some is due to several years of interaction with the workstation vendors pointing out the problems. Figure 1 indicates the evolution of the performance of IP over ATM host adaptors over several years. The main point is that even though we still cannot get

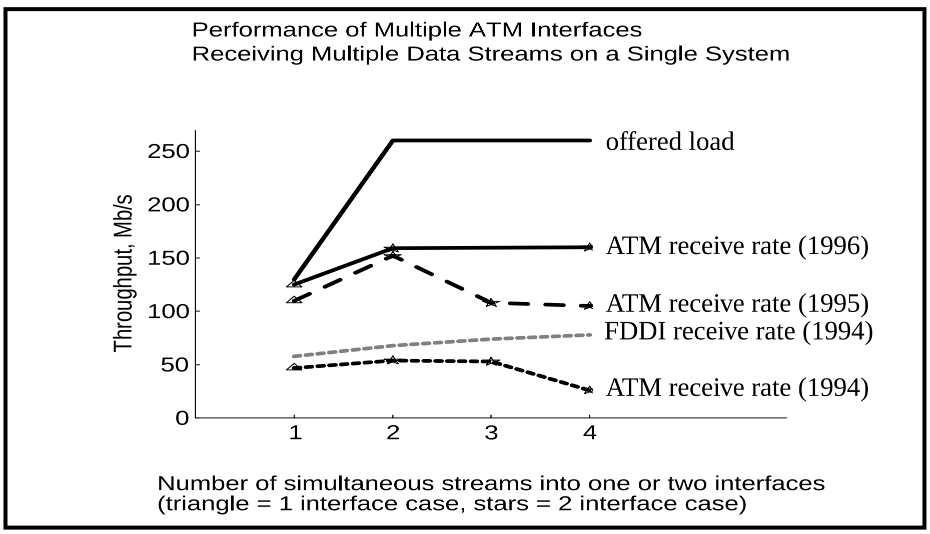

Figure 1

Evolution of the Performance of Workstation ATM Network Interfaces

linear speed-up by parallel use of multiple host interfaces, at least we see some increase (as opposed to the decrease seen two years ago!). In general, we can characterize the evolution of IP over ATM infrastructure supporting computer applications as follows.

- 1994
  - single interfaces were slow
  - multiple interfaces were no faster (poor independence of system data paths – nothing was multi-threaded)
  - switches dropped cells silently
- 1995
  - single interfaces become faster
  - multiple interface data path independence improved
  - switches still dropped cells silently

- 1996
  - single interfaces and systems are approaching wire speed
  - multiple interface performance is better still, but not ideal (so some striping is possible)
  - switches now report cell loss

3.3 ATM Network Performance

3.3.1 Uncongested Behavior

In order to determine suitability of the Pacific Bell Metropolitan Area ATM for a high data-rate application (video-angiography transmission that required sending 4 GBytes in 20 minutes with a duty cycle of about 50% – see [Johnston96b]), we ran a series of throughput tests. The experiments consisted of sending 25 Mbits/sec, with a duty cycle of 40 minutes on, 20 minutes off, 24 hours a day, and monitoring the end-to-end throughput. These tests were conducted continuously over several weeks.

Although this network was basically uncongested, several of the experiments were timed to correspond to some high resolution video and audio multicast sessions that generated 5-10 Mbits/sec of “continuous media” (steady state) traffic in the network, in addition to whatever other activity might have been occurring in the network.

The experiments were done by putting the source host ATM adapter into “continuous bit rate” mode, but sending data via IP over AAL-5. This ensured that we transmitted a “true” 25 Mbits/sec, and did not inject high-bandwidth “bursts” of data into the network (as would have been the case if we had rate-limited the TCP writes to an “average” of 25 Mbits/sec).

The results were that, to the level of about 1%, there was no variation in our ability to send 25 Mbits/sec of data for any of the several week-long periods during which the throughput experiments were run.

3.3.2 Congested Behavior

The impact of congested ATM networks on TCP transport (which is used for all of the data transfers described here) can be very complex. Today’s ATM networks and components are still relatively immature and with every new generation of switch software and hardware the congested behavior can change. The most severe impacts occur in circumstances when the characteristics of the various ATM links and network components conspire to defeat TCP’s otherwise well-tuned congestion control mechanisms. TCP has a suite of congestion detection and response algorithms (see [Stevens]) that have evolved over the past 10-15 years in the Internet environment, and in that environment they work very well. (The Internet operates under conditions of very high load and high potential congestion, but congestion is generally avoided and the links usually operate at high utilization efficiency.) The problem in ATM networks is that the TCP control algorithms are based on certain assumptions about the relationships among parameters like the packet switch buffer sizes, link minimum transmission unit sizes, link bandwidth, etc. It is very easy for ATM networks to violate these assumptions, and when this happens both the throughput and the link utilization efficiency can be very low in the presence of congestion. (In current implementations of ATM networks, this congestion is frequently caused by small output port buffers in ATM switches. Fortunately, most new switches are being designed with large output port buffers.) For a detailed analysis and case study of this situation, the reader is referred to “Performance Analysis in High-Speed Wide-Area ATM Networks: Top-to-Bottom End-to-End Monitoring” ([Tierney]).

4.0 Criteria and Recommendations for Successful Testbeds

This section was motivated by the author's participation in the "International Forum for APII Testbed" held in Seoul, Korea in June, 1996, during which the first steps were taken for the planning of a set of interconnected network "testbeds" in the Pacific-rim countries. (See, for example, [APII].) This section is written with the intent of providing some small contribution to the success of the APII goals.

4.1 Characterizing Successful, High-Speed ATM Computer Network Testbeds

In computer network testbeds there are essentially two general technology areas involved, each with its own criteria and issues for success.

Networks are enabling infrastructure, and are successful to the extent that they enable applications and access. It is hard to generate "excitement" for networks alone – no matter how successful they are – outside of the technology community.

Successful testbed *applications* provide new and useful middle- and high-level services – especially capabilities that enable new approaches to problems.

Applications generate "excitement" when they provide for visualizing, connecting, and analyzing data in new ways that provide new insights into solving problems. That is, they enable problem solutions that: are more useful; are simpler or more cost-effective; provide new insights into, or completely new uses for, information; provide new ways to express creativity.

4.2 Success Factors for Testbeds

The network aspects of the testbeds are successful when they:

- Evolve to provide simultaneously more flexible and robust communication
- Provide higher bandwidth IP delivered with higher connectivity
- Produce sustained collaboration between the telecommunications sector (carriers and equipment providers), the Internet service providers (the people and organizations supporting IP networks) and the application developers
- Connect to other testbeds
- Turn into carrier-based production computer networks.

Testbed applications are successful when they:

- Stress the network infrastructure (to identify problems and shortcomings)
- Provide for detailed monitoring of all levels of the data transport (to accurately characterize the network)
- Are useful in principle:
  - potentially useful in the future (if everything works as planned)
  - provide useful "public relations"
- Are useful in fact:
  - address a "real" problem, the solution of which will result in "rapid" migration to a prototype production service

- Promote software infrastructure to support future applications
  - e.g., portable and generally useful information services that will make it easier to build other useful applications
- Promote new ways of looking at problems
  - e.g. in the way that the World Wide Web provided a new way of publishing information

4.3 Some Specific Elements of Successful Testbeds

A computer network testbed consists of three elements: a telecommunications infrastructure, an IP infrastructure, and applications. These components are typically provided by three different groups, and in the "testbed" environment (typically where new telecommunications services – e.g. ATM – are being used) it is very important that the three groups work together as a routine part of the functioning of the testbed. The mapping of IP services onto ATM services presents many issues that are still only partially resolved, and that require close cooperation of the ATM and IP communities. Developers of applications, especially those requiring high bandwidth, have to work with the IP and ATM providers to identify and resolve network-level issues that impede the operation of the applications. IP computer networks are very different from telecommunications networks, and it is essential that people familiar with IP technology work with the telecommunications providers.

The telecommunications infrastructure should initially be provided by R&D groups from the carrier organizations rather than their operations groups. After success operation of the testbed has been demonstrated, and the issues identified and resolved, then the carrier organizations can, and should, train their operations staff and transfer the testbed operations. Operations groups do not typically have the flexibility of view and procedure to be successful in dealing with an experimental service. One criteria for the success of most testbeds is that they eventually lead to new operational services, but they cannot start out that way.

"Real" applications that make use of unique features of the network (e.g. bandwidth) are both more impressive and more likely to "stress" the network in ways that will identify significant issues. The TerraVision application in MAGIC is an example of such an application, as is the video-angiography application in the Pacific Bell ATM testbed. (See [TerraVision] and [Kaiser].)

Testbeds should typically start as a single administrative domain (e.g. effectively a single IP network or subnet that is administered by a single "group" – though the group is typically the result of consensus among the participants rather than truly a single group). This approach facilitates the cooperation needed to get all of the participants educated in the theory and operation of the testbed technology.

Initially, interoperation of testbeds with other testbeds or with production IP nets can be done relatively easily by using IP routing without end-to-end ATM connectivity. The longer term goal for a high performance infrastructure is to interoperate ATM "clouds" (multiple IP networks) at the ATM level, which is considerably more difficult. (See, e.g., [RFC1932].)

If "appropriate or acceptable use" policies are needed, keep them very simple. E.g.:

- One end of a connection must reside in the testbed, or
- Non-commercial transit traffic is allowed, or
- The testbed may not replace a commercial connection, etc.

The more open and more used the testbed is, the more will be learned.

One year is probably the minimum useful duration of a testbed. Two years will be required if significant applications are to be developed, deployed, and experimented with. Two years will also typically be required to move from a single IP network (first year) to an interconnected set of IP networks (second year) assuming that those involved do not have a lot of experience with IP over ATM networks.

Involvement of academic computer science departments is very useful. Students are a fertile source of ideas for uses of the testbed. (Almost no use of the testbed should be discouraged, whether "serious" or not. The more different kinds of applications and traffic the better.)

An "Information Infrastructure" testbed is an IP (Internet) testbed, and the applications will be IP based. This should be the focus of the testbed.

If high performance applications (i.e. those that can generate data rates comparable to the full link capacity) are part of the testbed environment, then application-level traffic shaping (rate limiting) will be required for those applications. At least initially, "ordinary" (low bandwidth) applications should not be burdened with traffic shaping. Host network adaptors (interface devices) should do traffic shaping transparently to the application.

Testbeds should obtain and deploy the MBone teleconferencing tools vic, vat, wb, and sd ([MBone1]). At least one, and preferably more, audio and video capable workstations should run the MBone tools at each testbed site. It should be a goal to hold testbed technical meetings using these tools (that is, do the meetings as network based teleconferences). Installing and operating these multimedia tools will teach the testbed community a great deal, will provide a very useful way to foster cooperation among geographically dispersed sites. Further, IP-based multimedia computer network conferencing is an increasingly important application that usually works well in testbed environments.

ATM point-to-multipoint capability should be configured to support IP multicast in the testbed. (Otherwise, set up Mbone tunnels between the testbed sites.) Care should be taken to understand the interface between the testbed Mbone environment (especially if the testbed is high bandwidth) and the local or external Mbone environment (which is liable to support only low bandwidth Mbone traffic) so that testbed use of Mbone does cause problems for production networks.

A general testbed management scheme that is know to, and accepted by, all testbed participants. Guidelines and advertised scheduling should be provided for switch reconfiguration, special (e.g. disruptive) experiments, important demonstration, etc. The guidelines should provide for prior notification, scope and limits of changes, etc. The intent is not to inhibit experiments, but rather to make sure that some experiments don't interfere with others, that demonstrations are successful, etc.

Monitoring and performance tools that are available and used by all of the sites are very important. The minimum set includes: throttled ttcp (ttcp modified so that it can transmit at various specified rates), traceroute (the basic IP network route monitoring tool), and a "ping matrix". (A ping matrix is a continuously updated connectivity matrix that is periodically posted to testbed Web servers, and describes the accessibility of connections to all of the testbed end nodes. See, e.g. the "technical information" section of the BAGNet Web page [BAGNet1] that describes the "bagping" tools.

Well designed and implemented ATM switches should be selected and used. The following characteristics are important:

- Accurate reporting of output port buffer statistics (especially per VC cell dropping) to applications (e.g. via SNMP)

- No arbitrary restrictions (e.g. on the number ranges of PVCs, forced association of CBR with PVCs, etc.)
- Remote configuration and operation
- Point-to-multipoint capabilities

Similarly, host adaptors should provide:

- Support for multiple IP nets on a single adaptor
- Class D IP addressing (IP multicast)
- Setable, per VC, CBR rates is a useful tool for performance experiments.

The testbed participants should be organized so that more experienced groups can assist less experienced groups. Provide sample switch configuration setup information, etc.

Every site should maintain a Web page for publishing both public and private testbed information. All aspects of the configuration of the testbed should be readily available to all participants. A central, or testbed, home page should link to all of these site pages. (See, for example, [BAGNet1].)

The personnel responsible for operating the telecommunications and IP networks should all have a good understanding of the material and background described in [RFC1932].

5.0 Conclusions

There is a great deal of work to be done in order to make ATM a well integrated transport mechanism for the Internet protocols. In the process of doing this work both IP and ATM will be improved technically (and already have been), and more knowledge will be built up about the higher level question of universal transport models. (An important contribution to this higher level model has been made in the book "Realizing the Information Future" [NAS].)

6.0 Acknowledgments

The work in BAGnet was done by many people from the fifteen constituent organizations. The BAGNet executive committee was co-chaired by William Johnston, LBNL (wejohnston@lbl.gov), Marjory Johnson, NASA Ames (mjj@riacs.edu), and Dan Swinehart, Xerox PARC (swinehart@parc.xerox.com). The IP-over-ATM group was initially chaired by Mark Laubach (then at HP Labs) and then by Berry Kercheval, Xerox PARC (kerch@parc.xerox.com). Many important contributions were made by other participants as well, especially performance analysis work by Lance Berc (DEC-SRC), Dave Wiltzius (LLNL) and Helen Chen (Sandia). Rick Hronicek of Pacific Bell developed the business strategy that enabled Pacific Bell to support the CalREN program.

7.0 Notes and References

APII Choi, B. "APEC Ministerial [statement] on the APII and Seoul Declaration"

(See http://gopher.nectec.or.th/bureaux/nitc/bilchoi.html).

"The First APEC Ministerial Meeting on Telecommunications and Information Industry was held in Seoul on 29 and 30 May 1995, following the APEC Economic Leaders' initiative at Bogor. Ministers recognized the critical role of telecommunications and information infrastructures in achieving the goal of liberalizing trade and investment and shared the view that they should take a full responsibility in advancing an Asia-Pacific Information Infrastructure (APII) initiative and establish the APII. The APII was considered to be an integral part of and can contribute to the development of the Global Information Infrastructure (GII).

The APII represents more than the simple construction of a physical information superhighway in the APEC region. APII is a cooperative framework of the APEC member economies, the primary goal of which is to establish a seamless APEC-wide information infrastructure, by interconnecting the individual national information infrastructures of the APEC members. APEC is designed to connect the APEC as one, in order to dissolve information gap among member economies and to establish a simultaneous-information living zone within the APEC region."

BAGNet1 "Bay Area Gigabit Testbed (BAGNet)" See http://www-itg.lbl.gov/BAGNet.html

Fuller Fuller, B. and I. Richer, "The MAGIC Project: From Vision to Reality". IEEE Network, May/June 1996, vol. 10, no. 3.

Johnston96a Johnston, W., G. Jin, G. Hoo, C. Larsen, "Distributed Environments for Large Data-Objects: Broadband Networks and a New View of High Performance, Large-Scale Storage-Based Applications", Internetworking '96. (Also see http://www-itg.lbl.gov/~johnston)

Johnston96b Johnston, W., G. Jin, G. Hoo, C. Larsen, J. Lee, B. Tierney, M. Thompson, "Distributed Environments for Large Data-Objects: The Use of Public ATM Networks for Health Care Imaging Information System". Available at http://www-itg.lbl.gov/~johnston .

Johnston95a Johnston, W., "BAGNet: A High-Speed, Metropolitan Area, IP over ATM Network Testbed", IEEE Compcon'95, March, 1995, San Francisco, CA. Also see http://www-itg.lbl.gov/BAGNet.hm.pg.docs/CompCon.95.3.html .

Kaiser Kaiser - LBNL - Philips CalREN project. See http://www-itg.lbl.gov/Kaiser/LKP/homepage.html .

MBone1 The vic, vat, wb, and sd MBone tools are developed in the LBNL Network Research Group and are described at http://ee.lbl.gov. .

MBone2 M. Macedonia and D. Brutzman, "MBone Provides Audio and Video Across the Internet", IEEE COMPUTER magazine, pp. 30-36, April 1994. (Available from ftp://taurus.cs.nps.navy.mil/pub/mbmg/mbone.html .)

MAGIC MAGIC (Multidimensional Applications and Gigabit Internetwork Consortium) is a gigabit network testbed that was established in June 1992 by the U. S. Government's Advanced Research Projects Agency (ARPA). The testbed is a collaboration between LBNL, Minnesota Supercomputer Center, SRI, Univ. of Kansas, Lawrence, KS, USGS - EROS Data Center, CNRI, Sprint, U. S. West, Southwest Bell, and Splitrock Telecom. More information about MAGIC may be found on the WWW home page at: http://www.magic.ne .

NAS "Realizing the Information Future", National Academy Press, 2101 Constitution Avenue, NW, Box 285, Washington, DC 20055, telephone: 800-624-6242. Also available from http://www.nas.edu:70/1/nap/online/rtif/, or from ftp://ftp.nas.edu/pub/reports/realizing_the_information_future.

RFC1483 J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", Internet Engineering Task Force, Request for Comments (RFC) 1483, July, 1993. (Available from http://ds.internic.net/rfc/rfc1483.txt .)

RFC1577 M. Laubach, "Classical IP and ARP over ATM", Internet Engineering Task Force, Request for Comments (RFC) 1577, Jan., 1994. (Available from http://ds.internic.net/rfc/rfc1577.txt .)

RFC1932 Cole, R., D. Shur, C. Villamizar, "IP over ATM: A Framework Document", Internet Engineering Task Force, Request for Comments (RFC) 1932, April, 1996. (Available from http://ds.internic.net/rfc/rfc1932.txt .)

> "The discussions of the IP over ATM working group over the last several years have produced a diverse set of proposals, some of which are no longer under active consideration. A categorization is provided for the purpose of focusing discussion on the various proposals for IP over ATM deemed of primary interest by the IP over ATM working group. The intent of this framework is to help clarify the differences between proposals and identify common features in order to promote convergence to a smaller and more mutually compatible set of standards."

Stevens Stevens, W. "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms". IETF Draft. See ftp://ietf.cnri.reston.va.us/internet-drafts/draft-stevens-tcpca-spec-01.txt .

TerraVision A terrain visualization application that uses the Distributed Parallel Storage System (http://www-itg.lbl.gov/DPSS) operating in MAGIC to supply high-speed data streams for realtime, "virtual reality" type navigation of real landscape. (See http://www.ai.sri.com/~magic/terravision.html .)

Tierney Tierney, B., W. Johnston, G. Hoo, J. Lee, "Performance Analysis in High-Speed Wide-Area ATM Networks: Top-to-Bottom End-to-End Monitoring", IEEE Network, May/June 1996, vol. 10, no. 3. (Also available from http://www-itg.lbl.gov/DPSS/papers.html .)

Wiltzius Wiltzius, D., L. Berc, and S. Devadhar, "BAGNet: Experiences with an ATM metropolitan-area network", ConneXions, Volume 10, No. 3, March 1996. Also see http://www.llnl.gov/bagnet/article.html .