SAND2001-1345C

Systems Integration of the Supercomputing Infrastructure and User Environment among the ASCI National Laboratories

James A. Ang, Pete W. Dean, Martha J. Ernest, Arthur L. Hale, Leonard Stans
Sandia National Laboratories[†]

Jean Shuler, Dave Wiltzius
Lawrence Livermore National Laboratory[‡]

Stephen E. Turpin
Los Alamos National Laboratory[§]

ASCI DisCom^2 Objectives/Recent DisCom^2 Milestone Goal

The Accelerated Strategic Computing Initiative (ASCI) is a U.S. Department of Energy program among Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL) and Sandia National Laboratories (SNL). For several years the ASCI program has worked with the High-Performance Computing and Communications community to develop terascale computing technologies and systems to help ensure the safety and reliability of the U.S. nuclear weapons stockpile without full-scale testing. While many program elements of ASCI have focused on technology development, the impact of that work is felt through the successful integration and deployment of these technologies into the ASCI user environment. The Distance and Distributed Computing and Communications (DisCom^2) program element of ASCI has two key objectives: (1) extend the environments required to support high-end computing to remote sites, and (2) develop a complex-wide integrated supercomputing environment to support stockpile stewardship. The effort to meet these objectives is focused on providing highly capable environments for classified computing with appropriate need-to-know separations. A major challenge associated with these DisCom^2 objectives is coordinating and integrating technology deployments among the existing Tri-Lab production computing centers and maintaining their interoperability.

A recent DisCom^2 program element milestone goal was to demonstrate the ability to ensure that technologies from across the Tri-Lab ASCI program elements were deployed and integrated in a manner that allows the Tri-Lab user community to access and make effective use of the ASCI White platform *as if they were local*. The focus for this milestone is the systems integration of institutional capabilities among the three ASCI Labs to provide a seamless distance computing user environment. Users at LANL and SNL have significant ASCI-scale resources that must be seamlessly integrated with the ASCI White platform at LLNL.

† Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.
‡ This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.
§ Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

Determination of System Requirements

System requirements have been identified through a deliberate series of steps during the last few years. The first step was to develop use cases through the application of the Unified Modeling Language (UML) methodology to collect data via interviews, discussions, and surveys of the computational code teams, analysts, and the systems support personnel within the ASCI community. An initial analysis of the interview and survey data resulted in a synthesis of usage modes into 22 representative ASCI use cases. These use cases were documented in *Distance and Distributed Computing and Communication FY99 Service Model*, Beiriger, et al., SAND99-0683, April 1999. The resources and use cases for the ASCI System Infrastructure provided a focal point for more detailed surveys of 20 application code development teams who were asked to quantify the resource needs for their specific application in both routine and hero modes and in both current and future usage scenarios. An initial analysis of the ASCI resource requirements was documented in *Quantification of ASCI Application Resource Needs and Requirements*, Ernest and Hodges, SAND2000-1123, May 2000. The analysis behind these two documents, combined with detailed input from LLNL's production computing center, established the foundation for the current specification of system requirements in *Distance Computing Usage Model for the ASCI White Environment*, Wiltzius, et al., October 2000. This document is distinct in that while it reflects user requirements, the system requirements it specifies do not meet the specific needs of all ASCI users. Rather, it identifies the specific performance objectives that DisCom^2 committed to meet for the ASCI White Distance Computing Milestone. Within the constraints of these system requirements, other performance parameters were traded off (tempered by the pragmatic considerations of the production computing centers).

The *Distance Computing Usage Model for the ASCI White Environment* identified seven high-level usage model elements that characterize distance user interactions with the ASCI White platform:

1. Gather information about the environment
2. Data migration, manipulation, and storage
3. Application code debug
4. Problem setup
5. Running the application and solving the problem
6. Analysis of the results (primarily by visualization)
7. Tri-Lab coordinated operational support

Achieving this user environment depended upon the combined efforts of the ASCI Platform, Problem Solving Environment (PSE), Visualization (VIEWS), and DisCom^2 program elements. DisCom^2 made significant contributions to usage model elements 1, 2, 5, and 7 and provided tools and infrastructure to facilitate elements 3 and 6 for remote users.

System Description: Scope of Work

Table 1 describes the technologies and concepts that were utilized in the DisCom^2 milestone qualifications and that are currently available to enable usage model elements 1, 2, 5, and 7. The table also lists the test cases that were conducted in connection with each of the elements of the usage model.

Table 1. Technologies and concepts used and test cases conducted in the DisCom^2 milestone qualifications.

Usage Model Element	Technology/Concept	Test Case
Gather Information about the Environment	• World Wide Web • Distributed Computing Environment (DCE) Tri-Lab authentication • Distributed File System (DFS)	I: Tri-Lab Information Access
Data Migration, Manipulation, and Storage	• Parallel data transport • Parallel local/wide area networks (LANs/WANs) • High-speed ATM encryptors	II: Remote I/O Data Movement
Running the Application and Solving the Problem	• Globus/Grid technologies • Distributed Resource Management (DRM) • Local job submission tools	III: DRM Services IV: Remotely Controlled Capabilities
Tri-Lab Coordinated Operational Support	• Capability Release (CR) • Configuration Management Board (CMB) • Enhanced Priority Run (EPR) Committee	V: Tri-Lab Coordinated User Support Also Tests I, II, III and IV

Gather Information about the Environment

Users must understand the procedures and policies (from acquiring a login account to abiding by the operational policies) of the facility hosting and directly supporting an ASCI platform. Once users avail themselves of these resources, it is helpful for the users to remain informed of planned and unplanned down times, software and hardware updates and fixes, system usage, and system problems. Furthermore, the users benefit from a responsive and knowledgeable customer support (e.g., hotline) service.

Both classified and unclassified (where appropriate) World Wide Web servers located at the Labs provide information about the environments associated with use of the ASCI White platform. Often this information is customized for remote users. Particularly for the information available on the unclassified network, Web servers require additional user authentication via the Tri-Lab DCE for accessing certain pages.

The milestone test for this usage model element demonstrated that information was available, current, and accessible to the Tri-Lab user community. Specifically, the information identified for the milestone included the ASCI White User Manual, the *Distance Computing Usage Model for the ASCI White Environment*, *Access Protocols for ASCI "Most Capable" Systems*, the ASCI White Integration Plan, ASCI White machine status, and DisCom^2 WAN status between LLNL and SNL. The testing also demonstrated the use of the Tri-Lab DCE/DFS space as a software repository for sharing utilities and scripts.

Data Migration, Manipulation and Storage

Users of high-performance computers at the three Labs run applications that generate large data sets (for applications on the existing 3–10 teraOPS platforms, data sets on the order of 0.1–10 TB). It is common practice to archive these data sets to storage—such as the High-Performance Storage System (HPSS)—or to move all or parts of the data sets to visualization servers for postprocessing and/or analysis. In collaboration with VIEWS, new visualization tools are being investigated that may reduce the extent to which large data sets are moved off ASCI platforms in the future.

The necessity for effective movement of large data sets was addressed initially by acquiring an affordable, extensible (to 100 Gb/s) high-speed WAN interconnecting the three Labs (called the DisCom^2 WAN). Although innovative, low-cost, high-speed WAN solutions are readily available between some large metropolitan areas (e.g., Chicago, New York, Los Angeles, San Francisco), such services are not currently available in sparsely populated areas such as New Mexico. However, by understanding the user requirements and practices of the telecommunication industry, DisCom^2 identified parts of the WAN service level agreements that could be safely reduced to significantly lower the cost. As a result, DisCom^2 acquired an affordable OC-48c SONET service (2.4 Gb/s) between LLNL and SNL and between SNL and LANL.

Early work by PSE in collaboration with the Department of Defense and the NSA resulted in the development of OC-3c, OC-12c, and (with late CY01 availability) OC-48c ATM encryptors. These encryptors provide NSA Type 1 encryption, required for transmitting classified data over unclassified networks. Four OC-12c ATM encryptors are used at each site between the site router and the ATM switch connected to the DisCom^2 WAN.

DisCom^2 measurements show that utilizing the non-standard "jumbo frame" over Gigabit Ethernet results in a tripling of host-to-host (memory-to-memory) throughput (100 MB/s over a single Gigabit Ethernet interface from ASCI White to local SGI visualization server). Clearly, this feature is also highly desirable for use over the DisCom^2 WAN. Although this feature has not yet been enabled on the WAN, milestone performance targets were achieved. Further improvements in DisCom^2 WAN throughput are expected when this feature is implemented in the WAN edge routers later this year.

The data movement requirement for the Tri-Lab environment also imposes stringent requirements on data movement tools. Development of appropriate tools through a collaborative effort among DisCom^2 and the HPSS and PSE development teams resulted in the Tri-Lab deployment of a data movement tool based on File Transfer Protocol (FTP) that met the following DisCom^2 requirements:

- Clients and servers interoperate between machines at the three Labs,
- Utilizes parallel storage and network paths simultaneously to achieve higher performance,
- Clients would interoperate with HPSS servers,
- User login occurs with the Tri-Lab DCE service.

The milestone set aggressive goals considering the late schedule (February 2001) for the completion of the LLNL to SNL DisCom^2 WAN segment and the delayed availability of key technologies (e.g., Gigabit Ethernet jumbo frames) for WAN edge routers with OC-48c ATM interfaces. Using one multiprocessor login node on ASCI White and up to four Gigabit Ethernet interfaces, the milestone parallel FTP data movement rates were met. High-speed data transfers of ASCI-scale data (>200 GB) were transferred at over 100 MB/s from ASCI White to the visualization server Edison at SNL. Data was transferred from ASCI White to the HPSS at SNL at over 36 MB/s per second. Because the high-speed WAN between SNL and LANL was delayed (expected to go into operation in May 2001), smaller scale ASCI data were transferred from ASCI White to the HPSS and visualization platforms at LANL over SecureNet to demonstrate functionality rather than performance.

Running the Application and Solving the Problem

Shortly after ASCI began, the Tri-Lab user community became accustomed to accessing local machines using tools and services such as secure shell (SSH), DCE/DFS, the TotalView debugger, and FTP. Hence, a user requirement for the remote access of the Tri-Lab computers was the availability of these same tools. The milestone requirement was met by testing the availability and operation of a set of these tools when used remotely.

Each Lab also provides its users with local job submission and monitoring tools, allowing users to submit jobs to a machine and regulating how multiple users share access to that resource. Specifically, the job submission tool DPCS (Distributed Production Control System) for ASCI White must also enforce the access policies established by the Tri-Lab executives for this shared resource. The job submission and monitoring tools differ at each Lab because of different technologies supported by the different platform suppliers and because of individual site preferences. Hence, remote users must either become intimately familiar with the native services at each site or utilize new Grid services provided by DisCom^2.

The new ASCI Grid services are a significant first step towards providing a single interface for submission and monitoring on different resources. The ASCI Grid, which leverages the research and development efforts of the computational Grid community, uses a layered services approach. A modified Globus (DOE approved Kerberos version 5 replaces the public key infrastructure security protocols) provides the lowest level of services, communication, and job submission using the local resource management mechanism. This approach works side by side with non-Grid-submitted jobs and allows the local resource management mechanisms to enforce the access policies. The DRM services provide workflow management, brokering, and information services. For the top layer, a generic job submission tool, called the Production Wizard (a graphical user interface developed by ASCI DisCom^2), and Web-based monitoring have been developed and deployed as the initial grid-enabled user interfaces.

For this milestone the Production Wizard (PW) was used to remotely perform several tasks related to job submission and monitoring, such as submitting jobs to ASCI White and other machines, submitting jobs to DRM (which chose the most appropriate platform on which to run the job), monitoring the progress of the job, and collecting the job logs.

Tri-Lab Coordinated Operational Support

DisCom^2 recognized the need for coordination of implementations and maintenance of new capabilities with the organizations responsible for operations of the ASCI resources at each site. The CR initiative was created to address this need. Deployment of new features developed by DisCom^2, and other ASCI program elements that have developed relevant distance computing technologies and change management related to those features, is coordinated by this initiative. The CR facilitates communication among three communities (users, system management, and technology developers), three Labs, and across ASCI program elements to address technology integration and deployment issues.

Users desire a reasonably stable and uniform computing environment. At times, this conflicts with the necessity to provide additional capability for the user community by fixing bugs, upgrading software, deploying new or replacement systems, etc. The tension between a stable and uniform computing environment and tracking technology developments is difficult to mitigate with a single local user community. There is another dimension of difficulty when doing this between three computer facilities at three Labs with three user communities.

An acceptable solution must not unduly deter a computer facility from making important changes to the computing environment, but it should accommodate "reasonable" requests that would help minimize instability in the computing environments of the Tri-Lab user community. The focus is clearly on those operational aspects essential to the preservation of remote access. To a large extent, this can be realized by promoting extensive communication among the members of the computer facility staff.

DisCom^2 helped establish a Tri-Lab CMB to assist in this integration process and provide a mechanism to help anticipate and resolve Tri-Lab integration and configuration management issues. The CMB is made up of the ASCI production system managers from each Lab and the security officers responsible for the production computing systems (or equivalent security representatives) from each Lab. The expectation is that the members will coordinate planned changes to their local production computing environment that could raise issues or adversely impact the Tri-Lab user environment.

DisCom^2 also collaborated in the creation of a forum to address priority jobs. A classified Tri-Lab EPR videoconference is held every week to arbitrate scheduling conflicts of large jobs on the ASCI White platform. The purpose of the EPR meeting is to schedule the jobs whose operational parameters are outside the range of "routine" jobs for ASCI White.

All of the successes achieved in the testing are in some degree attributable to the good coordination that exists between the computing centers. A specific example is the test case that detailed, via real trouble tickets from January and February of this year, the effective and efficient manner with which remote users' problems and requests for assistance are handled through the joint efforts of the local and remote computer centers.

Testing a Distributed Implementation

DisCom^2 is responsible for the Tri-Lab infrastructure and features to support distance computing on the ASCI White platform. The testing will therefore demonstrate from a user perspective and from an operational standpoint that the technology has been coordinated, integrated, and deployed in the Tri-Lab environment to help ensure ASCI success.

The previously developed use cases helped define the global user requirements, and the *Distance Computing Usage Model for the ASCI White Environment* is a specific reflection of those user needs in system requirements for the Tri-Lab ASCI system infrastructure. This model is used as the structure and basis for the described features for this initial system implementation. As shown in Table 1, the elements from this model have been grouped into five test cases to depict specific areas of interest to the user. While Test Cases I, II, and V have a narrow but detailed focus, Test Cases III and IV have a broader end-to-end perspective of the computing environment. The test cases will be run from typical users' desktops at SNL and LANL, accessing ASCI resources at LLNL and (as appropriate) other local ASCI resources.

Test Case I: Tri-Lab Information Access

The main goal of Tri-Lab Information Access is to present how a Tri-Lab user could get started using the Tri-Lab environment and obtaining additional information about resources and services by accessing Web-based information. The information includes the ASCI White User Manual, the *Distance Computing Usage Model for the ASCI White Environment*, White Access and Allocations Policies from the *Access Protocols for ASCI "Most Capable" Systems*, and machine status and monitoring information. In addition, some Tri-Lab "behind the scene" type of information is available via Web and DCE/DFS, e.g., the software repository for the Tri-Lab, the CMB Meeting Notes, and the Configuration Management Plan for the Tri-Labs. This information demonstrates the extensive coordination that is required to achieve successful integration, deployment, and maintenance of the Tri-Lab user environment.

The metric associated with this test case is the ease of access to the Web- and DCE/DFS-based information. Ease of access means that credentials are accepted from the other Labs and that there are no significant delays in traversing the links or displaying the data as seen by the user.

Test Case II: Remote I/O Data Movement

The main objective of Remote I/O Data Movement is to ensure Tri-Lab coordinated, integrated, and deployed data movement capability. This includes high-speed data transfers via parallel network and application services. ASCI-scale data (>200 GB) was transferred from ASCI White to the HPSS and visualization platforms at SNL. Because the high-speed WAN between SNL and LANL is delayed, smaller scale ASCI data were transferred from ASCI White to the HPSS and visualization platforms at LANL to demonstrate functionality rather than performance. The parallel data transfer scripts used parallel versions of FTP to achieve performance over the parallel WAN between LLNL and SNL. These same scripts were also used to transfer data over SecureNet between LLNL and LANL.

Metrics gathered for this test case include data transfer rates, data integrity of transferred files, and common location of the Generic Security Service (GSS)-FTP scripts.

Test Case III: Distributed Resource Management Services

The focus of DRM Services is to ensure that remote job submission is available via PW. PW is a common interface that can be used to submit jobs to many of the ASCI resources. This common interface, along with DRM monitoring Web pages, helps manage the complexity of job submission, file transfer, and job monitoring. PW was used to submit a job locally at both SNL and LANL before submitting the job to ASCI White.

Metrics associated with this test case include using a common user interface. This means that the user-local and the user-remote job submissions both use the PW interface. Output files/logs are gathered from the user-local runs and user-remote runs for comparison.

Test Case IV: Remotely Controlled Capabilities

The focus of this test case is to ensure that a remote user who is logged into ASCI White has the same look and feel as if he or she were local. Capabilities that an LLNL user takes for granted will be demonstrated with the user logged in from SNL/NM and LANL. The demonstrated capabilities included logging into ASCI White via SSH, remote job submission via ASCI White's local job scheduler (DPCS), execution of an FTP script to transfer data to the local LLNL HPSS, remote access of codes or data via DCE/DFS, and use of a typical software development tool displayed from White back to the remote user (such as the TotalView debugger or a performance tool).

Metrics are the remote user's ability to log into ASCI White and directly submit a job via DPCS, the data rate for the transfer from ASCI White to LLNL HPSS, and the user-perceived performance of the software tool with respect to local use of the same tool. Performance is defined as comparable, acceptable, or unusable.

Test Case V: Tri-Lab Coordinated User Support

This test case demonstrates the coordination of the Tri-Lab production computing centers. It details the way remote users report problems and request assistance with their work efforts in the ASCI computing environment. It also documents how local support staff are kept current on the problems their local users are having in the distance computing environment.

Metrics are gathered to determine what problems are not resolved, how long (on average) it takes for a problem to be resolved, that user notification of trouble ticket occurs, and (if applicable) the user-local help is notified via e-mail. The following statistics are gathered: total number of problems reported, total number reported from remote users, and current type and/or rate of trouble tickets.

Conclusions

The DisCom^2 team made a number of significant accomplishments in usage model elements 1, 2, 5, and 7. The accomplishments to date integrate across the Tri-Lab capabilities that constitute an initial distance computing environment. Much work remains to bring the services and support to a generally available production state and then to deploy service improvements that address more complex usage scenarios (e.g., remote visualization).

Systems Integration of the Supercomputing Infrastructure and User Environment among the ASCI National Laboratories

While many program elements of ASCI are focused on terascale technology development, the impact of this development is made through the successful integration and deployment of these technologies into the ASCI user environment. This paper describes results of a recent ASCI milestone that focused on the deployment and integration of distance computing technologies to allow the Tri-Lab user community to access and make effective use of the ASCI White platform as if they were local. The focus for this milestone is the systems integration of institutional capabilities among the three ASCI Labs to provide a seamless distance computing user environment. Given the complexity of the elements of the Tri-Lab ASCI user environment, maintaining interoperability across the user environment is a significant challenge that is addressed through Tri-Lab coordination.