

Science-Driven R&D Requirements for ESnet

Report from ESnet R&D Workshop, April 23-24, 2007

Report Contributors:

*William Johnston, Joseph Burreasca, Michael Collins, Eli Dart, Chin Guok, Michael Helm, Stan Kluz,
Joseph Metzger, Michael O'Connor, Kevin Oberman, and Daniel Peterson - ESnet*

Brian Tierney - LBNL

Charlie Catlett - ANL

Nagi Rao and Bill Wing - ORNL

Les Cottrell - SLAC

Tom Lehman - USC/ISI

Mauro Campanella - GARR

Contents

1	Introduction.....	2
2	Requirements Derived from Science Application Drivers	3
2.1	<i>Guaranteed Network Bandwidth</i>	3
2.2	<i>End-to-End Monitoring</i>	4
2.3	<i>Federated Trust</i>	4
3	Research Topics	4
3.1	<i>Virtual Circuits.....</i>	5
3.1.1	<i>Control Plane Issues.....</i>	5
3.1.2	<i>Management Plane Issues</i>	7
3.1.3	<i>Data Plane Issues</i>	7
3.2	<i>Dynamic Wave Management</i>	8
3.3	<i>Automatic Network Management</i>	8
3.4	<i>End-to-End Monitoring Services.....</i>	9
3.4.1	<i>Measurement Framework</i>	10
3.4.2	<i>User Initiated Debugging of End-to-End Service</i>	11
3.4.3	<i>Middleware Extensions for Network Monitoring</i>	11
3.4.4	<i>Virtual Circuit Monitoring.....</i>	11
3.5	<i>General Network Issues.....</i>	12
3.5.1	<i>Reliable Multicast</i>	12
3.5.2	<i>Routing Issues.....</i>	13
3.5.3	<i>Defining Common Services.....</i>	13
3.6	<i>PKI and Federated Trust Issues</i>	14
3.6.1	<i>PKI Credential Ease of Use</i>	14
3.6.2	<i>PKI Credential Security</i>	14
3.6.3	<i>Trust federation issues</i>	15
3.7	<i>Network Security Issues:</i>	17
3.7.1	<i>Multi-site view of network probing/intrusions</i>	17
3.7.2	<i>Separation of network control and data planes</i>	17
3.7.3	<i>DNSSEC</i>	17
3.8	<i>Collaboration services</i>	18
3.9	<i>Other Topics.....</i>	19
3.9.1	<i>Effective Connections for Leadership Computing Facilities</i>	19
3.9.2	<i>Automatic Test Suites for Protocol and Application Testing.....</i>	19
4	Conclusions.....	19
5	References.....	21
6	List of Participants.....	22

1 Introduction

The Energy Sciences Network, (ESnet) is a high-speed network serving tens of thousands of Department of Energy scientists and collaborators worldwide. ESnet provides high bandwidth network connections to all major DOE sites and, via comprehensive connections to the world's research and education ("R&E") networks, to essentially all U.S. and international R&E institutions in order to support the science mission of the DOE Office of Science ("SC"). This science environment is very different from that of a few years ago and the changes are placing substantial new demands on the network. The distributed, large-scale collaborations and data analysis supported by SC programs requires vastly more bandwidth than in the past, as well as new services. As scientific instruments get larger and more expensive in order to solve more complex and more subtle science problems, the smaller number of these instruments results in them being used by very large and dispersed science communities. The large scope of these collaborations in turn drives significantly greater requirements for data movement, distributed analysis, and integration of simulation and instrument data than in the past.

The next generation networks and services that are needed to support the new science environment are being defined by examining three types of requirements: the data generating and use characteristics of new instruments and facilities; changes in the process of doing science in this new environment, and; examination of the historical trends in network traffic and then projecting forward.

Additionally, advances in the state of telecommunications and network technology over the past several years enable radical new approaches to providing national-scale network services. Concurrently, middleware and applications architectures are increasingly supporting wide area distributed systems, introducing fundamental changes in the nature and profile of network traffic as observed by network providers. Hybrid network architectures - involving both traditional Internet connectivity and scheduled, targeted (circuit-like) capacity - offer the potential to provide a radically different set of interaction modes between the network and applications and middleware.

ESnet4 - the next generation SC network - responds to both the need for dramatic increases in bandwidth and for new circuit based network services. Initially, ESnet4 will be a hybrid network that mixes a conventional IP packet network and a circuit oriented network. While work has been going on for several years to define the circuit services, control mechanisms, and AAA mechanisms, many questions remain as to how various aspects of this service should be introduced and how they should function in production. Additionally, there is increasing demand for services supporting end-to-end performance, such as monitoring on a per-application basis, as well as other new services that are needed to move the network from a communal, best-effort service to a schedulable, dedicated service that can be incorporated into the managed resource environment of large-scale science experiments.

To move beyond proof-of-concept demonstrations toward persistent, reliable services it will be necessary to harvest the best concepts shown to be feasible and systematically move them into production services. To support this new set of services, and these new application and middleware systems, there must also be deep understanding of the behavior of the systems and the interactions between applications, middleware, and network services.

In order to address both the near term and long term need of the science community, ESnet requires a carefully constructed "roadmap" that takes into account the maturity of various capabilities and technologies, and the requirements of applications and their communities, to lay out a set of milestones and supporting efforts to reach those milestones. This will involve an

analysis of the progress of individual components and technologies being developed and/or tested by the ESnet team and others, followed by a prioritization and sequencing of technologies to address new requirements.

In April, 2007 the DOE Office of Science organized a workshop to bring together a small group of experts to work with the ESnet team to examine the current roadmap, roadmaps of similar enterprises, user requirements, and new technology options. The objective of the workshop was to create a new, multi-year technology roadmap for ESnet that identifies milestones and partners. The workshop is summarized in this document.

2 Requirements Derived from Science Application Drivers

The R&D topics addressed at this workshop fell, for the most part, into the major areas that emerged from the ESnet requirements gathering process [3], which is summarized in a series of reports [4][5].

- Guaranteed network bandwidth
- End-to-end monitoring
- Federated Trust

An overview of the requirements for each of these is described in the following sections.

2.1 Guaranteed Network Bandwidth

Large-scale collaborative science – big facilities, massive amounts of data, thousands of collaborators – is a key element of DOE’s Office of Science. The science community that participates in DOE’s large collaborations and facilities is almost equally split between SC labs and universities, and has a significant international component. Very large international (non-US) facilities (e.g., the LHC particle accelerator at CERN in Switzerland and the ITER experimental fusion reactor being built in France) and international collaborators participating in US based experiments are now also a key element of SC science, requiring the movement of massive amounts of data between the SC labs and these international facilities and collaborators. Distributed computing and storage systems for data analysis, simulations, instrument operation, etc., are becoming common.

Bandwidth guarantees are needed, for example, by:

- on-line analysis that involves remote elements – users, storage, experiment components, etc., which involve time constraints such as the experiment-data analysis-science-adjust experiment cycle of magnetic fusion [4];
- distributed workflow systems such as those used by high-energy physics data analysis and the need to process a certain amount of data hour-by-hour or day-by-day so that the data is not lost due to overwhelming parts of the system that will never be able to catch up (such as is the case with the LHC data processing);
- workflow systems where the inability of one element (computer) to adequately communicate data to another will ripple through the entire workflow environment, slowing down other participating systems as they wait for required intermediate results, thus reducing the overall effectiveness of the entire system.

Traffic isolation is required because today’s primary transport mechanism – TCP – is not ideal for transporting large amounts of data across large (e.g., intercontinental) distances. There are protocols better suited to this task, but these protocols are not compatible with the fair-sharing of TCP transport in a best-effort network, and are thus typically penalized by the network in ways

that reduce their effectiveness. A service that can isolate the bulk data transport protocols from best-effort traffic is needed to address this problem.

In order for these large-scale collaborative science projects to succeed, networks must provide such communication capabilities in a service-oriented way – i.e. so that they are configurable, schedulable, predictable, and reliable.

2.2 End-to-End Monitoring

In order to build a large-scale, widely distributed system that must operate reliably in order to perform complex data analysis or computational simulation tasks, the system must be able to learn, in real-time, about unexpected changes in the state of the communication between all of its components. The network must provide information that is sufficiently specific and timely that applications can adapt their behavior to reduce the impact of the outage. Without such information the human users or operators are left trying to intuit what has gone wrong with some components of the system by debugging some problem with the overall system that may well be nothing more than an indirect manifestation of an unreported communication failure in some very different part of the system, or, indeed, a problem in the application's use of the network. With a reliable and cogent notification service that describes the state of the inter-application component communications, an application can be designed to use that information about the network to adapt its behavior to the changed circumstances, or at least to fail gracefully.

An essential change in network services over the next five years will be to provide reliable, comprehensive, timely, and interpretable information about the state of the network in ways that can be meaningfully interpreted and used by user-level applications. Monitoring information is also needed to set expectations, for planning, and to set and audit formal or informal service level agreements.

This ability to report must, of course, be accompanied by a corresponding capability in the middleware and application systems to be able to accept the communication services monitoring results and do something intelligent with those results: adapt the functioning of the system to the changed / diminished communication service capability, graceful shutdown of the system, notify the user what is happening (in terms that are useful to the users involved), etc. That is, the reporting must be with information, and in a context, that has meaning to the user view of the network.

2.3 Federated Trust

The cross-site and international nature of DOE Office of Science collaborations demands a well managed, scalable, flexible, and federated approach to authentication and authorization. Problems in the areas of federated trust are heavily affected by the tension between usability and security features. It is not the case that a more secure system must be less usable, or that an easy-to-use system is an insecure one. However, additional research and development is often required to overcome usability issues introduced by security, and security issues created when a service becomes easier to use.

3 Research Topics

To address the needs of guaranteed network bandwidth and end-to-end monitoring, a number of research and development topics were identified.

3.1 Virtual Circuits

There are a number of topics that need to be addressed to achieve the goal of guaranteed network bandwidth through the use of virtual circuits. These include:

- control plane issues
- management plane issues
- data plane issues

A number of research projects have begun to address some of these issues. These include the DOE funded OSCARS[16] and UltraScience Net [19] projects, the NSF funded CHEETAH [21] and DRAGON [23] projects, Internet2's BRUW [20] and HOPI [17] projects, CANARIE's UCLP [22] project, and GEANT's BoD(SA3) [18] activities. But much work remains to be done before multi-domain virtual circuits are ready for production.

3.1.1 Control Plane Issues

A critical enabling technology to realize this vision is a control plane which allows for the provisioning of services in a hybrid network: a multi-service, multi-layer, multi-domain, multi-vendor environment. The multi-service aspect refers to the capability to provide a variety of connection modalities such as Ethernet, SONET, or InfiniBand. The multi-layer aspect refers to the fact end-to-end service may be instantiated via a data plane path that traverses heterogeneous network elements that belong to different technology layers. The multi-domain aspect refers to establishing services across multiple administrative domains to provide the largest value to end users and applications. The multi-vendor aspect recognizes the diverse set of hardware and associated capabilities that are and will be deployed. It also pronounces the need for interoperability standards.

These key control plane technologies require a community wide effort to develop the needed inter-domain agreements, standards, and interoperable implementations. The focus on inter-domain communications is rooted in the realization that different networks and administrative domains will implement different network data plane and control plane technologies that best suit their situation based on factors such as performance, cost, available physical resources (such as optical fiber plants), current equipment, vendor relationships and user requirements. Another key driver of control plane interoperability across administrative boundaries is the fact that scientific data flows traverse multiple administrative domains in almost all cases. To effectively peer and interoperate such diverse networks, the key capability of the control plane is the definition of inter-domain communications.

For example, the setting up of inter-domain guaranteed bandwidth circuits typically involves the virtual circuit extending across five to seven autonomous networks: the lab/campus network at each end, the lab/campus service provider (e.g. ESnet, a US RON (Regional Optical Network), and a European NREN) and the US national or pan-European transit network (e.g. ESnet, Internet2, GÉANT) or SINet (Japan). Differences in network infrastructure (e.g. hardware, link capacity, etc.) must be addressed at the inter-domain boundary in order to provide consistent service characteristics (e.g. bandwidth, delay, and jitter) across domains, as must the issues of different policies, such as Acceptable Use Policies (AUPs), Service Level Agreements (SLAs) and security requirements. None-the-less, inter-domain circuits are essential, especially between ESnet, Internet2, and GÉANT.

Topology Schemas and Exchange

An exportable representation of the topology of the applicable underlying network is needed in order to perform inter-domain setup of virtual circuits. If multi-layered provisioning services are

supported (i.e. layer 1, 2, 3 circuits), the information in the topology may vary substantially. For example, IP addresses are necessary for a layer 3 service, but not for layer 2.

For traffic engineering of multi-domain end-to-end circuits, inter-domain exchange of topology information must occur. With the heterogeneity of the participating networks, a common topology schema and distribution mechanism must be designed and standardized.

Advanced Path Computation and Scheduling

As network topologies increase in complexity, path computation for traffic engineering VCs becomes exponentially more difficult. As such, various path semantics must be taken into account to generate valid and functional paths and also to optimally utilize available bandwidth across the infrastructure. In addition to scheduling the individual network paths, efficient methods are needed for: (a) scheduling paths in groups to accommodate multiple remote users accessing a single site and parallel connections to support simultaneous data and control channels for computational steering and visualization applications, and (b) co-scheduling of network paths along with time-allocations on facilities such as supercomputers, visualization walls, compute clusters and scientific instruments such as APS and SNS. It is important that the underlying path computation and scheduling methods guarantee that a solution will indeed be found if one exists using bandwidths anywhere in the network and that it is semantically correct and optimal in a specified sense. Systematic approaches are needed to design, implement and test these advanced path computation and scheduling methods to take into account complex path semantics and to ensure optimal allocation and utilization.

Service Selection

Provisioning an entire lambda circuit for a 100Mbps virtual circuit reservation would not be efficient use of the network. This could be addressed by intelligently selecting the appropriate layer service (e.g. MPLS LSPs, L2VPNs, GMPLS lightpaths) based on the user requirement. Extending this capability to bridge multiple administrative domains would require the exchange of service capabilities.

Secure Control Plane Exchange

The secure exchange of control plane messages between domains remains an unresolved issue. This is complicated by the differences in the underlying network structures of the domains participating in the exchange. Some networks may use MPLS, others GMPLS, and yet others TLI via a proxy host. Collaborative research in this area is necessary to ensure a seamless control plane message flow for end-to-end virtual circuit provisioning.

Circuit Reliability

Network outages, scheduled or unscheduled, can cause disruption to existing circuits. In some cases a simple reroute may resolve the issue if there is availability on the alternate path. For advance reservations, this can add significant complexity to the management of provisioned paths. For example, if an active circuit is disrupted due to an unscheduled network outage, an alternate path may have the bandwidth capacity to service the reroute, but only until a pre-existing reservation on that link becomes active and consumes the entire bandwidth. At this point a decision must be made to either displace the rerouted circuit, or deny/reduce the soon to be active reservation accordingly. Mechanisms for automatic, policy-based decision making must be developed.

3.1.2 Management Plane Issues

Authentication and Authorization Infrastructure (AAI)

To prevent inappropriate provisioning and usage of scarce resources, such as bandwidth on shared networks, proper mechanisms to enforce authentication and authorization polices must be in place. For inter-domain provisioning, the AAIs must be compatible.

Service Level Agreements (SLA)

An end-to-end multi-domain virtual circuit requires a consistent level of service across all involved administrative domains. This highlights the need to investigate:

- What would the SLA look like (e.g. bandwidth, latency, jitter bounds, resilience/redundancy)
- How would the SLA be enforced/verified

Charging for Circuits

In the event that there is costing involved in the provisioning of the circuits, how are the charges recharged back to the user/program making the reservation.

3.1.3 Data Plane Issues

Data Plane Exchanges at Domain Boundaries

To support multi-layer circuit services, the data plane handoffs at domain boundaries need to be coordinated (e.g. handing off an IP packet when a VLAN tagged packet is expected would not work). In addition, even within the same layer service, some coordination is required (e.g. negotiate the correct color lambda, or VLAN tag ID).

Vertically Integrated Multi-Layered Circuits

As services are provided at higher network layers, additional overheads are added to the transport layer, (e.g. IP (layer 3) over Ethernet (layer 2) over GFP (layer 1)). A potential method of making end-to-end virtual circuits more efficient would be to use the lowest network layer transport required. This could be applied to multi-domain circuits whereby the data plane exchange between domains would be the lowest network layer transport that met the circuit requirements.

Currently, packet switched networks are built using router-based infrastructure, and switched networks are built using Ethernet, SONET and optical switches typically to provide dedicated connections. In general, the routed infrastructure is more expensive to build but is more flexible in that dedicated connections can be setup using MPLS-tagging. On the other hand, switched networks are less expensive but mainly provide a limited number of dedicated connections with limited routing capability. While several switched connections have been demonstrated to peer smoothly with routed networks, these efforts have been at the demonstration stages. Research efforts are needed to provide cost-optimal design of routed-switched networks that maximally utilize the lower layered infrastructure while meeting the requirements of routing. Since dedicated connections are somewhat limited to a small number of large-scale flows, one option is to operate the infrastructure in the routed-mode by default, and peel out the dedicated connections as needed and put them back after use. Research efforts are needed for designing both data-plane and control-plane technologies for and optimal operation of such routed-switched connections. In addition, the stability of routing in such networks must be ensured while the circuits are

provisioned in and out of the infrastructure; for example, the topology as seen by the routed infrastructure must remain constant to prevent link state computation churn.

Securing Multi-domain Circuits

Circuits are typically perceived to be secure due to their end-to-end nature. Securing circuits within a domain is reasonably well understood. However once a circuit extends beyond the bounds of a single domain (i.e. traverses a DMZ), it is less clear how the overall security of the circuit can be evaluated (i.e. common security model).

Summary

There are a number issues that must be solved to enable the production use of virtual circuits. The control plane must allow for the provisioning of services in a multi-service, multi-layer, multi-domain, multi-vendor environment. Control plane issues include topology schema design and exchange for inter-domain setup of virtual circuits, finding and scheduling optical network paths, intelligent service selection, secure control plane exchange, and mechanisms for gracefully handling circuit outages. The management plane must support an authentication and authorization infrastructure that includes support for service level agreements and a circuit charge model. The data plane should support coordinated data plane handoffs at domain boundaries, efficient end-to-end virtual circuits using the lowest common network transport layer, and secure multi-domain circuits.

3.2 Dynamic Wave Management

ESnet4 operates on an optical infrastructure shared with Internet2. The majority of the optical circuits are used individually by each network, however some are expected to be used jointly by the two networks – that is, ESnet will have optical circuits dedicated to its use, Internet2 will have optical circuits dedicated to its use, and there will be some circuits that are shared in order to support the new dynamic provisioning services.

However, dynamic joint management of a collection of waves that are designated for this purpose will be an important capability for

- traffic engineering – the automatic provisioning of additional capacity on heavily (perhaps transiently heavily) used paths
- managing scheduled 10G circuits where entire waves will form the links in a circuit path

In order to accomplish the R&D needed to build and test the control plane for dynamic wave management it is essential to have a testbed environment where the R&D can be performed on the DWDM equipment in a way that is guaranteed to be non-interfering with the production waves on the optical network.

3.3 Automatic Network Management

R&E networks in general, and ESnet in particular, have become much more complex in recent years, and this trend will continue into the future. Better tools for network management will be essential in order for these networks to scale without greatly increasing the operations staffing levels.

Network management systems should be able to anticipate problems based on a continual analysis of monitoring data in order to detect potential problems before they happen. The goal of such a system would be to anticipate problems based on knowledge-based analysis and projections of network state. For example, if packet rate counters for a router

interface were steadily – but irregularly enough to escape manual inspection – growing over time, an automated analysis system would predict that the router would become saturated and warn the engineering staff. Another example is the case where a relatively low bandwidth interface, say with a commercial peering partner, were to rapidly saturate and thus disrupt traffic to that peer. This circumstance should be automatically detected in order to allow for defensive action such as automatically rerouting traffic to an alternate path in order to avoid problems.

Analysis across many geographically diverse interfaces could also be used to detect coordinated stealth attacks against the network, against a group of sites, or against a particular network application; allowing cybersecurity actions to be initiated.

Such an automated analysis and management capability would allow network engineers to specify ‘what if’ scenarios, and then specify what to do if such a situation starts to develop. Even if this did not result in the “permanent” or production fix of the problem, it would give engineers breathing room to better analyze the situation in detail and design a long-term solution. This type of system could also be used to detect paths with particularly large number of high-speed flows where a dedicated wavelength could be used to better manage the network traffic. In the long term this could be integrated with the virtual circuit management system to automatically move such traffic off of the IP network and into the circuit infrastructure.

In order to create such an automated system several issues must be addressed. Improved networking monitoring data and monitoring data archives such that described in section XX above will be required; signal/anomaly detection techniques will have to be identified or developed; ontologies will have to be developed to describe the semantics of network functionality; rule-based systems will need to be interfaced with the signal/anomaly detection system in order to generate knowledge-base actions; the consequential actions will have to be passed to systems that interact with the network control plane to make changes in the network configuration; and so forth.

3.4 End-to-End Monitoring Services

Currently there is no standardized way to determine what performance levels the network is capable of delivering, what portions of the network are up and working correctly and what parts are broken or down for maintenance, that works across multiple domains. This leads to problems distinguishing between network issues, problems in the applications, problems with the protocols, and congestion caused by other network users.

Network measurement and monitoring services must be developed in order to facilitate effective use of advanced networks by complex distributed applications. These services should allow users, applications and other network middleware to determine realistic performance expectations, document the services that are delivered, verify that the capabilities committed are actually provided, and easily discriminate between network problems and limitations, and application problems. These services need to be provided in a seamless fashion to support paths that cross many different domains in support of globally distributed applications.

End-to-end monitoring services include the following R&D areas:

- development of a measurement framework
- end-to-end monitoring
- integration with Grid Middleware
- monitoring tools for virtual circuits

Each of these areas is described below.

3.4.1 Measurement Framework

To match the needs of worldwide collaborations and Grid computing a new generation of network monitoring is required. The initial required step is to research and build robust middleware for a sustainable, multi-domain, deployable, infrastructure with community ownership that enables both *measurement points* and *measurement archives*.

Measurement points provide both passive (e.g. Netflow, Simple Network Monitoring Protocol (SNMP) data from routers, switches etc.) and active end-to-end measurements (e.g. delays, loss, jitter, reachability, routes, throughput etc.). Measurement archives and measurement points provide cross-domain services (e.g. access to measured data, on demand measurements) that are registered, published and can easily be located by analysis services. All transactions between measurement points, measurement archives, and analysis services need to use standard protocols and published application interfaces and standard authentication and authorization methods to protect resources.

Therefore it is critical to establish a standardized framework including well-defined interfaces, protocols and data schemas. An instance of such a middleware is the PerfSONAR project. PerfSONAR is actively being pursued by European National Research and Education Networks (NRENs) and end sites, and R&E networks and sites in the US and Brazil. The PerfSONAR collaboration is starting to produce results in this area but there is still significant work to be accomplished in defining the standards, developing the applications and systems that will implement them, and getting the system deployed on a global scale.

A measurement and monitoring infrastructure such as PerfSONAR will enable a number of important new services. For example, such a network measurement system can be used to provide believable forecasting of future network performance with confidence estimates and easy access for distributed applications. This is needed for applications needing to choose between multiple instances of a service, and for aiding in detecting “anomalous” events where for example, performance has changed significantly for a sustained interval. This service will reduce the need to manually look at hundreds or thousands of performance time series to detect problems, by providing automatic anomalous event detection with low false positives and high detection probability.

Such a system will provide improved event diagnosis by enabling the end user/network administrator to query devices in multiple Autonomous System Domains (ASDs) without human intervention, thus dramatically reducing event diagnosis time. It will also allow analysis of the data in anomalous events to be filtered and provide useful diagnosis from multiple related sources such as routers along a suspected path, information from other end-to-end measurement paths with sub-paths shared with a suspect path, multiple metrics, on-demand measurements to attempt to divide and limit the scope of the problem, and related measurements of affected end hosts. Finally, it can be used to provide a flexible, reliable, alert mechanism.

Beyond the basic PerfSONAR framework there are additional research topics. Research will be needed to explore, evaluate and validate measurement tools that will work above 10Gbits/s. A *topology discovery system* that is deployed in as many administrative domains as possible is needed to provide a standard view of physical layer 1 links, layer 2 link protocols and IP layer 3 network topology. A topology service is necessary in order correlate information between the static and dynamic circuit services and the measurement infrastructure. For example, the dynamic circuit services may utilize a summarized topology database for the scheduling and reservations. Measurement and monitoring services will need the topology to know what

router/switches to interrogate the results from in order to diagnose problems, to know what is in common between different pairs of end hosts, to find bottlenecks or performance problems, and so on. Creation of such a topology service will involve designing new tools and integration of existing topology database(s) and/or discovery methods into a consistent framework.

3.4.2 User Initiated Debugging of End-to-End Service

Network problems can often be intermittent or remote (in parts of the network that are seemingly unrelated to the user-perceived path) as seen by the end users. This will be frustrating for the user when network operations report that they see do not report any problems, only for the symptoms to occur again. This can be exacerbated if the end-to-end path transits multiple network domains. Most end host systems do not have the operating systems and network stack tuning mechanisms necessary to distinguish between end host problems and network problems. In addition they typically do not have robust tool sets installed and are not setup to perform automated or remotely driven tests. A user initiated debugging service that provides intelligent information is needed to isolate the fault and directing the problem report efficiently.

3.4.3 Middleware Extensions for Network Monitoring

Once a network measurement framework exists, tools are needed to help the currently used middleware services such a Storage Resource Manager (SRM) (<http://sdm.lbl.gov/srm-wg/>) use this monitoring data for troubleshooting, planning, and optimization. Tools are needed to help determine when a virtual circuit should be requested, or when a best-effort service is adequate. While this work would mainly be done by the middleware developers, and not by ESnet, developers of the monitoring framework need to work closely with middleware developers to design APIs that are useful.

3.4.4 Virtual Circuit Monitoring

Connections on next generation networks may be provisioned using different technologies and their hybrid combinations, such as default IP connections, switched connections on layer-2 networks as in SDN, layer-3 MPLS tunnels as in ESnet, and their hybrid concatenations. An understanding of the performances over these connections in terms of throughputs and jitter properties is essential in assessing if certain applications can be optimally supported, such as monitoring and steering of a computation on a supercomputer. For example, high jitter connections are not suitable for on-line remote control of sensitive instruments, but have very little effect on bulk data transfers. Such performance assessment requires a careful design of experiments and sound statistical analysis of the measurements; the approach of running standard tools and tabulating measurements by itself is not likely to provide the deep insights needed to support high-performance applications. While the collection of usual measurements such as *iperf* and *tcpmon* is essential, it would only constitute a small portion of such performance assessment.

Performing end-to-end host measurements in a structured fashion on a virtual circuit can be difficult. Maintaining the security assertions provided on circuits may require that the network operators cannot inject measurement traffic into the circuit without requiring additional functionality from the end systems.

Isolating measurements to specific segments of the circuit (e.g. a single domain in a multi-domain circuit) is also a hard problem. The variations in network devices (e.g. what counters are available for a specific platform), and measurement methodologies (e.g. polling intervals) may result in incompatible results between domains.

Determining that a circuit is down can be easily done by hosts at either end of the circuit. However determining the specific segment or link that caused the outage is significantly more complex. In an IP (or layer 3) network, a traceroute (which can be done by the end user) can generally provide some indication of where the failure occurred. This is often not the case for circuits. For example, a multi-domain layer 2 circuit has no notion of IP addresses and therefore a traceroute would be ineffective.

Summary

We need a standardized way to determine what performance levels the network is capable of delivering, what portions of the network are up and working correctly and what parts are broken or down for maintenance, that works across multiple domains. This is needed to distinguish between network issues, problems in the applications, problems with the protocols, and congestion caused by other network users. As advanced services such as virtual circuits are deployed this problem will only get harder.

Network measurement and monitoring services, including a network topology discovery service, must be developed in order to facilitate effective use of advanced networks by complex distributed applications. Further research and development is needed in the areas of the deployment of a network measurement framework, improved end-to-end monitoring tools, better integration with Grid Middleware, and tools for monitoring virtual circuits.

3.5 General Network Issues

A few other general networking topics came up at the workshop that did not fall into the categories above.

3.5.1 Reliable Multicast

Multicast by its very design relies exclusively on User Datagram Protocol (UDP), which is lighter weight and often faster than Transmission Control Protocol (TCP), but is not connection oriented or acknowledge based. Multicast is forced to use UDP for reasons of scalability where many receivers could overwhelm a single source with acknowledgement packets or cause synchronization issues while the source waits for each listener to respond, as TCP does. Another complication when using multicast is that it can be viewed as a network application that actually changes routing state in the network based on what the source and sink nodes are doing at a particular instance in time. Beacon tools attempt to address this dynamic state issue by maintaining a persistent session between participants in an effort to create and maintain multicast routing state that can be observed for troubleshooting purposes without having to pull in end system users each and every time the multicast configuration needs to be verified.

Reliable multicast flows at line rate are problematic not only because they are UDP. Since the network routes copy each packet to multiple outbound interfaces they end up taking paths that in all likelihood have widely varying line characteristics. This is especially true in for inter-domain multicast paths.

The current focus of emerging multicast protocols is primarily Single Source Multicast (SSM) based and the only option available for IPv6. SSM is directed at solving the problem of distributing broadcast media over cable television networks or the Internet. The SSM approach differs from the Any Source Multicast (ASM) architecture employed by the Access Grid conferencing system where many sources participate simultaneously in remote conferencing.

An actively maintained beacon application and diligence at the participating sites and intermediate networks is essential to reliable multicast transport.

3.5.2 Routing Issues

Network routing protocols today are optimized to make next-hop routing decisions. They are not designed for end-to-end path computation. A standardized method by which a site can send information about which of several paths should be used for traffic to the site would be useful. This would allow, for example, an experimental cluster to signal a preferred path different than another portion of a campus.

3.5.3 Defining Common Services

Fundamental to the provisioning of dynamic allocation of end-to-end data paths across dedicated network resources is the communication between the network service consumer and providers. Application experts, end system experts, and network experts tend to use different terminology to describe service requirements thus making it difficult to translate into a supportable network service. Through a common services definition, it is envisioned that resources can be configured to meet service expectations. Such a common service definition would clearly describe the capabilities and performance characteristics of the available services.

The goal of a common service definition is to enable the creation of circuits by assembling resources through multiple administrative domains and across different network technologies to actually create an end-to-end path that is predictable, repeatable, and consistent with the users' expectations. A request from an application should be both simple and clear, so that both at the intra-domain and inter-domain level all interested parties can validate whether they can or cannot fulfill the request.

It is expected that different administrative domains will provide different services. It often unclear to the requesting user or application which services can be combined into an end-to-end service. For instance a lightpath delivered over a next-gen SONET/SDH section across one network may be linked to an Ethernet VLAN service or a routed IP service on another network. If services can be linked together into an end-to-end service, it may not be immediately clear what the constraints of the resulting service are. For instance, if next-generation SONET/SDH-based lightpaths in separate domains are linked through an Ethernet VLAN Service through a third domain, the resulting service may not have the properties the requestor expects. Such a constraint should be clear to the requesting process, which can then determine if the resulting service will meet its needs. Provisioning guaranteed bandwidth circuits require that each provider in the path address differences in network infrastructure in order to provide consistent service characteristics.

Advanced applications require a broad portfolio of network services be incorporated into the managed resource environment of large-scale science experiments. The vision for hybrid network services is to formally define network characteristics and to allow applications the flexibility to dynamically acquire and integrate required services into the applications environment.

In order for applications to request a path through a hybrid network a common services definition needs to clearly and formally define all aspects of the service. Following is a partial list of the parameters necessary to fully specify service characteristics for a circuit/service: end points definition, maximum sustainable capacity, MTU, protocol framing, connection profile, maximum latency, jitter, loss, sensitivity to packet reordering. The multi-service aspect of the emerging hybrid network refers to the ability to provide a variety of connection modalities such as: basic Ethernet service, Ethernet VLAN service, Ethernet LAN service, basic SONET/SDH service,

basic FiberChannel service. A network service can be described at various levels of abstraction, with high detail at the low abstraction level and low detail at the high abstraction layer. Tools are needed for generating the descriptions at the right level of detail for the intended consumer.

Summary

Other general networking issues that need further research and development include better support for IP multicast, exploration of network routing protocols that are based on end-to-end path computation instead of next hop information, and the development of a common services definition that would make it easier for application experts, end system experts, and network experts to clearly describe service requirements necessary to deploy a particular network service.

3.6 PKI and Federated Trust Issues

The cross-site and international nature of DOE Office of Science collaborations demands a well managed, scalable, flexible, and federated approach to authentication, authorization, and the creation of virtual organizations to manage the collaboration resources. Current approaches are ad hoc and impose a high overhead on the scientists – and security professionals. What is needed is a system for cross-site authentication and attribute-based authorization to allow sites and communities involved in cross-site collaboration to meet the policy needs of federal government computing sites, without imposing unmanageable procedures on users of the resources. The goal of such a system would be to build on federation of identity, by allowing a user's identity to be initially sourced from a trusted organization, such as a university IT department or DOE laboratory, which performs authentication and vetting of the user. ESnet is a logical trusted third party to manage such a system.

There are a number of R&D topics in this area, which are described below.

3.6.1 PKI Credential Ease of Use

X.509 credentials, PGP keys, SSH keys - these different flavors of public key credentials (and PKI) support distributed computing in science in many ways. PKI credentials have proved very difficult for non-experts to manage. During the early days of PKI software development various kinds of human factor problems were overlooked, and implicit assumptions about how to secure and manage the PKI key pairs were made, that fail in the mass market. The split key algorithms in PKI are powerful and useful from a security and software engineering standpoint, but are non-intuitive and the qualities associated with each key are not understood by the typical user (and our typical user is a sophisticated scientist or engineer with an advanced technical degree and considerable technical experience in his or her field).

Key lifecycle - acquisition, management, and replacement - of PKI components must be improved considerably to make it more feasible for deployment and more practical to manage. At the same time, most aspects of PKI are unnecessarily part of the user experience.

3.6.2 PKI Credential Security

For users the focus has been on "software crypto-stores" i.e. keys stored as computer files. This allows a large number of attacks, but perhaps more importantly plants doubt in the minds of resource owners who are exposed to risk of financial loss and embarrassment when resources are used by illegitimate parties. They have little trust in these file-based stores because they know how weakly defended their own resources are, perhaps.

R&D is needed to develop secure credential stores - places where users could store & retrieve their PKI credentials, minimizing exposure to weak resources unsuited for software crypto stores. We also need R&D to integrate personal hardware tokens (smart cards, USB tokens, PIV-2 cards, &c) into our application and security services.

For hosts the factors that make resources a poor environment for software crypto-stores make hosts very vulnerable to identity theft. Typically, host keys are left in an unsecured (unencrypted) state on the host's file system, so that the host services that use these keys can be stopped and started or rebooted without any further intervention.

In a research environment, the compromise of a host key employed for server-side SSL or similar is typically not a very interesting attack. It is a side effect of a compromise of the whole host, which is itself a catastrophe. But the attacker is usually not able to do anything particularly interesting with the host credential. In a multi-tiered service environment, or in an environment where hosts are acting as proxies for abstract services, the compromise of a host key can be disastrous. Sophisticated token management and security techniques for hosts will help us overcome this problem.

3.6.3 Trust federation issues

Various kinds of identity, authentication, and authorization 'federations' have been developed and/or are essential in multi institution science environments (the norm today). In general these federations need to integrate pre-existing services and moderate or map between them without affecting the operation of those services; in practice they all have significant pushback and interoperability issues with existing services.

Grid Authentication/Authorization Portal

Authentication and authorization issues in Grids for DOE science must be scoped in such a way that a simple credential management portal for scientists can be provided. Current efforts are focused on integrating a federation technology (Shibboleth) with an X.509 certification authority and a Grid trusted authorization service (VOMS), to produce short-term certificates for Grid users. This is a requirement for grids, producing a "just in time" certificate with just the right authorization attributes, with a short lifespan, limiting the risk to long term credentials and user identity.

1000-Host Certification Problem

The certificate infrastructure needs to provide certificates to blocks of hundreds, perhaps growing to thousands, of hosts at various sites. These blocks represent clusters of machines, or sometimes simply collections that are managed in a uniform manner. Both the certification authorities issuing these certificates, and the Grid or VO management services that oversee them lack the tools to manage "herds" of hosts or services.

Credential Validation Services

The various types of credentials (particularly the PKI-based ones) need to be validated by the resource owner. Validation starts with a simple evaluation of the cryptographic integrity of the credential, followed by an evaluation of the trust anchors supporting the credential, a check for revocation, and a check for suitability of the credential for the intended use (e.g. my library card is valid, but not useful for paying my registration fee). The current arrangement involves shipping files from trust anchor to every possible resource provider and relying party, which is hopelessly unscalable. Services are needed that can be deployed globally, deployed by virtual

organizations, sites, and even locally, in tiers, in order to provide a robust evaluation and validation infrastructure. Validation services must also support the X.509 bridge CA architecture, in order to integrate this complex X.509 feature with current limited client-side software.

Provisioning

Complementary to Credential Validation described above, a robust provisioning system is needed to move the necessary trust anchors and related policy objects to the various consumers requiring them, e.g. credential validation services, relying parties, user browsers and operating systems, and perhaps other targets. Some initial work on this has been done in some Grid projects, but considerable work needs to be done to expand the coverage and scope of these services.

Abstract Services

Abstract services are services that are instantiated on multiple individual hosts. X.509 certification authorities and VOMS attribute authorities are examples of abstract services. There are many others, including network measurement services, SAM data movers, and DNS. Current technology forces the identification of individual hosts and services, but this is not useful. Applications (or users) need to be able to find the abstract service from a trusted network information service (itself an abstract service), and then query it or update it. In general, in the complex workflows found in large-scale experiments, there are going to be chains of multi-tiered applications, supporting agents of multiple abstract services, interacting with each other and supporting delegation of rights from various other entities.

Users need to be able to trust they are talking to a legitimate instance of the service, and that they have found the service they want and not a fake service. This has to be extended to cover all the cases implied by multi-tiered and multi-agent workflows as well. Extensive development of trust mechanisms are needed for digital signatures and appropriate deployment strategies.

Shibboleth vs. VOMS

Identity federation technology such as Shibboleth, VOMS, and perhaps some other technologies such as Liberty have developed in separate domains in pursuit of similar but not identical security goals. The inability to interoperate between these different authentication models represents an obstacle to high-level federation. Techniques to interoperate must be developed; perhaps in some cases this can evolve to convergence.

X.509 Bridge Architecture

Extensive software development is required to provide the components of X.509 CA bridge architecture to PKI in use in science. This includes extensive work to the underlying PKI software, client robustness, validation services, and network directory infrastructure.

Levels of Assurance (LoA)

Extensive work developing LoA has been done in NIST and similar agencies. It is probably useful to develop standards appropriate for scientific use, and either adopt some mapping from NIST, or create appropriate designations of LoA for our use. NIST has provided a model for evaluating credentials associated with human identity; it currently has nothing for hosts or services.

One Time Password (OTP) Credentials

One time password (OPT) credentials are useful in an environment that is very untrustworthy. While OTP cannot prevent “man in the middle” attacks, it still helps prevent identity theft a great deal. Currently many users complain about the large number of OPT tokens they must carry in order to log into various recourses across multiple sites. Research and development effort is needed to federate OPT tokens and integrate OTP into other authentication services.

Summary

The cross-site and international nature of DOE Office of Science collaborations demands a well managed, scalable, flexible, and federated approach to authentication, authorization, and the creation of virtual organizations to manage the collaboration resources. Current approaches to this problem are ad hoc, difficult to use, and easy to mis-configure. Research and development topics in this area include making PKI easier to use, development of secure ‘crypto-stores’ for user PKI credentials, and a number of trust federation issues.

3.7 Network Security Issues:

The following security related R&D topics need to be addressed.

3.7.1 Multi-site view of network probing/intrusions

Examining traffic data from a single point in the network is typically ineffective in determining a diversely sourced attack against ESnet infrastructure. To circumvent this, a distributed, intelligent network of intrusion detectors could monitor the traffic at key network ingress points. The information about wide-spread, multi-site intrusion attempts detected by the IDS system could then be used to defend ESnet network infrastructure.

Some particularly significant forms of attacker activity can be much better identified when analysis can draw upon a view that spans multiple vantage points. Correlating connection attempt information across multiple sites allows for much more interesting forms of analysis and attack detection. A network provider of the size and scope of ESnet is in a unique position to notice new types of large-scale probes or attacks, and potentially block the offending traffic. Such connection attempt data would also be extremely valuable for forensics and incident response.

3.7.2 Separation of network control and data planes

By separating the network control and data planes, a network removes security risks associated with the interaction of the traffic carried by the network and the operational state of the network. Most of the large commercial carriers have separated the control and data planes of their networks – this means that it is impossible for traffic carried by the network to change the control state of the network. The one exception to this is that the network can still be overwhelmed by traffic volumes larger than its capacity (denial of service attack) but even DoS attacks are limited to resource exhaustion since they cannot target the control or management interfaces directly. Separating the control and data planes presents unique challenges to research and education networks, since R&E networks provide enhanced services such as virtual circuits and IP multicast where this separation is considerably more difficult.

3.7.3 DNSSEC

DNSSEC is being mandated by the Office of Management of Budget (OMB) for use on many government systems in 2008 and is expected to expand to all government systems in the future.

Research is needed into the impact that a large number of systems requesting validation from large and intermediate sized server will be. The validation process is compute intensive and it is not well understood how it will scale. Further, the current intent is to only have validation done by name servers, but it is reasonable to expect this to expand to stub resolvers to provide complete, end-to-end validation of the DNS name resolution and this may severely impact site and facility name servers

Summary

There are a number of network security-related R&D topics that should be addressed, including detection of distributed, coordinated attacks, separation of network control and data planes, and deployment and testing of DNSSEC.

3.8 Collaboration services

ESnet currently provides collaboration support via its telephone and video conferencing services. The following were identified as areas where additional services are needed.

Cross-domain calendaring

Organizing a meeting, multi-media conference, or similar event with participants from various institutions is very challenging. Most institutions have enterprise calendaring, and some public

online service providers like Google and Yahoo provide some level of calendaring service, but utility and coordination between services is lacking. A standard exists to support the basic interchange of calendaring information (see <http://en.wikipedia.org/wiki/ICalendar>). What is needed is development of appropriate middleware to support interchange between different calendaring services, security solutions, and interaction with appropriate white pages services.

H.350 - Directory Services Architecture for Multimedia Conferencing

The H.350 recommendation describes a directory services architecture for multimedia conferencing using LDAP. Standardized directory services can support the association of persons with endpoints, searchable white pages, and clickable dialing. Directory services can also assist in the configuration of endpoints, and user authentication based on authoritative data sources.

The use of a common, authoritative data source for call server, endpoint, user, authentication and white pages information is an important aspect of large-scale multimedia conferencing environments. By standardizing the LDAP schema used to represent the underlying data, products from different system vendors can be deployed together to create an overall application environment. Effort is needed to develop the necessary schemas and for the deployment of the LDAP server.

Summary

Additional development effort cross-domain calendaring and a directory services architecture for multimedia conferencing are needed to improve ESnet's suite of collaboration services.

3.9 Other Topics

A couple other miscellaneous topics were identified at the workshop, and are described in this section.

3.9.1 Effective Connections for Leadership Computing Facilities

Simply connecting NICs on supercomputing systems to networks is unlikely to yield optimal throughput, as evidenced by 5 Mbps TCP throughputs when the ORNL Cray X1(E) is connected to a dedicated 1Gbps circuit. Systematic methods are needed to investigate the cross-connects between supercomputers and network connections to ensure effective throughputs. This work typically involves designing interconnection systems that convert the native data framing (such as Infiniband or Fiberchannel) of these systems to those of wide-area networks such as Ethernet or SONET. Furthermore, research efforts are needed to ensure that the I/O stacks on the end hosts are properly aligned and tuned to ensure high network throughputs. In another direction, more work is needed to investigate the effectiveness of non-standard wide-area technologies, such as Infiniband over Ethernet or SONET, Fiberchannel over IP and other data-movement technologies. Efficient data movement will be the defining characteristic of the Science Data Network and that data will reside on storage at one or both ends. Efficient coupling of that storage to the network will be an absolute requirement for completing DOE's mission. In addition to storage and file transfers, effective technologies for memory transfers would be needed to support computational monitoring and steering as well as interactive remote visualizations.

3.9.2 Automatic Test Suites for Protocol and Application Testing

Newer protocols and application modules are being developed to make use of the newer network connection capabilities, such as dedicated bandwidth and controlled jitter, to support a variety of tasks including file transfers, memory transfers, computational monitoring and steering and remote visualization. For example, transport methods are being developed to achieve high-utilization of dedicated connections and dynamic stability of closed-loop control flows for instrumentation and computational steering. Also, the applications to support the monitoring of on-going computations over wide-area connections and interactively steering them are being developed. The ability to test the viability, scalability and performance of these protocols and applications would be of immense value. Research efforts are needed to develop automated tools that will enable testing of these technologies through automated scripts that lessen the burden on the user in getting involved in network details. For example, effectiveness of a transport protocol can be assessed by an automated script that sets up connections of various types and lengths, collects the measurements, and carries out analysis. The design and development of such tools is currently in its infancy and would require sustained research efforts to mature them to be provided transparently to the users.

Summary

Additional research and development is needed on how to better connect DOE's Leadership Computing Facilities to ESnet, and to provide services to help test experimental network protocols on new network capabilities such as virtual circuits.

4 Conclusions

DOE Office of Science supported programs continue to become more and more distributed. These distributed, large-scale collaborations require huge amounts of bandwidth and advanced network services. Participants of this workshop agreed that the research and development topics

described in this report are all important to the success of the DOE mission. We recommend that a Network R&D program be developed, and that ESnet partner closely with a set of R&D projects to address the issues described in this report.

5 References

- [1] “High Performance Network Planning Workshop,” August 2002
<http://www.doecollaboratory.org/meetings/hpnpw>
- [2] DOE Science Networking Roadmap Meeting, June 2003 <http://www.es.net/hypertext/welcome/pr/Roadmap/index.html>
- [3] Science Requirements for ESnet Networking, <http://www.es.net/hypertext/requirements.html>
- [4] “Science-Driven Network Requirements for ESnet: Update to the 2002 Office of Science Networking Requirements Workshop Report,” February 21, 2006.
<http://www.es.net/hypertext/requirements.html>
- [5] 2007 Network Requirements Workshop for the DOE/SC BES Program Office,
<https://workshops.es.net/2007/bes-net-req/wiki/bin/view/BES-Net-Req/WebHome>
- [6] Open Science Grid – <http://www.opensciencegrid.org/>
- [7] LHC Computing Grid Project <http://lcg.web.cern.ch/LCG/>
- [8] CMS - The Compact Muon Solenoid Technical Proposal. <http://cmsdoc.cern.ch/>
- [9] The ATLAS Technical Proposal.
<http://atlasinfo.cern.ch/ATLAS/TP/NEW/HTML/tp9new/tp9.html>
- [10] The Relativistic Heavy Ion Collider at BNL. <http://www.bnl.gov/RHIC/>
- [11] <http://neutrons.ornl.gov/>
- [12] <http://www.nersc.gov/>
- [13] <http://www.ccs.ornl.gov/nlcf/>
- [14] <http://www.alcf.anl.gov/>
- [15] “ESnet4: Advanced Networking and Services Supporting the Science Mission of DOE’s Office of Science” 2007 SciDAC review report at www.es.net, presentations and publications tab.
- [16] C. Guok, D. Robertson, M. Thompson, J. Lee, B. Tierney and William Johnston, *Intra and Interdomain Circuit Provisioning Using the OSCARS Reservation System*, Proceedings of the GridNETS 2006 workshop, Oct. 2006, LBNL-60373
- [17] H. Boyles, *Recent Results from Internet2’s Hybrid Optical and Packet Infrastructure Project (HOPI)*, Presented at TERENA Networking Conference, Rhodes, Greece, 2004.
- [18] A. Patil, *Advance multi-domain provisioning system*, Presented at TERENA Networking Conference, Catania, Italy, 2006.
- [19] N. Rao, W.R. Wing, S.M. Carter, and Q. Wu, *Ultra-Science Net: Network testbed for large-scale science applications*, IEEE Communications Magazine, vol. 43, no. 11, 2005.
- [20] B. Riddle, *BRUW: A bandwidth reservation system to support end-userwork*, presented at the TERENA Networking Conference, Poznan, Poland, 2005.
- [21] M. Veeraraghavan, X. Zheng, H. Lee, M. Gardner, and W. Feng, *CHEETAH: Circuit-switched high-speed end-to-end transport architecture*, in Proceedings OptiComm 2003, Dallas, TX, Oct. 13-17, 2003.
- [22] J. Wu, M. Savoie, S. Campbell, H. Zhang, G.V. Bochmann, and B. St. Arnaud, *Customer-managed end-to-end lightpath provisioning*, International Journal of Network Management, vol. 15, no. 5, pp. 349-362, 2005.
- [23] X. Yang, T. Lehman, C. Tracy, J. Sobieski, S. Gong, P. Torab, and B. Jabbari, *Policy-based resource management and service provisioning in GMPLS networks*, in First IEEE Workshop on Adaptive Policy-based Management in Network Management and ContIEEE INFOCOM 2006, Barcelona, Spain, Apr. 2006.

6 List of Participants

Joe Burrencia (joeb@es.net)
Mauro Campanella (mauro.campanella@garr.it)
Charlie Catlett (catlett@mcs.anl.gov)
Les Cottrell (cottrell@slac.stanford.edu)
Cees de Laat (delaat@science.uva.nl)
Dennis Gannon (gannon@cs.indiana.edu)
Chin Guok (chin@es.net)
Dan Hitchcock (Daniel.Hitchcock@science.doe.gov)
Thomas NDousse (tndousse@er.doe.gov)
William Johnston (wej@es.net)
Tom Lehman (tlehman@isi.edu)
Joe Metzger (metzger@es.net)
Harvey Newman (Harvey.Newman@cern.ch)
Kevin Oberman (oberman@es.net)
Larry Peterson (llp@CS.Princeton.EDU)
Nagi Rao (raons@ornl.gov)
Jerry Sobieski (jerrys@maxgigapop.net)
Rick Summerhill (rrsum@internet2.edu)
Brian Tierney (bltierney@lbl.gov)
Linda Winkler (winkler@mcs.anl.gov)
Bill Wing (wingwr@ornl.gov)