Red Storm Infrastructure

Robert A. Ballance, John P. Noe, Milton J. Clauser, Martha J. Ernest, Barbara J. Jennings, David S. Logsted, David J. Martinez, John H. Naegle, and Leonard Stans

Sandia National Laboratories*
P.O. Box 5800
Albquuerque, NM 87185-0807

May 13, 2005

Abstract

Large computing systems, like Sandia National Laboratories' Red Storm, are not deployed in isolation. The computer, its infrastructure support, and the operations of both have to be designed to support the mission of the organization. This talk and paper will describe the infrastructure and operational requirements of Red Storm along with a discussion of how Sandia is meeting those requirements. It will include discussions of the facilities that surround and support Red Storm: shelter, power, cooling, security, networking, interactions with other Sandia platforms and capabilities, operations support, and user services.

1 Introduction

Stewart Brand, in *How Buildings Learn*[Bra95], has documented many of the ways in which buildings — often thought to be static structures – evolve over time to meet the needs of their occupants and the constraints of their environments. Machines learn in the same way. This is not the learning of "Machine learning" in Artificial Intelligence; it is the very human result of making a machine smarter as it operates within its distinct environment.

Deploying a large computational resource teaches many lessons, including the need to meet the expectations and anticipation of a large user population, to centralize data collection and analysis, to automate system administration functions, and to dovetail the system inputs and outputs with a highly evolved, and continually evolving infrastructure.

Today, we are mid-way through the process of integrating Red Storm into Sandia's computational infrastructure. In another year, at the next Cray User's Group Meeting, we will be able to say even more about how Red Storm has been adapted to its environment as well as how its environment has co-evolved to Red Storm.

The Red Storm infrastructure consists of four primary components:

1. The physical setting, such as buildings, power, and cooling,
2. The computational setting which includes the related computing systems that support ASC efforts on Red Storm,
3. The networking setting that ties all of these together, and
4. The operational setting that provides management services

2 Facilities

Two new buildings were constructed for Red Storm. The Super Computer Annex (SCA) houses Red Storm proper. The Computer Utility Building (CUB) houses the water chilling equipment. The two buildings are located next to each other, and essentially adjacent to two other buildings that house (i) other ASC and Sandia computing resources and (ii) the network monitoring teams. The lay-

*Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company for the United States Department of EnergyÕs National Nuclear Security Administration under contract DE-AC04-94AL85000.

out of the buildings ensured that the physical, computational, and operational settings for Red Storm would all be located in close proximity.

Both the SCA and CUB are designed for unattended operations. Environmental monitoring is in place and monitored 24x7 by facilities personnel.

2.1 Super Computer Annex (SCA)

The Super Computer Annex is essentially a large, dedicated, machine room. Construction began in June, 2002 and was completed in October, 2003. The main room is $20,250$ square feet, with a clear span of 150 feet. Equipment bays, offices and other rooms occupy an additional $4,500$ square feet. A 50 foot ceiling provides ample space for air mixing and return.

Based on Red Storm's potential for expansion, the SCA is designed to support a much larger system. For example, the East wall of the main floor is not a load-bearing wall; the main floor can easily be doubled.

The facility delivers 3.5 megawatts of electrical power, although the current configuration of the machine requires just less than 2 megawatts. Transformer pads and wiring are in place in case a power expansion is needed.

The entire machine floor has a 36 inch plenum underneath. Only power, ground, and water pipes for the air handlers are under the floor. All wiring (network, fiber, etc) are routed in cable trays above the floor. The entire floor is pressurized to 0.5 inches in order to provide cold air to the Red Storm Cabinets.

Physically, the main floor of the SCA contains five equipment sections. Figure 1 shows a schematic. One row of power distribution units is arranged down the side of the room. In the center of the room is a corral of air handlers that surround the Red Storm cabinets. Inside the corral are the actual Red Storm equipment cabinets. Outside the corral are two banks of RAID storage.

The corral of air handlers draws air in from above, chills the air, and forces the air into the plenum. The Red Storm cabinets draw air from beneath and force it through the blades to an exhaust port on the top of the cabinet.

2.2 Computer Utility Building (CUB)

The Computer Utility Building (CUB) provides chilled water to the SCA. Construction began in June, 2002 and was completed in October 2003. At present it provides 2000 tons of chilled water to support Red Storm. The CUB can also be tied into the cooling water loop for building 880, which houses other SNL computational resources. Like the SCA, the CUB is designed for expansion; it can be expanded to 4000 tons capacity when needed, either by adding additional chillers or by replacing the existing chillers with larger units.

3 The Computational Setting

ASC Red Storm is being deployed into a complex computational milieu with an existing, skilled, user base. As the replacement to ASCI Red, Red Storm is required to function in ways that parallel ASCI Red. Innovations are encouraged, so long as they are improvements in the ways that the system operates or interacts with users and other systems. In addition, Red Storm is not entirely standalone. Key functions such as visualization services and archival storage are being provided by new platforms that are tightly coupled to Red Storm. One challenge is to make the integration of these plaforms and their functions as seamless as possible.

Deployment of Red Storm is guided by the *Red Storm Usage Model* [rsu04], a requirements document jointly negotiated among the ASC Tri-Labes users. The usage model defines the activities and tools to be supported on Red Storm. Conformance to the usage model constitutes a major milestone in the deployment of the system.

This section provides an overview of Red Storm's computational milieu. Section 4 provides details on the networking. For now, suffice it to say that the pieces are, in fact, tied together with a variety of 1Gb/s and 10Gb/s networks.

3.1 Computational Infrastructure

The core of the Red Storm computational infrastructure consists of Red Storm and its supporting visualization and storage clusters. Along with those clusters are a number of computing systems that support compilation, debugging, mesh generation, remote job management, visualization, and other activities directed by Red Storm users.

Red Storm is a dual-headed machine with connections to two separate networks: one for unclassified work and one for classified work. The two networks are similar, though not identical. In general, any system or service provided on the unclassified network will also appear on the classified network.

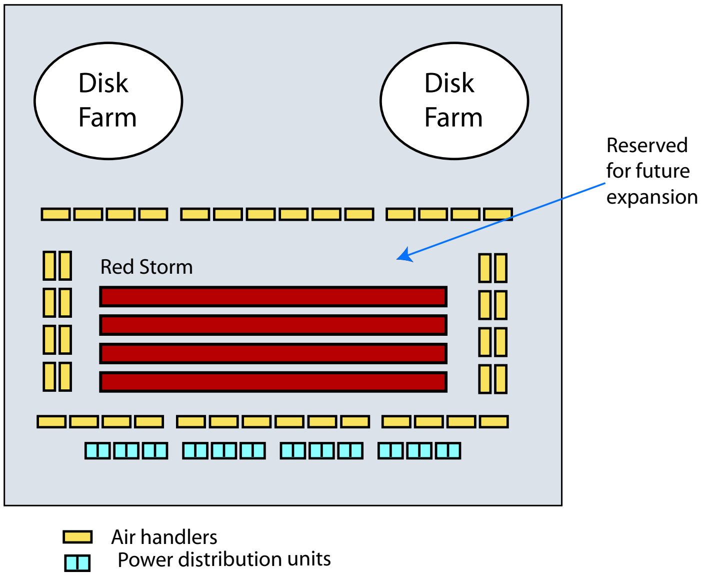

Figure 1: SCA equipment schematic

3.2 The RoSE Clusters

The Red Storm Environment (RoSE) clusters are Linux-based clusters that provide

1. Interactive data analysis, visualization and archiving of output from Red Storm, Sandia's mid-range computer systems, and other ASC platforms such as ASCI White, Q, Purple and Blue Gene/L, and

2. High-performance institutional file system providing high-bandwidth shared access to terascale data sets and direct read-write access by Sandia's mid-range computers.

The published performance requirements include:

Vis Power for highly interactive visualization and analysis of large (300 mega-cell and larger) datasets including at least 4 X present capability of Feynman[1] Vis nodes.

- 10^{10} (peak) triangles/sec render power. and
- 4 Teraflops (peak) compute power ($\approx 1/10$ of Red Storm).

I/O Power for accessing terascale data within RoSE cluster at interactive rates including

- 25 GB/s parallel file system,
- 1 second to access one time-step of a 300-mega-cell calculation for moving terascale data from Red Storm, and
- 25 GB/s (90 TB/hour) parallel file transfer, to minimize impact on Red Storm file system.

Two different clusters are being deployed: Black RoSE for unclassified work, and Red RoSE for classified. Both are Linux clusters providing visualization nodes, HPSS tape storage silos, and Lustre file systems. They differ in that Red RoSE uses Infiniband for its high-speed network, while Black RoSE will initially be deployed using Myrinet. Figure 2 shows a the Red ROSE architecture. Black RoSE follows a similar architecture.

The RoSE systems are tightly coupled to Red Storm's 10Gb/s ethernet I/O modules, as shown in Figure 3. The key to the connections is the existence of a private 10Gb/s network backbone that is primarily dedicated to data movement between the two systems. The *pftp* client is being modified to exploit the parallel paths between the systems in order to achieve the mandated 25GB/s throughput. A secondary effort is looking at how capabilities already designed into the Lustre file system can be adapted to allow Red Storm compute nodes to mount Lustre file systems provided by RoSE.

The data paths between Red Storm and RoSE shown in Figure 3 are indeed parallel. Each system provides 50 nodes equipped with a 10Gb/s ethernet connection. The

[1]The Feynman cluster is a predecessor to Black RoSE.

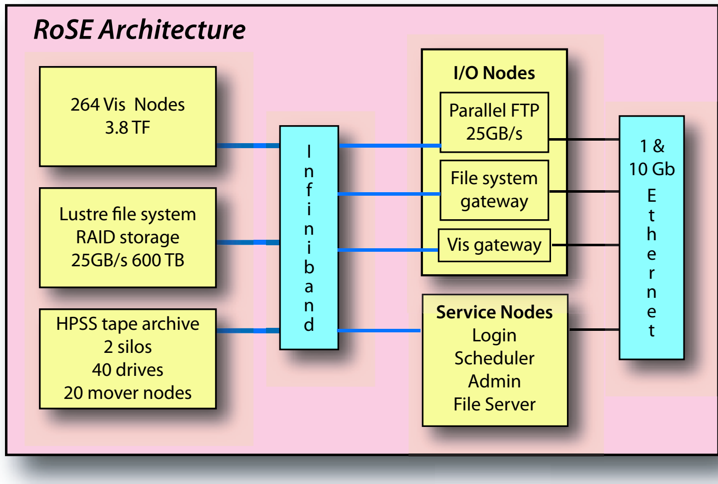

Figure 2: Red RoSE schematic

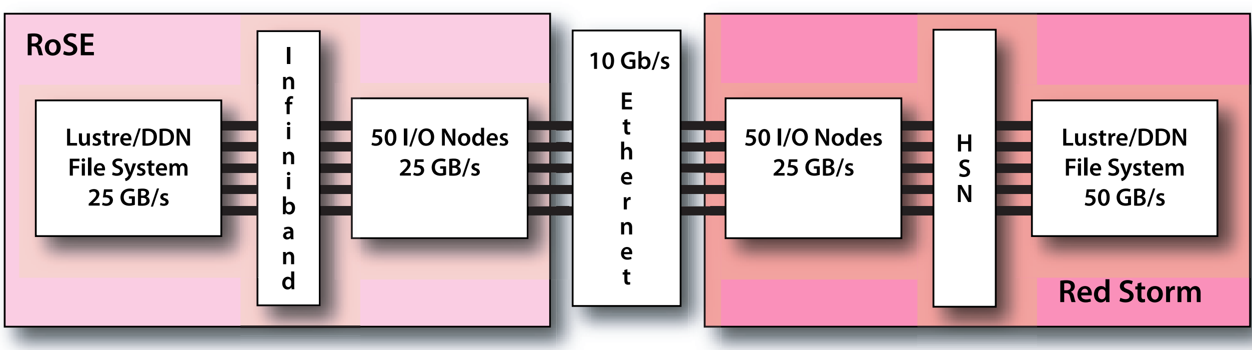

Figure 3: Red Storm and RoSE interconnections

private network between the two systems provides sufficient capabilities to support the full aggregate bandwidth of 25GB/s.

3.3 Data Services

Several methods for connecting the Red Storm and RoSE file systems have been discussed. At present, we are planning to use the parallel FTP (PFTP) tool as the primary data movement vehicle.

Sandia supports two tools on Red Storm for transferring files to and from the SNL HPSS tape archives attached to the RoSE clusters. The *dsacp* tool provides transfer of directory hierarchies of data to and from HPSS. A *dsabrowse* utility provides ftp-based browsing and management of data on the HPSS. In addition, the *zephyr* tool is available to move directory hierarchies of data between supported systems.

The SNL HPSS tape archives attached to Red RoSE and Black RoSE each have one silo. Future capability of the silos include a capacity of 2 petabytes and a bandwidth from disk to tape of about 0.5 GB/s. Black RoSE has 8 mover nodes connected to 16 HPSS tape drives. Red RoSE has 10 mover nodes connected to 20 HPSS tape drives. Each mover node has dual 2.4-GHz Xeon processors with 2.0 GB of memory. They have two 2-Gb/s Fibre-Channel connections to their respective tape drives. The parallel file systems on the RoSE clusters provide staging areas that allow data to be transferred (via *dsacp*) rapidly from Red Storm, and then more slowly to tape.

3.4 Other Computing Platforms

Sandia operates a variety of job preparation and analysis platforms alongside the RoSE clusters. These in-place systems are being used during the migration to the RoSE systems. In addition, they provide specialized capabilities such as large, shared-memory 64-bit processing, that the clusters do not yet challenge.

Other large Linux clusters, using Myrinet or Infiniband interconnects, also provide large volumes of capacity computation to Sandia's users. Key applications run on all of Red Storm, ASCI Red, and the clusters. A crucial factor in the successful transition to Red Storm is to provide a user environment that can leverage the experience and training of our users.

3.5 Compilation platforms

Sandia is providing Opteron-based workstations that will provide cross-compiling environments for Red Storm. These workstations augment Red Storm's capability to compile applications on login nodes. One compilation server is dedicated to unclassified work and another is for classified applications. Compilation servers also act as the license servers for the Red Storm PGI family of compilers.

3.6 Vis Labs

Sandia's large computational resources such as Red Storm are designed for manipulating massive datasets. Viewing the resulting multi-terabyte visualization data on a standard computer monitor would be rather like conducting a geological survey of the Grand Canyon by peering through a straw. While removing the straw and opening one's eyes to take in the full grandeur of the Grand Canyon is straightforward, opening the field of view to visualize the electronic data is more involved.

In recent years, Sandia National Laboratories has begun to "remove the straw" through the creation of Data Visualization Corridors (DVC), informally known as Visualization Labs, or Vis Labs, for short. First and second generation facilities have been deployed at Sandia National Laboratories, while new facilities are on the drawing board. The following provides some detail on facilities deployed in the past, those recently set up, and future facilities:

VIEWS The ASCI VIEWS Corridor Facility, deployed Fall 2000 at SNL/NM

- 48 tiled DLP projectors, each at 1280x1024 resolution, arranged in a 12x4 array create a 62.9 megapixel display
- Three 10.2ftx12.7ft screens joined at 45 deg angles provide a 38ft wide pseudo-immersive environment
- 3500 Lumen projectors are bright enough for use under full-lighted room conditions
- Image tiles are aligned to within 1-pixel (0.75mm)
- A 64 node, off-the-shelf PC, Linux render cluster drives the high-resolution display
- Data is rendered by commercial and custom in-house software

DISL Distributed Information Systems Lab Interactive Design Center, deployed Summer 2004 at SNL/CA.

- 27 projectors in a 9x3 curved array provide comprise a 35.4 megapixel display

- A 7.5ft x 30ft segmented screen
- Six ancillary smart displays for collaboration, videoconferencing, and system control
- Two 80-processor GraphStream graphics computing clusters, two 16-processor Silicon Graphics servers, and 18 off-the-shelf Windows-based PCs drive the visualization environment

JCEL Joint Computational Engineering Laboratory, deployed Spring 2005 at SNL/NM. This facility is similar to the VIEWS Corridor with the following distinctions:

- 24 tiled "dark-chip" DLP projectors, each at 1280x1024 resolution, arranged in a 6x4 array, create a 31.5 megapixel display
- A single 10ft x 20ft screen
- A 12-node, dual graphics card PC Linux render cluster drives the high-resolution display

MESA Microsystems and Engineering Sciences Applications Microlab Design & Education Center, scheduled deployment Spring 2006 at SNL/NM.

- Design is similar to the JCEL Vis Lab (above)
- To enable real-time display of nanotechnology, the design includes a microscope with the ability to output images for projection onscreen.

MESA/WIF Microsystems and Engineering Sciences Applications Weapons Integration Facility, deployment date for SNL/NM not yet scheduled.

- A 60-person theater-style visualization facility is planned based upon the high-resolution, multiple image tile model used in previous facilities detailed above.

These visualization facilities strive to open a clear unrestricted window into the massive modeling and simulation data sets used for weapons engineering. By leveraging the increased graphics performance realized in the computer gaming market, within a scalable graphics render cluster linked to a high-resolution display, a transparent sizeable window into massive simulation data approaches practical reality.

4 Network

The computational platforms, servers, and workstations that provide services or access to Red Storm are all located on networks either within Sandia, or connected to Sandia via external networks. Both unclassified and classified networks are supported. Services on the unclassified networks are generally mapped into equivalent services on the classified networks (e.g. both support login, email, and Web services), although the details may differ. Figure 4 shows a high-level diagram of the Red Storm data networks.

Red Storm's login nodes are accessible via 1Gb/s switch connections. The system management workstations are on a distinct private network that can be configured and monitored separately. Finally, a third private network allows outbound traffic from an administrative node designated as a "status" node. This connection allows remote management and monitoring tools, such as a Web site that displays current PBS queue status, to receive information directly from the PBS server.

In preparation for the Red Storm deployment, a significant portion of the high-end users were converted from the existing 1Gb backbone to a modern high-performance infrastructure. The network on both the classified and unclassified environments now consists of a multi-10Gb backbone that supports 1G to the desktop and 10G to large servers and a few user desktops as well. Figure 5 provides a high-level diagram of the backbone networks. This user network has been in production for over a year. When Red Storm and the data services cluster come on-line, user desktops will be an integrated component of the high performance computing network environment.

Red Storm and the Data Services machines (including the RoSE clusters) are the only one that can currently take advantage of the very large bandwidths provided by 10's of 10Gb/s interfaces.Ê Sandia currently has over 80Gb/s of bandwidth available to the rest of the production network from Red Storm and Data Services. This can be scaled up as other machines become available that can utilize more bandwidth.

Red Storm is also supported by one of the fastest encrypted long-haul production networks known to open literature.ÊConnections to our other Tri-Lab partners; LLNL and LANL, run over encrypted OC-48 (2.5 Gb/s) links. We have demonstrated 200 MBytes/s of sustained user data transfer to/from remote Tri-Lab sites over this 1100 mile connection.

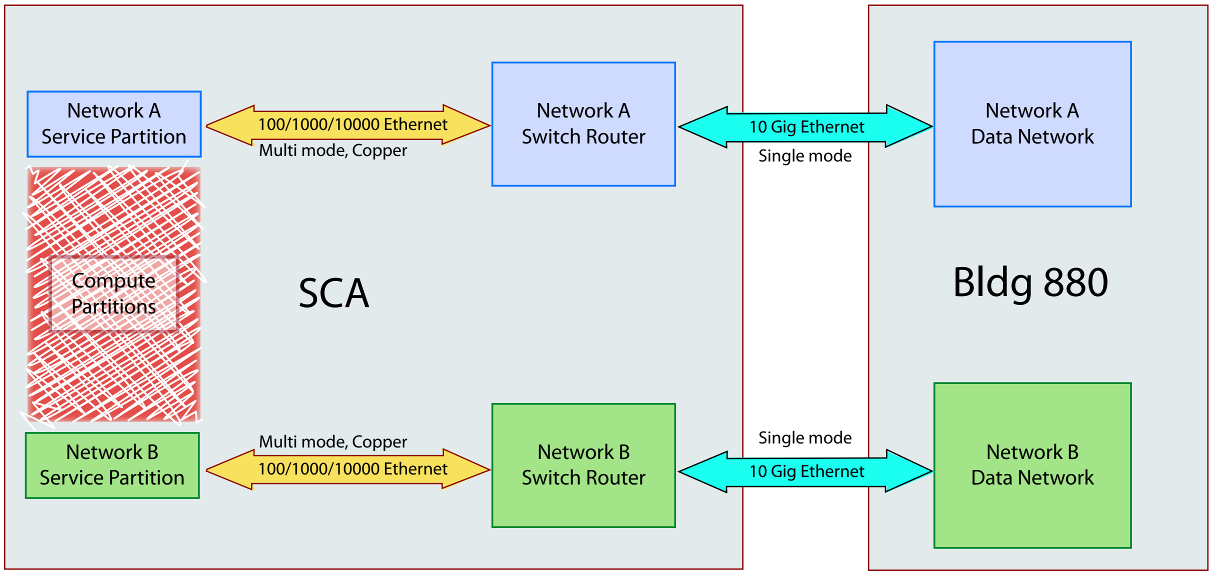

Figure 4: Red Storm Data Networks

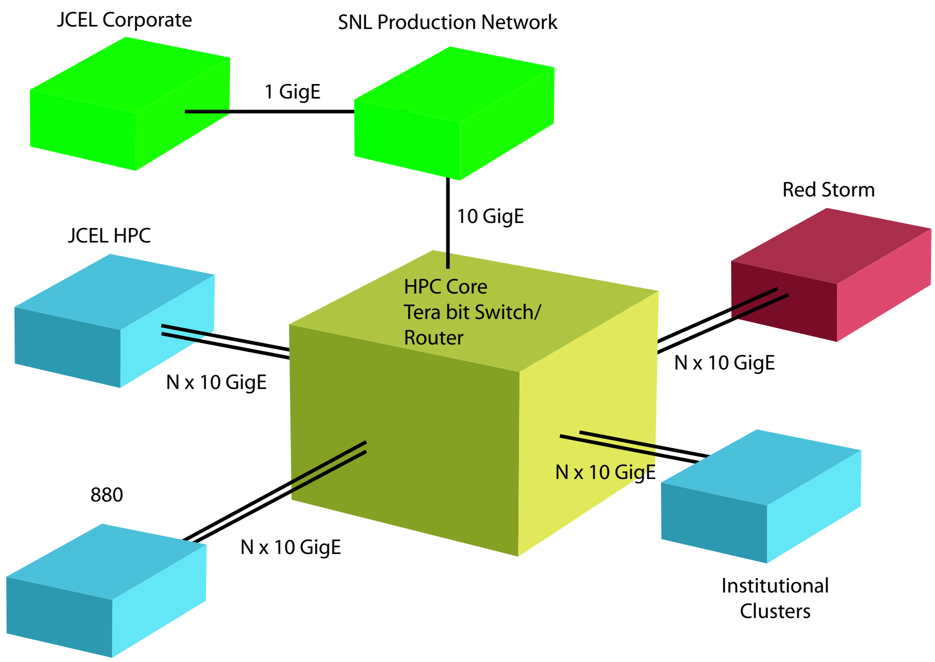

Figure 5: Sandia Backbone Network (Unclassified)

5 Operational Setting

The operational setting for Red Storm consists of the services, procedures, and servers that underpin the daily operations of a large-scale computing center.

5.1 Network Services

Sandia Corporate Services provide the usual networking services in a distributed environment: dhcp, DNS, NTP, and electronic mail. However, not all of the Red Storm network connections are visible to the external network environment.

For example, parts of the network connections between Red Storm and the RoSE clusters are configured in such a way that dhcp cannot be used to allocate IP addresses. Statically assigned addresses are required. During deployment, the first method chosen for assigning addresses proved to be impractical — it required node-specific specializations on too many nodes. The site-local Cray support team properly diagnosed the situation, and developed a fix that would allow us to administer all of the statically assigned interfaces via a single file consulted during boot time.

DNS has also proven to be idiosyncratic, since any round-robin scheme requires either an adaptation of the centralized servers or the provisioning of a Red Storm-specific server. Adding a new server to the network that appears to subsume or conflict with the centralized servers is discouraged in our network environment. Hence the need for an ongoing discussion about how best to provide DNS services for Red Storm.

Sandia relies on Kerberos for network authentication. External logins rely on access to the Kerberos daemons for their authentication tokens. Internal to Red Storm, we rely on passwordless *ssh* authentication. This means that we have to support both regular and Kerberized versions of the *ssh* daemon.

Red Storm relies heavily on email and Web communications for sharing information and knowledge about the system. The email servers and list servers are managed via the central corporate services; Red Storm is just another client. Parts of the Web-based services, including the CLIK knowledge community (Section 6.2) are hosted centrally; other portions are hosted by the organization charged with operations of the system.

5.2 User Accounts

User accounts are administered at Sandia using a collection of tools and databases collectively known as the Network Information System (NWIS). This is not to be confused with the old NIS databases, though it serves a similar purpose with respect to to account management. Built atop the NWIS structure are Web-hosted tools for handling account requests (local and external) and group management.

All ordinary user accounts are represented within NWIS, along with information about which systems are accessible to each user. An automated process downloads user and group information on a regular basis; and the system's local password and group files are updated appropriately.

Red Storm caused two ripples in this workflow. First, a Red Storm account carries with it privileges on a number of peripheral systems such as Black RoSE or the compilation server. Second, the actual files used on the system reside in the shared root file system managed by the boot node and the *xtopview* command. Peripheral accounts are being handled by modifying the update process to recognized tightly-coupled system. The file update process is being adapted locally by Cray to properly log and install new files, such as the groups file, via scripts.

5.3 Job Accounting

Many of the systems already mentioned, including Red Storm and the RoSE clusters, are managed by a single Center (9300) within Sandia National Laboratories. We use the AIRS [BGH03] software system to provide daily job accounting information to a central database. From this database, we can provide usage information tracking users, projects, groups, and systems to the users, project leads, and management. AIRS data can also be used to demonstrate utilization trends for procurements.

Red Storm's internal accounting package, delivered as part of the system, is being integrated into the AIRS database. This is primarily a matter of adding a new data feed into AIRS that uses the Red Storm internal database tables, and the augmenting the Red Storm information with information drawn from the PBS scheduler logs in order to track queue wait times and other behavior. Since AIRS already supports a number of data formats, this integration is expected to be a low risk task.

5.4 Remote Job Management

The usage model for Red Storm assumes that a user will be connected to a login node where jobs will be submitted to the batch queuing systems. Despite a foray into Grid computing, this has been the dominant model of job management at the Lab. However, this model does not integrate smoothly with high-level job management functions in which a large number of individual application runs have to be coordinated over space (different platforms) or time.

The DART project at Sandia is working to address a number of the issues involved with engineering analyses that require a number of HPC job runs. We are working with DART to provide interfaces to the Red Storm queueing system that will ensure accurate real-time job management and tracking without burdening the PBS server.

5.5 Monitoring

The XT3 product line, which includes Red Storm, has been designed with a powerful, event-based RAS management and monitoring system. The RAS system is able to monitor the health of Red Storm, to detect key process failures, and to provide failover operations for essential services.

At Sandia, the Cyber Enterprise Management (CEM) System Operations Center (SOC) is staffed 24x7 to monitor and diagnose issues arising either on our networks or on key machines located on the network. For example, The CEM System (CEMS — The Tools) monitors the mail servers, and the alerts from the servers bubbles up to Event List in the SOC. The Ops-Analysts are then able to notify system administrators of outages.

The Ops-Analysts are the point of contact for contacting system administrators in the event of failures or outages. From a system administration perspective, this means that Red Storm has to provide enough information to properly identify critical events, without overwhelming the team with data that can be routinely evaluated during normal working hours.

A key tool in the CEM approach is to define a "Business Service View" which could be thought of as a "user perspective" for the system. This perspective correlates failure events (e.g. a daemon quits working) with user-visible behavior. This allows the SOC Ops-Analysts to work backwards from problem reports ("The mail server won't upload my mail") to probable causes. Development of a user perspective model for Red Storm is currently underway.

How do you integrate the RAS system with the CEMS? This question is the focal point of a discussion between the Red Storm administrators and the CEM team. On one hand, Red Storm provides powerful diagnostic and failover capabilities. At times, the two groups have similar capabilities. For example, CEM uses a standard probe to check the viability of scheduler queues on Linux Clusters. Red Storm can do the same thing, using the RCA daemons in the RAS system. The problem is to get the RAS system to talk to the CEM system in a way that allows CEM to get the data they need without additional overhead on Red Storm administrative or login nodes.

Pursuing the PBS example, Red Storm not only monitors the PBS server, but can restart it if a failure is detected. All CEMS needs to know at this point, is that the failover was complete. However, it would be desirable if CEMS could detect a sequence of failure/recovery events in PBS, since such a sequence might indicate problems requiring rapid attention.

CEMS also monitors the network connections to Red Storm, along with the health of key support services such as Kerberos. However, the data network between Red Storm and RoSE is not so easily monitored as the login network. A team that includes CEM, networking, and Red Storm is currently looking at ways to expose network traffic monitoring data on the data networks for both realtime and retrospective analysis. The performance requirements on the network require a small footprint for the monitoring tools. Solutions are currently being evaluated.

The state of the RAID disk systems provides a similar challenge: the DDN RAIDs can talk over the internal RAS networks to agents on the management workstations. We are now starting the process of determining which status and events need to be passed on to CEMS, and which can be safely held on the SMW.

6 User Support

Red Storm is an ASC capability platform, like White or Q. The ASC community has evolved a set of standards for user support and communications designed to make it easy for ASC users to find information about particular systems and to move their applications among ASC platforms. The Red Storm user support teams are following that model.

In addition to the ASC requirments, Organization 9300 is evolving a support model based on a collaborative learning model.

6.1 ASC Support Model

Key aspects of the ASC user support model include

- Web-based reference information,
- Near real-time status information about platforms, queues, and jobs,
- Phone and email support,

- Classroom and web-based training, and
- Printable documentation that conforms to standards.

The standards for ASC support follow a "pattern-based" approach. Web sites are not required to have identical structures or typography, but are expected to follow a general template so that users can anticipate how and where to find information. This approach means that each new instance, like Red Storm, can build on an experience base while integrating new approaches to communication.

At Sandia, the ASC model is being followed with the following new additions:

- The Analyst Home Page provides a view of all the Sandia systems from the reference point of an HPC engineering user.
- A collaborative knowledge environment is used to capture all electronic interactions with users, as well as to store and present information gained during the operations of Red Storm.

The Analyst Home Page is a classic Web reference site: it recasts the information provided by many systems into a common format that is easy to navigate. Most links point back into the individual system's web sites.

The collaborative environment is described next.

6.2 Scientific Computing Support

Sandia Scientific Computing Support (SSCS) is being developed as part of the total support that is required for users of Red Storm the next super computer in line for release as well as all High Performance Computing (HPC) resources at Sandia National Laboratories (SNL). The support required for the HPC environment is quite unique in that: 1) the dynamic landscape is constantly changing and actually being created on a daily basis, 2) the customers are often the experts, 3) problem solving requires research into areas that may or may not have previous record, 4) the users are local as well as remote and 5) support must be interoperative among the support teams from local sites as well as the remote host locations.

To meet the needs of this unique environment, the Scientific Computing Department at SNL has designed a model for HPC customer support that incorporates an online collaborative support environment and is also supported by advanced academic research. In brief, this environment is a shared wisdom community, making available to all, the experience and knowledge of the users, designer, application developers and system administrators, vendors, and our sister labs. This requirement mandates secure access by local and remote customers to multiple previously established information resources and the creation of a resource to capture ongoing knowledge as it is developed.

A custom designed web based application has been created to allow for the required latitude in user access (secured by need-to-know and platform independent) while incorporating existing support practices to have minimal impact on current user expectations and support operations at each of the locations.

Applications for support at the corporate level were reviewed and researched to determine the allowance for customization required by the design. Our research resulted in the determination that proprietary solutions were not flexible enough for the required shared environment. Further any customization that was incorporated would have to be redone for each version change or upgrade required by the proprietary solutions.

The resulting design is configured to allow for interaction with the current support operations that are provided for HPC and desktop users at the local and remote locations. Further this design incorporates role based user access to allow for information control within the current Need-To-Know regulations and incorporates several existing corporate information resources including NWIS, WebCARS, and the metagroups access control application.

This design, based on information gathering, dissemination, and management results in a collaborative knowledge resource that doubles as a tool to be utilized by users who need to research information and consultants who need to research and compile available problem solving techniques. This web based tool is referred to as CLIK[Jen05], or Collaborative Learning, Information, and Knowledge. CLIK interfaces with support provided by telephone, email and web access and enabling user transformation from novice to expert while providing consultants the standard problem tracking techniques of vendor supplied applications customized to the SNL environment at one quarter the cost of the proprietary solutions that do not completely fulfill the identified need for shared resources.

7 Thanks

Integrating Red Storm into AIRS, or modifying the RAS system to send new kinds of traps to our monitoring center, are excellent examples of how a system like Red Storm gets adapted to its operational environment. Cray

has provided many excellent resources for starting the adaptation. In cooperation with Cray, Sandia engineers are working to develop other extention points. Together, we can grow Red Storm into the operational system that its designers envisioned.

References

[BGH03] Robert A. Ballance, Jared Galbraith, and Roy Heimbach. Supercomputing center management using AIRS. In *ClusterWorld Conference and Expo*, San Jose, CA, June 2003. 8

[Bra95] Stewart Brand. *How Buildings Learn: What Happens After They're Built*. Penguin Press, 1995. 1

[Jen05] Barbara Jennings. Distributed cognition of as a solution to high performance computing customer support. In *Cray User's Group*, 2005. 10

[rsu04] The Red Storm usage model. Internal document, 2004. 2