ANL-97/1

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

ARCHITECTURE OF THE MULTI-MODAL ORGANIZATIONAL RESEARCH
AND PRODUCTION HETEROGENEOUS NETWORK (MORPHnet)

by

Robert J. Aiken, Richard A. Carlson, Ian T. Foster,
Timothy C. Kuhfuss, Rick L. Stevens, and Linda Winkler

ECT and MCS Divisions

January 1997

Argonne National Laboratory, with facilities in the states of Illinois and Idaho, is owned by the United States government, and operated by The University of Chicago under the provisions of a contract with the Department of Energy.

This technical report is a product of Argonne's Electronics and Computing Technologies and Mathematics and Computer Science Divisions. For information on the divisions' scientific and engineering activities, contact:

Director, Electronics and Computing Technologies Division
Argonne National Laboratory
Argonne, Illinois 60439-4815
Telephone (630) 252-7586

DISCLAIMER

CONTENTS

i

Abstract

The research and education (R&E) community requires persistent and scaleable network infrastructure to concurrently support production and research applications as well as network research. In the past, the R&E community has relied on supporting parallel network and end-node infrastructures, which can be very expensive and inefficient for network service managers and application programmers. The grand challenge in networking is to provide support for multiple, concurrent, multi-layer views of the network for the applications and the network researchers, and to satisfy the sometimes conflicting requirements of both while ensuring one type of traffic does not adversely affect the other. Internet and telecommunications service providers will also benefit from a multi-modal infrastructure, which can provide smoother transitions to new technologies and allow for testing of these technologies with real user traffic while they are still in the pre-production mode. Our proposed approach requires the use of as much of the same network and end system infrastructure as possible to reduce the costs needed to support both classes of activities (i.e., production and research). An initial step is to define multiple layers of production services (i.e., at the physical, network media, network bearer, middle, and application layers) that can be made accessible for concurrent use by the network researcher, manager, or application programmer. Breaking the infrastructure into segments and objects (e.g., routers, switches, multiplexors, circuits, paths, etc.) gives us the capability to dynamically construct and configure the virtual active networks to address these requirements. These capabilities must be supported at the campus, regional, and wide-area network levels to allow for collaboration by geographically dispersed groups. The Multi-Modal Organizational Research and Production Heterogeneous Network (MORPHnet) described in this report is an initial architecture and framework designed to identify and support the capabilities needed for the proposed combined infrastructure and to address related research issues.

1.0   Introduction

The research and education (R&E) community has a continuing need for persistent and scaleable network infrastructure supporting production and research applications as well as network research. This infrastructure is essential if researchers are to advance the state of the art both in advanced applications (for which reliable "production" network capabilities are required) and in the networking technologies that will provide the infrastructure of the future (for which crashable "research" network capabilities are required). The continually shortening cycle associated with the evolution of network research to production status only fuels the demand for advanced production networking capabilities and further strains the ability to provide it. Historically, the very different requirements of production and research have led to the use of distinct physical infrastructures for these two purposes. Yet, as the demand for increased bandwidth and capabilities continues to increase, the R&E community will have difficulty paying the high costs associated with acquiring and supporting parallel networks. Hence, we propose a new approach that will allow the use of the same physical infrastructure for both research and development purposes. As we explain in this report, this new approach poses significant challenges that will require a major research effort to overcome, but promises substantial benefits in terms of cost savings and enhanced research and production capabilities. In fact, we argue that the economics of network infrastructure associated with this approach are essential if the R&E community is to continue large-scale networking.

The need for integrated production and research infrastructure arises because, while network technologies, bandwidth, and capabilities continue to rapidly improve, enhancements to the resolution and scale of existing multimedia, collaboration, and database applications (and entirely new applications) are increasing demand at an equal or greater rate. So we can expect to see competition for scarce network resources for the foreseeable future. The R&E community cannot financially afford to support both a high-speed production and an extremely high-speed experimental network infrastructure. Neither can it afford to conduct network research at the expense of the scientific application researcher, or favor a plan that stagnates the network research required to meet the constantly increasing applications requirements by funding only production networks. Internet service providers (ISPs)[1] face similar problems: they can ill afford idle bandwidth, even for short periods, and therefore must seek new and innovative methods to utilize the infrastructure. A successful implementation of an adaptive

---

[1] ISPs in the United States include the inter-exchange carriers (IXCs), Regional Bell operating companies (RBOCs), cable companies, alternate access providers, commercial and private providers, and any other entity that provides telecommunications and Internet services to its constituency on a wide-area basis. ISPs in other parts of the world include similar service providers as well as national PT&Ts.

multi-modal network infrastructure and architecture will not only address the requirement for concurrent production and experimental infrastructure, but also holds promise for quick deployment of research and development (R&D) infrastructure to address national crises[i].

These considerations lead us to conclude that the grand challenge in networking is to implement and concurrently support both advanced production network services (e.g., vBNS[ii], ESnet[iii]), which applications can use with little risk, and a persistent experimental service (e.g., Dartnet, CAIRN[iv]) over as much of the same infrastructure as possible. In building such a shared infrastructure, we must endeavor to ensure that R&D network traffic and experiments do not adversely affect production traffic (and vice versa). This sharing of infrastructure can occur at numerous layers in the network, including the hardware, media, network bearer, transport, and application layers. The efficient sharing of resources will also occur on and within different network scopes, including the local (e.g., Campus), regional (e.g., Gigapop, MREN), and wide area (e.g., vBNS, ESnet, CAIRN) levels.

In addition to increasing networking bandwidth and capabilities, we must become smarter and more efficient users of network technologies because the demand for network capabilities always exceeds the available resource or the user's ability to pay for it. To overcome the physical limitations of traditional supercomputers, we adopted the use of massively parallel machines. Similarly, we need to become more innovative with router, switch, and overall network architecture design to take advantage of parallelism in switches, multiplexors, and routers. Adaptive temporal use and reuse of segmented network infrastructure must also be explored. Some router and Asynchronous Transfer Mode (ATM) switch vendors are already experimenting with such models, as evidenced by dual fabric switches. Active network technologies, as well as quality of service (QoS) support, can also support concurrent virtual networks with radically different technical requirements (e.g., production and R&D networks) and dynamic policies.

The benefits claimed for multi-modal network infrastructure in the R&E community also apply to telecommunication and Internet service providers, who must support concurrent virtual infrastructure for both production and experimental purposes, as well as multiple policy-based virtual networks on the same infrastructure. The benefits are especially applicable if these providers wish to make more efficient use of network resources in addition to being able to strain and test new network capabilities and features in the experimental mode using real applications; even if only on a temporary basis. Telecommunications service providers are currently seeking new and innovative ways to make use of untapped and underutilized infrastructure in the last mile (e.g., local loop) as well as in their own clouds and switching fabrics. ADSL, in fact all nDSL technologies, as well as ATM are perfect examples of these attempts. An adaptive, active infrastructure will greatly enhance the ability of these providers to tap underutilized bandwidth

by allowing them to dedicate network resources on a finer granularity in both time and capability. It is important to note that, although the adaptive, multi-modal network infrastructure that we propose will support the R&E community by separating production and experimental traffic, this model can easily be adapted to support any number of two or more virtual networks with heterogeneous and sometimes conflicting requirements and policies. For example, these capabilities can be used to separate traffic based on security, business, or acceptable use policies.

2.0   Benefits and Risks

Concurrent support for production and experimental network traffic will benefit the research community by providing more convenient access to large-scale testbeds. While small testbeds and localized pilots are useful for laboratory testing and exploration, their small scale does not normally strain and test new network protocols, tools, and architectures in a manner consistent with the demands of large numbers of users or advanced applications. Production applications, as well as a large number of participating end nodes, are required to thoroughly test new protocols and infrastructures. For example, the experimental R&D Dartnet network was used to develop and test new network protocols (e.g., Multicast IP and RTP) with a small number of nodes and participating researchers. Afterwards, the researchers sought out larger-scale networks (e.g., NSFNET and ESnet) to demonstrate and validate these protocols on a larger scale. Modeling and simulation may be of some use in analyzing and testing new protocols and architectures as long as they are not strictly based on Poisson models. Paxton and Floyd[v] have demonstrated that Poisson models, commonly used to design regular telephony services, do not reflect or represent data network traffic accurately. Therefore it is imperative that networking models and simulations be validated via wide-scale implementations and experiments using real user applications.

Concurrent support for both production and experimental network traffic will also have benefits outside the R&E community. Telecommunications and Internet service providers can use a multi-modal network infrastructure to provide "production-level" services concurrently with experimental or evolutionary network services. This will satisfy their requirements for incremental upgrades as well as customer requirements for both production and R&D facilities, large-scale stress testing of targeted infrastructure, and the introduction of new technologies and services as they evolve. Businesses can use a dual-mode environment to run their production applications while simultaneously experimenting with and evolving their use of new network infrastructures and capabilities. The Internet and, generally speaking, most enterprise networks are haunted by the demands and spirits of networks past (e.g., Decnet, SNA, and other proprietary networks), networks present (e.g., IPv4), and networks of the future (e.g., IPv6[vi]). The multi-modal network model provides us with virtual networks, concurrently supported at

4

various layers, that help us cope with this cyclic development and deployment of networks, systems, and applications.

The phased deployment model for new technologies is still valid for initial experimentation (i.e., performing fairly risky experiments such as new protocols that must first be tested in constrained environments). This model subsequently requires that the scope of the experiment be expanded to fully test these new capabilities. The challenge that lies before us is determining how to use as much of the same infrastructure as possible for concurrent and efficient use by both R&D and production traffic after the initial constrained testing is complete; this challenge becomes greater when we seek to stress test new protocols and architectures and benchmark their capabilities under real traffic. Not only is this multi-modal use and support of networks required for supporting the R&E community's combined R&D and production infrastructure requirements, but it is also useful for (1) introducing incremental upgrades (version upgrades or enhancements) to switches and routers in deployed infrastructure, and (2) providing a transition path for applications eager to exploit new network capabilities; e.g., quality of service (QoS) signaling from an application layer. This multi-modal approach does not necessarily invalidate the use of separate network infrastructures, such as separate switches or links, when the concurrent shared use of some or all of the infrastructure cannot be safely achieved.

Some risk is associated with all new technologies, even "pre-production" services offered by ESnet and vBNS, for example. Users and applications need to accept this fact and plan accordingly. One method for dealing with this issue is to perform a risk analysis of the proposed architecture and identify the portions or layers of the infrastructure that lend themselves to shared use. The "comfort levels" associated with this sharing will most likely vary depending on institutional culture and financial factors. However, wise use of adaptive, multi-modal infrastructure is necessary if we are to further enhance our ability to provide for advanced network research and production networks in the face of dwindling financial resources, as well as for more efficient use of infrastructure by the telecommunications and Internet service providers.

3.0 Virtual Production Network Services (VPNS)

A shared infrastructure can use the concept of a variable "bar" of production-level service to facilitate both the smooth introduction of new capabilities and the concurrent support of production and experimental activities. This concept also supports on-demand experimental use and manipulation of network infrastructure, bandwidth, and quality of service. The bar is virtual in that it can be temporal (i.e., exist for short, medium, or long periods of time) or spatial (exist at various levels of network services at the same time), while concurrently providing for multiple levels of production and R&D-level services depending on the requirements and perspectives of the applications and the network R&D experiments.

3.1 VPNS Bar

One issue in providing for both production and R&D experimental network services (the former supports R&D applications) is the definition adopted for the "production layer." A desirable environment would allow for a certain amount of concurrent elasticity where the production layer is perceived on a per application or virtual network basis. For example, when using this approach to support Asynchronous Transfer Mode[vii] (ATM) experimentation over a shared production hardware media, we might see a production ATM service composed of the ATM switch and local loop for the computer scientist experimenting with a network bearer service such as IPv6. Application scientists (e.g., physicists), though, view the IP layer and below as the production layer as they experiment with RSVP[viii] or reliable multicast for their message passing interface (MPI[ix])-based application. Each of these models has been provided separately in the past; i.e., a dedicated network for each scenario, with the possible exception of tunneling, which we will address later. We believe that the challenge is to provide concurrent support of these virtual production networks, as viewed by the applications and network researchers, on the same infrastructure. Each layer would provide the opportunity and concurrent support for network research and production network services at the next layer up. Each layer depends on the production bar of the services below it.

3.2 Hardware Layer

The first level of providing a "production bar" is the hardware level. We can multiplex both production and network R&D traffic on the same hardware by implementing a hardware multiplexing scheme such as Wave Division Multiplexing (WDM) or Sonet block multiplexing. A portion of the service or circuit (i.e., local loop Sonet, or WDM colors) could be physically split off to a set of production switches; the other portion(s) could be physically split out to yet another distinct set of R&D switches. This model allows for the sharing of a local loop while keeping the production and R&D traffic physically separate on the local loop and in the switches. Whether one multiplexes the two types of traffic over the same infrastructure on the local campus or in the carrier cloud (i.e., on either end of the multiplexed local loop) is determined by the entities in control of those infrastructures and any agreements they have come to with the end user. The carrier may indeed carry both the production and R&E traffic over the same set of switches and links or it may provide separate sets inside the cloud so as to separate the two types of traffic inside their cloud. Either solution provides the end user with the view of one access and local loop to the cloud to support both types of traffic.

The hardware layer can be further exploited if it is composed of distinct objects (switches, links, routers, multiplexors) that can be assembled by an application or network manager on either a real-time basis (i.e., milliseconds to seconds) or on a scheduled basis within hours or days in advance of its anticipated use. For example, an OC-12 pipe could be provided by using four OC-3 links and associated multiplexors and switches. The initial allocation can have two OC-3 links dedicated to production use and two OC-3 links used specifically for network research. If the network researchers are not using their portion of the network (i.e., their OC3s) at any given time, it makes sense to allocate those resources to the production traffic. This assumes that the network segments and components in question can easily transition to production-quality status and back to experimental status at the conclusion of use. Conversely, if the network researcher could use three OC3s for a short- term test of new protocols, and the production traffic is not using its share of the infrastructure, the experimental network project could temporarily make use of a specified amount of the production infrastructure for a short time and then restore it to production status after the experiment is completed. The production portion of the infrastructure may choose not to allocate all of its share of the infrastructure to the network R&D experimenter. Even during off-peak hours when the networks can make use of all of the available infrastructure, the production component may choose to keep a small portion of the production infrastructure available for non-real-time production traffic. This temporal, elastic, on-demand control of hardware layer infrastructure can greatly reduce our need for costly redundant services, circuits, switches, and routers.

3.3   The Media Layer

The model that provides a production bar of services at the media layer (e.g., ATM) assumes that the hardware layer is of production quality and takes the model of infrastructure sharing one step further by supporting both R&D and production services over the same physical media. For example, one can provide an ATM virtual path or circuit for the production traffic as well as a separate and distinct ATM permanent path or circuit dedicated to the experimental network research (e.g., implementing both IPv4 and IPv6 in native mode). A single switch, if appropriately designed and implemented, can satisfy both the R&D and production requirements by supporting experimentation with ATM signaling and QoS at the same time production traffic is passing through the same switch. It is important to provide mechanisms in a switch that ensure that one type of traffic (e.g., experimental) does not bring down the switch or trample the other type of traffic (e.g., production). The use of a redundant, yet separate, internal fabric within the switch is an example of such a mechanism.

3.4 IP

The normal mode of operations employed by today's Internet providers relies on IP as the production network bearer service. In this example, the IP bearer service and all the infrastructure underlying it (i.e., the media and hardware layers) are considered production quality for IP-based applications. Applications may experiment with new middleware capabilities and services, such as RSVP for IP based QoS, but they expect that the IP bearer service is of production quality and will not be used for experimentation by network researchers. Any network experimentation at the bearer layer is accomplished by either using a separate infrastructure (e.g., Dartnet for RTP) or by using tunneling. The use of IP as the production bar provides as solid a production bearer network service as IP can deliver while allowing for experimentation with RSVP and other advanced IP-based capabilities.

Tunneling is a powerful tool that can be used to (1) minimize some of the need for duplicative infrastructure on a wide-area IP bearer service basis, and (2) reduce risk to the production bearer service layer. However, because tunneling does not necessarily address the requirement of an application that wishes to test and utilize a new network layer or network to MAC layer capabilities and infrastructure in native end-to-end mode, it should not be viewed as the only tool for concurrently supporting both a production and network R&D infrastructure. Tunneling not only delays the traffic's end-to-end trip, but it also requires the manual configuration of the virtual tunnels; as we saw with the virtual Mbone overlay, this does not easily scale when large numbers of sites become involved. Although tunneling may be useful during the first stage of the experimentation process, it is only a short-term answer for coexistence and may not truly test the routers and switches as they would be tested when they are supported in native (non-tunneled) mode. The model that concurrently supports a native-mode production and non-production bearer service in the routers by no means contradicts the goal of one common bearer service as described in the often-referenced National Academy of Sciences (NAS) publication, "Realizing the Information Future."[x] Rather, it addresses the reality of overlapping time lines for the "network of the past," the "network of the present," and the "network of the future," evidenced today by legacy networks, IPv4 and IPv6 (respectively). These three phases will always be in existence on any given network, although the actual IP versions may change over time, and should be considered a normal state of affairs[xi] . We will always be improving the bearer service (e.g., Multicast in IPv4) as well as introducing new bearer services or versions (e.g., IPv6). Multi-protocol routers implement a version of the concurrent bearer services model when they support concurrent multiple protocols such as IP, IPX, and SNA in native mode.

3.5 Middle Layer

Applications require the existence of many production-quality middleware services to support experiments with new network technologies and to provide the enhanced distributed computing environment capabilities that are required if these experiments are to be tractable. For example, when RSVP makes it to production status, we will see many experiments in which application developers attempt to improve application performance by representing explicitly the varied array of network QoS associated with different application components. In this case, the production bar would be RSVP and it would simultaneously support both production and experimental networking at the application layer. Other middleware production services may include name servers, security key and certificate infrastructure servers and authorities, directories, session managers (e.g., SDR[xii]), advanced IP-based capabilities such as the Mbone, and resource information and scheduling services such as those being developed in the Globus project[xiii].

3.6 Applications

Many applications programmers are constantly in search of new technologies and will use any available technologies to advance their programming environments and capabilities. Many are more than willing to use experimental facilities and will make use of the varied array of production bars previously mentioned, either in a concurrent or temporal mode. Advanced application programmers require the ability to set QoS parameters, monitor infrastructure, and experiment with new network capabilities to support their advanced application and programming environments. One application may require raw access to the SONET or ATM infrastructure via relevant QoS activation and signaling techniques, while another application concurrently requires a production IP layer to support experimentation with RSVP. The infrastructure needs to be able to support both of these requirements simultaneously, on both a short-term (seconds to minutes) and long-term (hours to days) basis.

4.0 Scope of the VPNS

In order to deploy an infrastructure that supports both production and experimental network research, telecommunications service providers need to adopt a new customer-supplier model. In this model, the customer and service providers would work together to define the service elements, network management tools, and administrative models and architecture necessary to support the customers' requirements and their view of the network, as well as that of the telecommunication and Internet service providers. This model requires the telecommunications carrier and service providers to work with the customer in the

standards arena to define appropriate end user tool and access-to-information capabilities. It also requires the ISPs to be more open with respect to customer non-intrusive access to network and switch state information. This information includes QoS, circuit or access class information, traffic flows, error status, MIB variables and other state information on an end-to-end basis that the end user community requires to monitor and verify its network services. The customer may also require the ability to dynamically configure, reconfigure, and acquire network infrastructure resources based on end user QoS or policy requirements. This will involve the support of active network components (e.g., circuits, switches, routers, multiplexors) in the infrastructure as well as the signaling and op-code capabilities required to dynamically trigger a reconfiguration. In order to fully utilize these capabilities, applications will require state information and appropriate tools for determining what network infrastructure may be available to them at any given time and for reserving the appropriate network resources in a dynamic fashion, whether that be on a millisecond, minute, hourly, or daily reservation basis.

In addition to enhancing non-intrusive access to network state information on an end to end basis, the ISPs also need to work with network research experimenters to define what is necessary to support the network research on their infrastructure without interfering with the production traffic. This might include providing the researcher with the ability to dynamically alter configurations and settings in a dedicated R&D switch and add/drop multiplexors, or providing safe toggles and state changing tools in production switches to affect network management and monitoring tools. All of this is further complicated by the fact that different network management models and tools are required to support the different thresholds and levels of comfort associated with production and experimental traffic. An adaptive network application infrastructure (e.g., active network control over multiplexors, switches, circuits, and routers) programming interface (API) would make it possible for the end user to easily move between production and experimental modes and infrastructures, easing the pain of living in both policy worlds.

The end user may have agreements or contracts with various service providers, each with a different scope, ranging from the campus to the regional area as well as to the wide-area network (WAN). The continuing deregulation of the industry will blur the distinction between regional and wide-area providers, but the location of the actual physical infrastructure still favors regional economies of scale (e.g., major metropolitan areas), so collaboration between close physical or cultural institutions will prevail. In any event, the issue of supporting production and network research on the same infrastructure will need to be addressed on a campus, regional, and wide-area level. A customer's service may be provided by many nested layers of ISPs, some of whom obtain services from other providers. As a result, there is a need to ensure that the end user and network managers have the capabilities and tools necessary for navigating and monitoring the many nested

layers of ISPs, as well as peering points, so that the customers can support their applications on an end-to-end basis.

Regardless of the scope, the major focal point of the concurrently supported infrastructure will be at the customers' demarcation point, commonly referred to as the "edge," where the customer's equipment interfaces and peers with that of the service provider–whether it is at the campus, regional, or WAN level. In fact, the end user may be peering with each of these concurrently. The importance of the assumption regarding the provider's cloud demarcation point is that a service provider can support the production and experimental network traffic any way it chooses within its cloud or infrastructure. For example, an ISP may choose to use one switch and a single fabric, or use separate switches and lines as long as the access interface and expected or contracted services to the end user are met. QoS and network management capabilities rely on the ISPs implementing and supporting standards and tools on an end-to-end basis across the campus, regional, and WAN network infrastructures.

4.1   Crash and Burn Test Bed

The local "crash and burn" test bed is the simplest to envision and support because it can be built as a separate small network on a departmental basis. This is the "Bonneville salt flats" model for performing network research and development; it is usually the first choice for the alpha testing of experimental network protocols because if you crash while trying to break the speed record, you do not adversely affect the production applications. This model normally employs a separate, dedicated local network on a room, building, or campus basis whereby the R&D network never connects to or exchanges traffic with the production network. It is easy to manage, provides excellent access to the researcher, and is very flexible, but it does not scale well.

Many organizations also utilize a small number of demonstration or test routers and switches in a separate "sandbox" for the purpose of testing version upgrades and enhancements to network protocols and architectures. However, they normally cannot afford the number of routers or switches necessary to properly test these upgrades and enhancements under expected real-life traffic and stress. Regardless of the amount of testing that is done before deployment, when the upgrades or enhancements are finally enacted in the routers and switches, the production network becomes an experimental network until the modifications are demonstrated to have no ill side effects.

4.2   Shared Campus Infrastructure

The Shared Campus infrastructure is an attempt to share as much of a campus local area network (LAN) infrastructure as possible to support both the production traffic and the network R&D traffic and experiments. This model is attractive because it allows for the easy introduction of "guinea pig" user applications that not only test the new networking capabilities, but also allow the applications to adapt to the new infrastructure on a pre-production basis. These applications normally run on the production network. However, there are a number of users who are willing to test or stress the experimental network even though it may crash. Application programmers are willing to do this because they derive more benefit from the early adoption of the advanced capabilities or bandwidth offered by the experimental network than the cost or pain associated with the conversion of their codes to take advantage of the new capabilities. This model can be implemented with completely separate network segments for the production network and the experimental R&D network, or it can be built of separate segments that share some subset of gateways, routers, and switches. In a shared network, the traffic may "cross in the night" as it passes through the routers or switches (e.g., virtual LANs [VLANs], ATM private virtual paths [PVPs], or shared routers). The campus network manager may choose to support both types of traffic on the same regional or WAN link as described in Sections 4.3 and 4.4. The challenge on the campus level is how to operate and manage the shared gateways and switches, and how to define a campus network operation center (NOC) that is responsive to both the requirements and thresholds for production and research activities.

The campus LAN will continue to be a heterogeneous mixture of LAN technologies providing the "last foot" to the desktop, including ATM and non-ATM technologies, such as 100 Megabit and Gigabit Ethernet. Because of this heterogeneous mixture, applications will require the development and deployment of integrated solutions that map layer-three-based services (e.g., RSVP) to layer-two services (e.g., ATM or switched Ethernet), including those supporting QoS and network management. In order to take advantage of the QoS capabilities available in layer-two services, applications require the capability for some level of cross-layer signaling (e.g., RSVP to ATM). In situations where a high-speed server is located directly on an ATM network, the application will need to be able to directly view and control the layer-two QoS parameters. In addition, there will be situations where a high-speed server is located on a very-high-speed, non-blocking switched Ethernet segment, or it is the only node on a high-speed broadcast segment. Because these latter two scenarios carry no possibility of media collisions or contention, we need to explore ways to extend bona fide layer-two QoS (e.g., ATM) across these traditionally non-QoS supporting

media so that the applications can achieve end-to-end QoS in a heterogeneous media environment.

4.3   Shared Regional Infrastructure

Because the local loop usually accounts for approximately 30%-50% of the cost for connecting to either a regional or WAN ISP, major link/access cost savings can be realized by multiplexing a local loop to support both production and network R&D traffic and applications. This approach can generally be achieved in two different ways, depending on the user's level of trust that one type of traffic will not adversely affect the other type of traffic.

The "no trust " scenario, which might be invoked to support very experimental research, would use two sets of switches on either end of the local loop (see Section 3.2) with two switches located on the campus and two switches located at the loop demarcation point where the local loop enters the carrier's cloud. The traffic is separated on the local loop such that the only infrastructure shared by the two types of traffic is the local loop itself, not even the switches. It is important to note that the service access interface and agreements that users have with their carriers will determine whether both sets of traffic could eventually be carried over the same lines and switches inside the carrier cloud or carried on distinct infrastructure. The disadvantage of this approach is that extra switches are required to implement this scheme. On the other hand, the advantage perceived by some for separate infrastructure is that the two types of traffic are kept physically separate, which reduces the risk of any problems that may arise from the inadvertent confluence of the two types of traffic. The support of both the production and R&D environments may be achieved through the aggregation of various network infrastructure segments and components, which may be dynamically combined and configured to produce a temporary production or experimental network.

The "guarded trust" model entails one set of switches on either end of the loop in addition to the sharing of the physical local loop. The separation of R&D and production traffic at this level can be easily accomplished via the use of ATM PVPs or Permanent Virtual Circuits (PVCs), assuming that there are guarantees that no bleed-over from one type of traffic to the other occurs or that no errant application can adversely affect the other type of traffic due to congestion control, buffer management, QoS management, or any other policy enforcing algorithms implemented in the switches. Because the separation of traffic, either based on type or policy, is not accomplished in hardware, users as well as network managers and providers require tools that they can use to monitor the network infrastructure and assure themselves that their requirements are being met.

13

Either production or R&D networks could make use of segmented network infrastructure, in which switches, routers, and muxs are assumed to be either for production or for experimentation purposes and can be dynamically aggregated into virtual networks on demand. It is also apparent that this capability can be easily adopted by the commercial sector for supporting end- user demands for temporary network requirements for trade shows, demonstrations, proofs of concept, and temporal use of additional bandwidth. This type of capability can be supported through the use of adaptive hardware devices and techniques such as end user on-demand control of Sonet drop/add multiplexors, aggregating/de-aggregating WDM color frequency multiplexors, or real-time manipulation and configuration of switches and routers. In order to support this capability, the telecommunications industry needs to alter its business and technical models to not only provide non-intrusive access to network state information but also to provide the ability for the end user to safely manipulate the network infrastructure to create either production or R&D networks as they need them, even if under special circumstances and for only a short time period.

We can extend the concept of regional sharing of infrastructure one step further by defining a network peering point–where multiple local entities and institutions can connect and peer with each other–and providing a common funnel and peering point with WAN ISPs such as Sprint, MCI, the vBNS and ESnet. The Network Access Points (NAPs)[xiv] were originally designed to support this model, but the implementations failed in this regard because they only provided ISP-to-ISP peering. The Gigapop is the latest iterative concept and attempt to support a communal sharing of infrastructure to peer local institutions with advanced production services and ISPs. We contend that the Multimode Gigapop (M-Gigapop) extends the Gigapop and NAP concept because it can concurrently support both production and R&D traffic on as much of the same infrastructure as possible and hand off the traffic to the appropriate commercial or R&D ISP, depending on the type of traffic. The research challenges again are how to ensure that one type of traffic does not adversely affect the other at the M-Gigapop and how to provide for distributed network management of the peering point(s) (i.e., what end user tools and management capabilities are required in the switches, routers, and multiplexors).

4.4   Shared WAN Infrastructure

When providing shared wide-area infrastructure, the telecommunications service providers (e.g., MCI[xv]) and ISPs (e.g., ESnet and vBNS) will face many of the same issues as the traditional regional carriers (e.g., Ameritech[xvi]) and ISPs (e.g., CICnet[xvii]). The major issues center on what

14

access interface and capabilities are provided to the end user and how experimental traffic, if any, is supported on the same or separate infrastructure as production traffic. For example, all experimental traffic may be provided over physically separate circuits and switches within the WAN ISP's cloud. The ability of the telecommunications carriers to provide multi-modal infrastructure may be hindered by the fact that some of their customers do not like to assume any risk. The federally funded private WAN ISPs (e.g., ESnet, NSI, vBNS, DREN[xviii]) may have a little more latitude in supporting some experimental network traffic and capabilities within their clouds, but they are also reluctant to assume much risk because some members of the application research community they support expect absolute production-level services. However, the challenge still facing all ISPs who expect to be solvent and viable service providers in the future will be how to support multiple varied policy (e.g., production versus experimental or guaranteed versus best-effort services) virtual networks because it is too costly for both the end user and the provider to support duplicative infrastructures (for the reasons already outlined in this report). Small amounts of calculated risk are critical in the evolution of networks and must be assumed by the end user and the service providers. Even when we test router or switch upgrades in a bounded environment prior to deploying these changes into production networks, we still assume some risk when we finally deploy the upgrades because any change to the running system or network in effect changes it from a production to an experimental network, albeit a controlled one. We all can think of many occasions where seemingly small upgrades or modifications have caused far-reaching problems. We need to develop networks that are more resilient and fault tolerant (i.e., can support experimental as well as production traffic and be dynamically configured to compensate for problems) on both a macroscopic and microscopic level. The on-demand use and re-use of network infrastructure components and segments will further enable the service providers to support both the production and experimental requirements, as well as the other varied and sometimes conflicting policy-based network requirements of its customers in a more efficient and cost-effective manner.

Because we can expect to see the use of ATM continue for provision of regional and WAN service, we need to address the issue of ATM QoS support in the ISP clouds as well as access to these capabilities by the end user. One approach is to treat the ATM cloud as only a raw bit pipe and to rely on techniques such as RSVP to provide end-to-end QoS across not only the non-ATM LAN technologies (see Section 4.2), but also the carriers' ATM clouds. This type of approach defeats one of the major reasons an end user would consider deploying ATM on the campus or explicitly request it for WAN services. One can argue that RSVP QoS is not the hard QoS some applications require, and therefore we should utilize ATM QoS whenever possible. In either case, the ability of the end users to use ATM QoS

signaling in a dynamic fashion to satisfy their dynamic application requirements is dependent on the availability of standards-based signaling implementations and APIs in the switches and end host systems, as well as admission control capabilities for both ATM and RSVP. The current state of deployment for ATM equipment that can support applications dynamically signaling and managing QoS in regional and WAN networks is fairly poor; this may impede the adoption of native ATM by the end user community. The lack of RSVP admission control tools available for use by the end user and network manager, as well as the lack of admission policies based on the application and campus network manager's perspective, may also impede the adoption of RSVP.

4.5   Impact of Shared Infrastructure on the End System

The concurrent support of production and R&D infrastructure must extend to and include the workstation. The current mode for supporting multiple-network use policies is based on the use of separate workstation and IP network addresses for the production traffic, and a separate workstation and IP address for the experimental R&D traffic. The R&D IP address must be garnered from a Class B, C, or Classless Internet Domain Routing (CIDR[xix]) address block that is different from the one used for the production network. The IPv6 address space is much larger than that utilized by IPv4; however, there is nothing in the IPv6 address or routing specification that will alter the need for using separate addresses from different address spaces in order to support multiple policies on the same end node. Hybrid solutions exist that involve using a workstation with two network interface cards (NICs), each having an address on different networks (e.g., different CIDR blocks). The reason for selecting addresses for the production and R&D NICs from different network address spaces, or for multihoming the two addresses on the same NIC, is to ensure that, when necessary, the production traffic takes a different route over the infrastructure than that taken by the experimental R&D traffic. Given the fact that current IP routing algorithms choose routes for traffic based on the network portion (e.g., top 24 bits of a Class C address) of the destination address, we have no option but to use two separate addresses to enforce the varied policies associated with production and R&D networks. This is an issue that mostly affects the end user, the workstation, and possibly the campus network because the regional and WAN clouds are treated primarily as switching engines at the IP level and will route any packet based only on its destination IP address and the associated routing table entry (which indicates which interface provides the next-best hop for the packet on its way to the destination).

The practice of using two different IP addresses on a workstation from different network address spaces or subnets is referred to as multihoming[2] and gives one workstation the ability to send and receive traffic over two distinct networks or subnets based on policy. Some workstations possess the ability to support two distinct IP addresses on a single NIC, thereby achieving the same result with only one NIC. One can bind the appropriate workstation source IP address when opening a socket for transmission (i.e., binding the production IP address as the source address in packets when the application is doing production work, and the experimental IP address when the application is performing network research). However, there is no way for the application programmer to know which IP address on the destination workstation or server belongs to the production or experimental subnet. Several methods can be used to solve this problem. The first method requires the user to possess a prior knowledge about which host IP addresses of the destination node are on the production or research subnets. The second method uses a local configuration file (i.e., a "hosts.exp.txt" file) that lists the domain names and IP addresses of all the experimental hosts and subnets. This method assumes those host addresses not listed in this file are used for production purposes. The third method involves making modifications to the Domain Name System[xx] (DNS) to identify experimental host addresses. This would allow for a site administrator to define experimental hosts in the DNS and thus leverage off an existing and scaleable infrastructure. The fourth method makes use of VLAN technologies to build experimental R&D subnets that extend across the campus and possibly regional or wide-area networks.

In the effort to reduce the amount of infrastructure required to concurrently support production and R&D environments, we would like to minimize the amount of hardware required by the end user to easily live within both a production and R&D environment. Ideally this would entail using only one workstation, one multihomed NIC, and one physical subnet. It would also allow applications to move between production and R&D environments on their screens simply by moving their mouses from the production window to the R&D window and vice versa. This requires that state information associated with that process be handled appropriately as part of a processes' normal context switch. The end user should be able to specify that a particular window and/or environment is either for experimental or production use and the kernel within the node must be able to determine

---

[2] Some administrators propose using separate workstations and network infrastructure to avoid the administrative issues associated with multihoming. However, it may prove to be more efficient to multihome the relatively small number of workstations that require both a production and research address, and to rely on DHCP to dynamically configure IPv4 production hosts and the use of IPv6 link and local address capabilities to dynamically assign addresses to IPv6 end nodes.

which mode is active so that it may act appropriately (i.e., set the correct source IP address in the outgoing packet). The kernels on both the sending and receiving nodes must verify that only experimental-to-experimental and production-to-production traffic flows occur.

5.0 R&D Challenges

The need for advanced programming environments for the application and end user domains is driving the need to support network research in the area of network management tools. Application programmers require the ability to monitor, analyze, and debug their applications, including the impact of network traffic conditions. Network managers require the ability to protect, ensure, monitor, analyze, and debug the network services that they are providing. To support concurrent production and experimental activities, the suggested R&D areas of focus are on network management as well as end user tools for utilizing a shared infrastructure that is as efficient and error free as possible. Providing dual-modality network capabilities (i.e., production and research) with sufficient safeguards requires advances for ATM and IP (both IPv4 and IPv6) in the areas of network management, QoS, admission control, cost accounting, and end station dual-modality support. It is important that the application programmer and network researcher be able to utilize network resources to meet their programmatic goals; the campus network manager and other service providers (MAN, WAN) must be able to manage and fully utilize scarce network resources. The adaptive, on-demand configuration and management of lower-layer network infrastructure (e.g., add/drop multiplexors, switches, routers, and network segments) greatly enhances the ability of service providers to support multiple policy and multimode virtual networks on the same infrastructure. Much of the experimentation with protocols, switches, and routers has been initially focused on the campus level. While the network researcher's focus will most likely be initiated on the campus level, it is important to focus on the end-to-end applications performance, which will undoubtedly include the campus to ISP demarcation point. It is imperative that ISPs and carriers support the QoS and non-intrusive end-to-end network management tools and capabilities that are required by the applications and the campus/LAN network managers to determine network performance characteristics. It is also crucial that ISPs and carriers support network research capabilities as part of their infrastructure because they derive direct benefit from the results, regardless of whether it is via dedicated or shared infrastructure.

5.1 Network Management

Applications programmers require real-time network diagnostic and analysis tools that can be utilized for monitoring services and debugging on an end-to-end basis across the multitude of campus, regional, and WAN network infrastructures. They also require tools to utilize QoS to dynamically adapt their application to utilize network services. While the traditional notion of

18

an NOC that monitors network activities remains important, advanced network capabilities call for new weapons in the network management arsenal.

There are some network management tools and capabilities that are commonly required and employed among the WAN, LAN, and campus network managers; however, there are also capabilities that may be unique to each one of these areas. In particular, the tools utilized by the ISPs providing regional and wide-area networking services will most likely intersect but not necessarily be a proper subset of those tools employed by the campus/LAN network manager. Many tools in the ATM environment to date have been proprietary. For ATM to be widely adopted, more interoperable management and debugging tools need to be available. Standards bodies such as the ATM Forum and the IETF need be lobbied to get vendors to adopt interoperable management and debugging tool suites. The following is a non-exclusive, initial list of basic capabilities that the campus/LAN manager will need to support a dual-mode infrastructure and provide for the applications' and network manager's requirements.

- Reference implementations of ATM device discovery, ATM ping and ATM traceroute across heterogeneous vendors' equipment.

- Reference implementation of ATM QoS traceroute that traverses each switch and returns the QoS on a per PVC, SVC, PVP basis.

- The ability to securely (i.e., authentication) manage admission control, cost accounting, and priority policies, as well as QoS support for RSVP, IPv6 and ATM (i.e., fair and efficient tools for allocating network resources, including priority bidding and cost accounting).

- Tools for debugging cross layer signaling, admission control and other QoS capabilities for RSVP, ATM, and any other QoS supporting protocols.

- The ability to support both production and network experimental R&D activities on the same infrastructure with varying degrees of thresholds, alarms, and required responses.

- The ability to debug IP and ATM networks concurrently, including the ability to capture, analyze, and display cells, packets, and flows.

- Distributed inter-NOC capabilities with other regional, WAN, and LAN/campus network NOCs, including the ability to exchange

trouble ticket information as well as to remotely view the state of the network from another NOC's point of view or point of presence.

- VLAN and switched Ethernet management and analysis tools that are integrated with or can be run concurrently with IP and debugging tools to promote a coherent multiple-layer view of the network.

- The ability to load beta versions of code into the LAN or WAN experimental infrastructure components to enable the testing and experimentation of new capabilities while concurrently running production traffic.

- Support of the negotiation of end-to-end QoS IP "best effort" services, non-ATM, switched Ethernet, and other broadcast technologies.

- The ability to dynamically configure, use, and re-use network infrastructure segments and components (i.e., Sonet drop/add multiplexors, ATM switches, routers, links, and circuits) from a pool of networking resources and objects to create on-demand virtual production or experimental networks.

Regional and WAN ISP managers require many of the same tools that the campus/LAN managers utilize (listed above); however, they also require the following additional tools and capabilities if they are to support the concurrent multi-modal use of infrastructure:

- Tools and capabilities that allow non-intrusive monitoring, analysis, state and data gathering, querying, and providing QoS support for the application programmer as well as the campus/LAN or regional network manager.

- The capability for the campus/LAN manager to support both a production and experimental network environment across WAN infrastructure with as little risk as possible.

- Capabilities for applications to easily request particular classes of traffic (e.g., UBR and ABR) burn rates as well as manage QoS for ATM and RSVP at the carrier demarcation points.

- Provision of multiple dynamic classes of QoS between endpoints.

- Provision of IP and integrated services protocol support for secure RSVP and IPv6 policy, cost accounting, and admission tools that will

provide the end user or campus network manager with the ability to securely control the use and management of their traffic on an end-to-end basis in native mode.

- The capability to support the loading of beta versions of code into the experimental infrastructure components at the request of or via issuance of commands by the campus network researcher to enable the testing and experimentation of new network capabilities.

5.2 Application

Application programmers require network management tools that they can use to determine the state of the network in real time in order to debug their distributed applications, determine whether the network is functioning up to expected levels, dynamically configure and manage virtual production network services, and query and request appropriate QoS. This last requirement includes the ability for cross layer (e.g., RSVP or IPv6 to ATM) signaling to affect the required environment as well as to bid for priority status when resources are scarce. These tools may be used directly by the programmer or accessed automatically by programs running on behalf of the programmer. For example, an adaptive parallel application might be constructed to use a research network when it is available–or the production network when it is not (or vice versa)–or to interact with the research network management system to tune system parameters. In all cases, a key issue will be providing tools that can translate between low-level network constructs (e.g., ATM QoS) to the higher-level tools and concepts used by application programmers.

The environment for the programmers can be greatly enhanced by providing them with the capability for migrating seamlessly between production and experimental status on one workstation with the mere movement of their mouses from the production window to the experimental window and vice versa. The application programmers may also wish to avail themselves of multiple levels of production network infrastructure. For example, they may implement production-quality IPv4 and experimental IPv6 services over a production ATM network while at the same time running both production and experimental applications over the production IPv4 services. Specific tools and capabilities required by the end user for making use of the dual-mode infrastructure include the following:

- Reference implementations of ATM device discovery, ATM ping, ATM traceroute, and other ATM tools across heterogeneous vendors' equipment that the user can invoke to ensure that the ATM pipe is functioning at expected levels.

- Convenient application-level interfaces to information provided by ATM tools, enabling applications to determine and then adapt to changes in both link-level and end-to-end performance.

- Reference implementation of ATM QoS traceroute that traverses each switch and returns the QoS on a per PVC, SVC, PVP basis.

- Application support for cross layer (i.e., IP to ATM and RSVP to ATM) QoS signaling, including querying and invocation.

- Direct view and control of layer-two QoS parameters by applications wishing to use high-speed, QoS-enabled servers.

- The ability for an application to negotiate end-to-end QoS over non-ATM, switched Ethernet, and broadcast network segments when the destination node is either the lone server on a broadcast network or is on its own virtual LAN through switched Ethernet or other VLAN technology.

- Seamless dual modality (i.e., production and R&D) application support on one workstation.

- Ability, when required, to bid securely (i.e., authenticated) for priority use of scarce network resources (assumes some cost accounting and bidding system capabilities).

- Secure (i.e., authenticated) admission control querying and reservation capabilities for RSVP, IPv6, and ATM.

- Session control tools that seamlessly integrate QoS tools, as well as the multimedia, directory, information agents, and labspace environments tools.

- Integrated analysis and debugging tools and capabilities to support the integrated QoS, multimedia, information agents, and labspace environments.

- Versions of application programming tools (e.g., MPI, CORBA) that are enhanced with QoS signaling capabilities and performance tools that can explain observed network performance in application terms.

- Scheduling tools that allow programmers to specify network QoS requirements as well as computing or data requirements when requesting the resources required for a particular computation.

- Distributed computing environments that securely utilize network-based admission control techniques for ATM, RSVP, and other QoS protocols and services.

- The ability, as well as the application programming interface (API), to dynamically configure, use, and re-use network infrastructure segments and components (i.e., SONET drop/add multiplexors, ATM switches, routers, links, and circuits) from a pool of networking resources and objects to create on-demand virtual production or experimental networks. A standard characterization of network infrastructure segments and components from an end user's and application's perspective is required so that the end user can correctly create, configure, and monitor virtual network infrastructure.

The validation and evaluation of these tools and concepts will require access to a suite of interesting applications that can be used to stress various aspects of the multi-modal network infrastructure. Examples of such applications include the following:

- Distributed collaborative engineering applications in which engineers at different sites collaborate on the design and evaluation of complex systems. The Argonne BoilerMaker[xxi] system is an example of such an application; this allows engineers to use virtual reality systems to guide placement of inlets in a simulated industrial boiler.

- Remote I/O applications in which, for example, programs running on a supercomputer access input datasets or create output datasets located on remote file systems. These applications require the ability to manage and monitor network QoS to achieve high supercomputer utilization when streaming data between supercomputer and remote file system. The Argonne RIO project[xxii] is developing infrastructure for such applications.

- Remote visualization applications, in which data produced on a supercomputer by a scientific instrument or read from a file are streamed to a display device at another location. Many such applications were demonstrated as part of the I-WAY project.

- Remote instrument control applications, which may feature multiple data streams with different characteristics, including time-critical control data, high-bandwidth video, and audio.

- Distributed computation applications in which large numerical computations are distributed over multiple distributed computing resources in order to solve larger problems or to solve fixed size problems more quickly. Computational chemistry and astrophysics are two examples of disciplines in which this approach has been applied successfully.

6.0   Conclusion

Application developers are rarely eager to invest a large amount of effort and time to convert their codes to "test drive" new network technologies, especially if the infrastructure is to be short lived. Yet the development and deployment of new architectures and protocols are extremely dependent on applications, without which it is not possible to test and stress the infrastructure or to validate that it works with real applications and can be deployed in production mode. For example, the I-WAY[xxiii] network developed to support Supercomputing 95 succeeded, by virtue of tremendous effort, in demonstrating the benefits associated with an advanced pre-production infrastructure; yet this infrastructure evaporated immediately after the close of Supercomputing 95, making it difficult for many of the principal investigators and institutions to continue their collaborations. Network researchers need real applications and traffic to use and stress their experimental and production networks, and application developers are constantly seeking new network capabilities to enhance their computational environments. Neither group can progress without a persistent high-end, advanced infrastructure and without addressing the daunting cost associated with concurrently supporting both a production and experimental infrastructure. We must endeavor, then, to find the technical, social, and political means necessary to share as much infrastructure as possible at the campus, regional, and wide-area network level to support both production and experimental R&D activities.

7.0   Acknowledgments

Electronics and Computing Technologies Division and Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Illinois 60439

8.0   References

[i] Computer Science and Telecommunications Board, The National Research Council, Computing and Communications in the Extreme, National Academy Press, Washington, D.C., 1996

[ii] http://www.cise.NSF.gov/ncri/nsfnet.htm

[iii] http://www.es.net

[iv] CAIRN is the successor to DARTNET - http://www.fnc.gov/cairn.html

[v] Paxton, V., Floyd, S., Wide-Area Traffic: The Failure of Poisson Modeling, IEEE/ACM Transactions on Networking, Vol. 3., No. 3, pp 226-244, June 1995

[vi] IPv6 is also known as IPNG and is defined in the following Internet Engineering Task Force (IETF) Request for Comments (RFCs); S. Deering, R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, 1/04/1996 (http://ds.internic.net/rfc/rfc1883.txt); Y. Rekhter, T. Li, An Architecture for IPv6 Unicast Address Allocation, 1/04/1996 (http://ds.internic.net/rfc/rfc1887.txt); A. Conti, S. Deering, Internet Control Message Protocol for the Internet Protocol Version 6 (IPv6), 1/04/1996 (http://ds.internic.net/rfc/rfc1885.txt); M. Borden, E. Crawley, B. Davie, S. Bastell, Integration of Real Time Services in an IP-ATM Network Architecture, 8/11/1995 (http://ds.internic.net/rfc/rfc1821.txt); S. Deering, R. Hinden, IP Version 6 Addressing Architecture, 1/04/1996 (http://ds.internic.net/rfc/rfc1884.txt)

[vii] htpp://www.atmforum.com/atmforum/atm_introduction.html

[viii] http://www.ietf.cnri.reston.va.us/html.charters/rsvp-charter.html

[ix] The Message Passing Interface (MPI) - http://www.mcs.anl.gov/mpi/index.html

[x] Computer Science and Telecommunications Board, The National Research Council, Realizing the Information Future, National Academy Press, Washington, D.C., 1994

[xi] Aiken, R., Cavallini, J., Standards: When are they too much of a good thing?, ACM StandardView, June 1994, Interop Connexions, August 1994, Harvard NII Standards Workshop Proceedings, MIT Press, May 1995

[xii] http://cs.ucl.ac.uk/mice/sdr

xiii Foster, I., Kesselman, C., Globus: A Metacomputing Infrastructure Toolkit, Intl. J. Supercomputing Applications, 1997 (to appear). See also http://www.globus.org/

xiv Aiken, R., Braun, H., Ford P., NSF Implementation Plan for the Interagency Interim National Research and Education Network (NREN), General Atomics/San Diego Supercomputer Center, GA-A21174, May 1992

xv http://www.mci.com/

xvi http://www.ameritech.com/welcome/

xvii http://www.cicnet.net/

xviii The Defense Research and Engineering Network, http://www.arl.mil/HPCMP/DREN/drenexe3.html

xix CIDR is defined in the following IETF RFCs; R. Hinden, Applicability Statement for the Implementation of Classless Interdomain Routing (CIDR) (ftp://ds.internic.net/rfc/rfc1517.txt); Y. Rehkter, T. Li, An Archirecture for IP Address Allocation with CIDR, 9/24/1993 (ftp://ds.internic.net/rfc/rfc1518.txt); V. Fuller, T. Li., J. Yu, K. Varadhan, Classless InterDomain Routing (CIDR): An Address Assignment and Aggregation Strategy, 9/24/1993 (ftp://ds.internic.net/rfc/rfc1519.txt); Y. Rehkter, C. Topolcic, Exchanging Routing Information Across Provider Boundaries in the CIDR Environment, 9/24/1993 (ftp://ds.internic.net/rfc/rfc1520.txt)

xx P. Mockpatris, Domain Names - Concepts and Facilities, 11/01/1987, (http://ds.internic.net/rfc/rfc1034.txt); P. Mockapetris, Domain Names - Implementation and Specification, 11/01/1987 (http://ds.internic.net/rfc/rfc1035.txt)

xxi Diachin, D., Freitag, L., Heath, D., Herzog, J., Michels, W., Plassmann, P., Remote Engineering Tools for the Design of Pollution Control Systems for Commercial Boilers,
Intl. J. Supercomputer Applications, 10(2): 208—218, 1996.

xxii Http://www.mcs.anl.gov/globus/RIO/

xxiii DeFanti, T., Foster, I., Papka, M., Stevens, R., Kuhfuss, T., Overview of the I-WAY: Wide Area Visual Supercomputing, Intl. J. Supercomputer Applications, 10(2): 123-130, 1996. See also http://www.iway.org/

DISTRIBUTION LIST
ANL-97/1

Mr. Allan H. Weis
President & CEO
Advanced Network & Services
200 Business Park Dr.
Armonk, N.Y. 10504

Mr. Caitlin Brown
Senior Account Executive
Ameritech
225 W. Randolph
HQ Floor 23C
Chicago, Illinois 60606

Patricia A. Caine
Vice President, Ameritech Advanced Data Services
Ameritech
225 W. Randolph
HQ Floor 3A
Chicago, IL 60606

Mr. Joel Engel
Vice President of Technology
Ameritech
30 South Wacker Drive, Floor 38
Chicago, IL 60606

Mr. Jeffrey Kipnis
Sales Engineer
Ameritech
225 W. Randolph, HQ Floor 6B
Chicago, IL 60606

Mr. Hardial Mann
ATM Service Product Manager
Ameritech
2000 W. Ameritech Center Drive, Suite 2B31C
Hoffman Estates, IL 60196

Paul M. Ross
Manager of Design and Sales Engineering/Illinois
Ameritech
Two Westbrook Corp. Center, Suite 600
Westchester, IL 60154

Mr. Andrew G. Schmidt
Internet, NAP, and Intranet Product Manager
Ameritech
2000 W. Ameritech Center Drive, Suite 2A01
Hoffman Estates, IL 60196

Mr. Rich Wilson
Senior Account Manager
Ameritech
1011 S. Second St., Suite B
Springfield, IL 62704

Lamont J. Young
Sales Engineer
Ameritech
225 W. Randolph, HQ Floor 6B
Chicago, IL 60606

Dr. Lee Holcomb
Director, Aviation Systems Technology Division
3000 E Street, SW, RC
Washington, DC 20546

Mr. Steve Wolff
Cisco Systems
Business Development
380 Herndon Parkway, Suite 300
Herndon, VA 20170

Dr. Robert E. Kahn
President
Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 20191-5434

Dr. Howard Frank
Director, Information Technology Office
Defense Advanced Research Project Agency
3701 North Fairfax Drive
Arlington, VA 22203

Ms. Hilarie Orman
Program Manager
Defense Advanced Research Project Agency
3701 North Fairfax Drive
Arlington, VA 22203

Mr. Michael Roberts
Vice President, Internet II
Educom
1112 16th Street, N.W., Suite 600
Washington, DC 20036

Dr. Michael R. Nelson
Special Assistant for Information Technology
Executive Office of the President
Office of Science and Technology Policy
Washington, DC 20502

Mr. Jim Williams
Executive Director
FARNET
11735 Tuttle Hill Rd.
Milan, MI 48160

Mr. Donald E. Scott
Vice President, Government Relations
GTE Government Systems Corporation
1001 19th Street, N. Suite. 1100
Arlington, VA 22209-1732

Dr. James Leighton
Deputy Head, Networking and Telecommunications
Computing Science Directorate
Lawrence Berkeley National Laboratory
1 Cylotron Road, MS-50B-4230
Berkeley, CA 94720

Dr. Stu Loken
Director, Information and Computing Sciences Division
Computing Science Directorate
Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS-50B-4230
Berkeley, CA 94720

Dr. William McCurdy
Associate Laboratory Director Computing Sciences Division
Computing Sciences Directorate
Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS-50B-4230
Berkeley, CA 94720

Dr. Alexander Merola
Deputy, Computing Sciences Division
Computing Sciences Directorate
Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS-50B-4230
Berkeley, CA 94720

Vint Cerf
Senior Vice President, Internet Architecture and Engineering
MCI Communications
2100 Reston Parkway, Sixth Floor
Reston, VA 22091

Mr. Richard desJardins
EOS Network Manager
Code 505
NASA Goddard Space Flight Center
Greenbelt, MD 20771

Randy Butler
Technical Program Manager
NCSA
605 E. Springfield Avenue
Champaign, IL 61820

Mr. Charlie Catlett, Associate Director
NCSA
605 E. Springfield Avenue
Champaign, IL 61820

Mr. Larry Smarr, Director
NCSA
605 E. Springfield Avenue
Champaign, IL 61820

Dr. Donald Austin
Assistant Director of Planning
National Coordinating Office
High Performance Computing and Communications Programs
4201 Wilson Blvd., Suite 665
Arlington, VA 22230

Mr. John Toole
Director, National Coordinating Office
High Performance Computing and Communications programs
4201 Wilson Blvd., Suite 665
Arlington, VA 22230

Dr. Thomas Kalil
Senior Director, National Economic Council
National Economic Council for the White House
Old Executive Office Building, Room 233
Washington, DC 20500

Ms. Marjorie S. Blumenthal, Director
National Research Council
Computer Science and Telecommunications Board
Room HA560
2001 Wisconsin Avenue, N.W.
Washington, DC 20007

Mr. David Clark
Chair, Computer Science and Telecommunications Board
National Research Council
Laboratory for Computer Science
Massachusetts Institute of Technology
545 Technology Square, NE43-508
Cambridge, MA 02139

Dr. Aubrey Bush
National Communication Research Infrastructure
National Science Foundation
4201 Wilson Blvd., Room 1175
Arlington, VA 22230

Dr. Melvyn Ciment
Acting Assistant Director
Directorate for Computer Information
Science and Engineering
National Science Foundation
4201 Wilson Blvd., Room 1105
Arlington, VA 22230

Ms. Darlene Fisher
National Communication Research Infrastructure
National Science Foundation
4201 Wilson Blvd., Room 1175
Arlington, VA 22230

Mr. Steve Goldstein
NCRI
National Science Foundation
4201 Wilson Blvd., Room 1175
Arlington, VA 22230

Mr. Mark Luker
National Communication Research Infrastructure
National Science Foundation
4201 Wilson Blvd., Room 1175
Arlington, VA 22230

Dr. George Strawn
Director, Networking Communication Research Infrastructure
National Science Foundation
4201 Wilson Blvd., Room 1175
Arlington, VA 22230

Mr. Terence H. Matthews
Chairman of the Board & Chief Executive Officer
Newbridge Networks
603 March Road
Kanata, Ontario, Canada K2K 2M5

Mr. Ed Oliver
Associate Director for Computing, Robotics, and Education
Oak Ridge National Laboratory
Bldg 4500-N, MS 6259
Oak Ridge, TN 37831

Mr. William R. Wing
Network Architect
Oak Ridge National Laboratory
Bldg 4500-S, MS 6144
Oak Ridge, TN 37831

Mr. Michael F. Sobek
Manager, Sales Engineering Government Systems Division
Sprint
8330 Ward Parkway
Kansas City, MO 64114

Mr. John S. Cavallini
U. S. Department of Energy
ER-30
19901 Germantown Road
Germantown, MD 20874-1290

Dr. Dan Hitchcock (30 copies)
Acting Director
Mathematical, Information, and Computational Sciences Division
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874

Dr. Fredrick A. Howes
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874-1290

Dr. Thomas A. Kitchens
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874-1290

Dr. Dave B. Nelson
Associate Director of Energy Research
for Computational and Technology Research
U. S. Department of Energy
Office of Energy Research, ER-30
19901 Germantown Road
Germantown, MD 20874-1290

Dr. Rodney R. Oldehoeft
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874-1290

Dr. Mary Ann Scott
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874-1290

Mr. George R. Seweryniak
U. S. Department of Energy
ER-31
19901 Germantown Road
Germantown, MD 20874-1290

John Morrison
Project Leader for Accelerated
Strategic Computing Initiative
U.S. Department of Energy
Los Alamos National Laboratory
528 35th Street, MS-B260
Los Alamos, NM 87544

Steve Tenbrink
Deputy Group Leader for
Networking Engineering
U.S. Department of Energy
Los Alamos National Laboratory
528 35th Street, MS-B255
Los Alamos, NM 87544

Dr. Andrew White
Director, Advanced Computing Laboratory
U. S. Department of Energy
Los Alamos National Laboratory
528 35th Street, MS-B287
Los Alamos, NM 87544

Dr. Gregory Jackson
Associate Provost
The University of Chicago
5801 South Ellis
Chicago, IL 60637

Dr. Joel J. Mambretti
Director, Academic Information Technologies and Networking Services
The University of Chicago
1025 East 57th Street, 2nd Floor, Culver Hall
Chicago, IL 60637-2745

Ms. Maxine Brown
Associate Director, Electronic Visualization Laboratory
University of Illinois at Chicago
1120 Science & Engineering Offices
Box 4348
Chicago, IL 60680

Mr. Thomas A. DeFanti
University of Illinois at Chicago, EECS Department
Director, Electronic Visualization Laboratory
Associate Director for Virtual Environments, NCSA
851 South Morgan Street, Room 1120 M/C 154
Chicago, IL 60607-7053

Mr. Dan Sandin
Co-Director, Electronic Visualization Laboratory
University of Illinois at Chicago
1120 Science & Engineering Offices
Box 4348
Chicago, IL 60680

Professor Deephinder Sidhu
University of Maryland-BC
Baltimore, MD 21228-5398