Chandra hardware and systems: keeping things running

Lisa Paton[1]
Smithsonian Astrophysical Observatory, 60 Garden St. Cambridge, MA 01749

ABSTRACT

System management for any organization can be a challenge, but satellite projects present their own issues. I will be presenting the network and system architecture chosen to support the scientists in the Chandra X-ray Center. My group provides the infrastructure for science data processing, mission planning, user support, and archive support and software development. Our challenge is to create a stable environment with enough flexibility to roll with the changes during the mission. I'll discuss system and network choices, web service, backups, security and systems monitoring. Also, how to build infrastructure that's flexible, how to support a large group of scientists with a relatively small staff, what challenges we faced (anticipated and unanticipated) and what lessons we learned over the past 6 years since the launch of Chandra. Finally I'll outline our plans for the future including Beowulf cluster support, an improved helpdesk system, and methods for dealing with the explosive amount of data that needs to be managed.

Keywords: computer, systems, architecture, clusters, security, user support, network

1. INTRODUCTION

1.1 The Chandra mission

NASA's Chandra X-ray Observatory, which was launched and deployed by Space Shuttle Columbia on July 23, 1999, is the most sophisticated X-ray observatory built to date.

Chandra is designed to observe X-rays from high-energy regions of the universe, such as the remnants of exploded stars. The spacecraft spends 85% of its orbit above the belts of charged particles that surround the Earth. Uninterrupted observations as long as 55 hours are possible and the overall percentage of useful observing time is much greater than for the low Earth orbit of a few hundred kilometers used by most satellites.

The Chandra X-ray Observatory is the U.S. follow-on to the Einstein Observatory. Chandra was formerly known as AXAF, the Advanced X-ray Astrophysics Facility, but renamed by NASA in December 1998. The Chandra spacecraft carries a high-resolution mirror, two imaging detectors, and two sets of transmission gratings. Important Chandra features are: an order of magnitude improvement in spatial resolution, good sensitivity from 0.1 to 10 keV, and the capability for high spectral resolution observations over most of this range.

1.2 Original science center computer and network architecture

The computer and network design for the CXC data system was first proposed in 1994. At that time current best practices for network and computing design were quite different than today. There was no widespread use of switched networks. Groups of systems were on shared bandwidth hubs. These hubs generally connected at 10Mbit speed. The World Wide Web was just catching on with the release of Mosaic in 1993. The first Linux kernel was written in 1991 by Linus Torvalds and hadn't yet established a wide following. Sun, DEC, and SGI were dominant players in academic computing.

---

1 Send correspondence to LP; E-mail: lpaton@cfa.harvard.edu

The first design for the CXC was based on these factors. The various groups were segregated into small subnets. Data was largely kept local with each individual system. A guest observer facility was envisioned that would allow users to visit the CXC and work on their data locally because sharing information and tools was difficult without web sites. We also planned to set up servers that observers could log into remotely to do data analysis since computer costs were high. Tools were developed that could be run on an observers desktop to reduce the dependence on the network since bandwidth was very limited.

From the beginning design decisions were made that would allow for ease of maintenance. The desktops for the CXC scientists were to be similar to those prevalent in the astronomical community. A single vendor and operating system was selected for all systems. A RAID based archive was designed to hold processed data in an easily searchable format. ANSI and other standards were followed wherever possible.

2. CHANDRA X-RAY CENTER (CXC) GROUPS

The systems group supports a diverse population within the CXC. It is important to meet the various requirements of the different groups and still allow the systems to be supported by a relatively small staff. By providing centralized copies of software we are able to supply a very rich environment. Each package only needs to be installed once for access by all of our various users. The advent of the web has also allowed us to centralize many services on our three main web sites; chandra.harvard.edu for public outreach; cxc.harvard.edu for scientific users; and icxc.harvard.edu for internal users.

2.1 User Support

User Support is responsible for all interactions with the Chandra User community. They coordinate the Chandra Peer Review process, answer scientific helpdesk queries, and keep the user community informed about what's happening with Chandra.

2.2 Mission Planning

Mission planning creates and maintains a catalog of approved targets. From these they prepare a long-term science plan and short-term schedules of about 1 week duration. It is their job to maximize the scientific output of Chandra.

2.3 Archive Support

The Chandra Archive is the repository of all telemetry and data products produced by the CXC data processing. The archive provides a secure, easily searchable location for data. It also feeds information from Mission Planning into the data processing system, and observing proposals from User Support to Mission Planning. The archive will be maintained by NASA's High Energy Astrophysics Science Archive Research Center (HEASARC) after completion of the Chandra mission to provide future data access.

2.4 Data Processing

The Data Processing Operations group (OPS) is tasked with processing Chandra telemetry once it is received from NASA. They utilize several pairs of systems to do the automated processing. The threads they manage are: standard data processing, reprocessing, special and custom processing, and software release testing. They are also tasked with Validation and Verification (V&V) and distribution of the data.

2.5 Science Operations Team

The OPS group also provides data to CXC Science Operations Team (SOT), which provides oversight for the health and safety of the spacecraft, monitoring of long-term trends, and assisting with issues involving spacecraft anomalies and instrument performance.

2.6 Software Development

The software development team is responsible for writing all of the software for proposal planning, pipeline processing, V&V, monitoring and trends analysis and scientific data analysis.

3. CURRENT CXC COMPUTER AND NETWORK ARCHITECTURE

3.1 Network

With the advent of switched networks, subnets are no longer required for efficient data flow. The CXC is served by a small number of extended subnets. Three of these are necessary because the CXC is spread out over three different locations. The scientists are primarily located within one building on Garden St.; operations is co-located with the Chandra Operations Control Center (OCC) at Hampshire St. and the software development group is in a third location at Cambridge Discovery Park (CDP). In addition to these subnets the CXC has a secure network used to restrict access to the Chandra Archive.

Links between the various sites are gigabit fiber connections. Within each building we provide gigabit connections for the large servers and 100Mbit connections to the desktop.

The networks are secured with a Cisco firewall, which provides protection from the outside as well as restricted access between groups within the Smithsonian Astrophysical Observatory (SAO). SAO also provides access to several unsecured subnets allowing for wireless and DHCP connections.

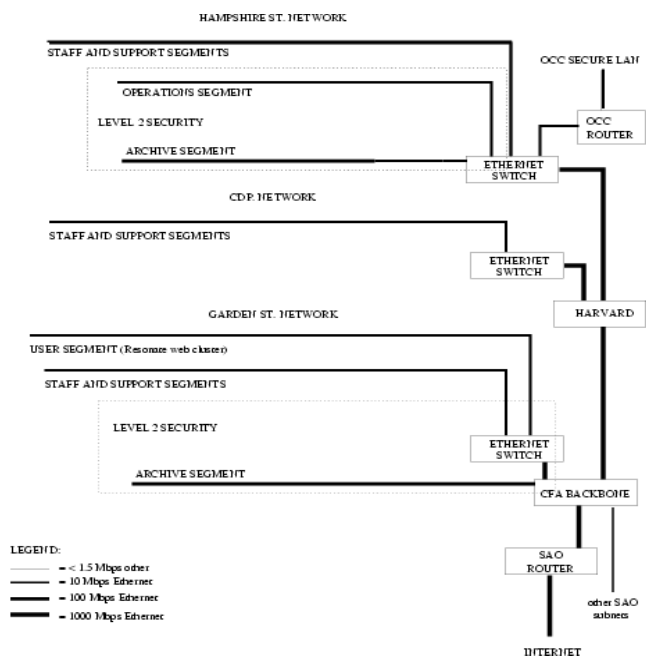

Table 1. CXC Network configuration

3.2 System Types

3.2.1 Sun Solaris

We initially selected Sun servers for their reliability and speed. These have been widely used throughout the CXC as both servers and desktops. Using a single architecture for all systems provided great economies of scale. A single set of software packages could serve scientific desktops, OPS, web servers, the mail server, and archive servers. All system staff could become expert in use and maintenance of these systems. A single maintenance contract could cover all hardware. The solution for an issue on one system could be used across the board. Systems managers could effectively automate many processes because the setup for all systems had been standardized.

For OPS reliability is critical to getting science data to the users in a timely fashion. For that group redundant hardware is provided. If any system were to fail, processing can be moved to identical hardware and immediately restarted.

The archive also has redundant hardware. Two sets of archive hardware, software and databases are available in two separate buildings. This provides for redundancy and disaster recovery. Because of this setup, archive upgrades can be done without interruption in access by the astronomical community.

Sun workstations still provide core infrastructure for the CXC by serving as email, web, DNS and NIS servers. These are areas where reliability is key and Sun servers have proven to be very robust.

Approximately 2/3 of the desktops used within the CXC are Sun Solaris systems. The majority of the rest are now Linux systems with many users now having both in their offices.

3.2.2 Linux

The future of the CXC lies with Linux. We are currently in the middle of a multi-year transformation from a Sun centered environment to one based on Linux running on Intel or AMD based systems. The primary driver for this is price/performance. Linux offers a very cost effective platform for doing all of the work currently performed by Sun systems. We are now buying very few new Sun systems. Because of the normal computer lifecycle we anticipate the majority of the desktops being Linux systems within the next three years. We migrate servers more slowly, but the majority of server systems should also be Linux based within five years.

There are many obvious as well as hidden costs associated with this migration. The transition obviously requires new hardware to be purchased. In addition we have needed to purchase new versions of software already provided on the Sun platform. The less obvious cost is in the time required to support an additional platform. A great deal of staff time has been spent building the infrastructure required to support Linux. We have had to build new install servers, write new scripts, install and maintain additional copies of software and learn new ways of administering the systems. Since Linux is similar to Solaris some much of the expertise of the systems staff applies to the new system, but many things are different.

3.2.3 Windows

Microsoft Windows systems are used primarily by our administrative staff and on dual-boot laptops. These systems are not well suited to scientific data analysis within the astronomical community and have not become widespread within the CXC.

3.2.4 Macintosh

Since the introduction of Mac OSX, Macintosh systems have been gaining popularity within the community. They are in use primarily as laptops and are popular because they allow the user to run common astronomical software as well as standard Microsoft tools such as Word, Excel and Power Point. Integrating desktop Mac's into the network has been challenging.

3.2.5 Network Appliance

At the core of the CXC infrastructure are powerful Network Attached Storage (NAS) devices. The CXC has selected Network Appliance as the preferred vendor. These devices allow for fast and reliable central storage and backup of important data. Centralizing data storage makes managing large volumes of data much easier.

4. RELIABILITY, AVAILABILITY, SERVICEABILITY (RAS)

Originally an IBM term, it refers to a computer system's overall reliability, its ability to respond to a failure and its ability to undergo maintenance without shutting it down entirely. (2)

4.1 Reliability

4.1.1 Hardware

Selecting appropriate hardware and software is critical to having reliable systems. No matter how well engineered your network is, no matter how responsive the helpdesk is, no matter how easy fail over is, unstable systems will take significant resources of an IT department. For the CXC selection criteria was heavily weighted toward stability in purchasing computer systems. Since we had strong reasons to stay with a single vendor (or single operating system) this choice is even more critical.

In the early days of Chandra Sun Microsystems was a well-known vendor with excellent credentials. Their systems were stable, often going long periods without crashing. SAO had been using equipment from Sun in other projects and their track record was good. Their engineering support was responsive in the few cases that it was needed. This was an easy and obvious choice. The OPS systems running Chandra pipeline processing have had very few failures in the seven years many of these systems have run. Upgrades and patches have been easy.

It is unfortunate that the SPARC architecture that is the basis for these systems has not kept pace with regard to price versus performance. Because of this the scientific staff of the CXC have been putting increased pressure on the IT managers to support Linux systems running on the Intel and AMD chip sets. Our challenge is to create as stable an environment on Linux as we have with Sun Solaris. We have been evaluating several possible hardware vendors to find one that meets our criteria for reliability. Our initial preference was to select a local vendor who could work with us to build a standard configuration for a set of Linux desktops and servers that could be deployed within the science center. It was our belief that if we could standardize the hardware platform the job of maintaining the operating system would be simplified. We found that many large vendors would switch out components so that a system you purchased today could have different components from one you purchased a month later. Unfortunately the local supplier didn't have the resources to handle issues that arose with several of their systems. Turnaround times of a week or more were seen. This was unacceptable and we continued to look for a hardware vendor. This remains an open issue. The majority of our Linux systems now come from either Dell (about half of all our Linux systems) or Sun (about one quarter). Sun has introduced a new line of systems based on the AMD processors. Some of their desktop systems had fan issues and they are slightly more expensive than other vendors but service and support remains excellent. Luckily the Linux development cycle has been so swift that many of the early issues with hardware support have already been addressed by changes in the Operating System (OS).

4.1.2 Operating System

When the decision was made to move away from Sun's SPARC based systems we needed to make the parallel decision of selecting a new OS. Sun's Solaris operating system was available for the Intel and AMD platforms, but it was not binary compatible with Solaris for SPARC. Our choices came down to Solaris for Intel, Red Hat Linux, Suse Linux, or other Linux varieties. We looked at survey data gathered by our user support team to see where the astronomical community was headed and did extensive research into the pluses and minuses of the various flavors of Linux. Red Hat Linux seemed to be the clear winner with the community. They also had strong support from the third party software providers that are important to our operations. In the midst of our decision making process Red Hat changed it's strategy and started offering a commercial version (retaining the Red Hat name) and a development version called Fedora. This split complicated matters greatly since the cost for Red Hat licensing needed to be factored into the equation. We also had to weigh the lower stability of Fedora against the faster cycle of bug fixes and new releases.

We have decided on a mixed approach. Our core systems, those given to software developers and those that will be used in operations, will use the more stable and widely supported Red Hat OS. Scientist’s desktops will run Fedora. The two operating systems are largely compatible and most software built for one platform will run on the other. This provides us with the required stability and compatibility with third party software, while keeping costs down and allowing scientists access to the latest enhancements.

4.2 Availability

No matter how reliable system's are, there will always be some level of failures. The CXC has done a great deal to eliminate single point failures where possible and to provide for easy or automatic fail over. Cost must be factored into this equation since having duplicates of all systems would be prohibitively expensive. Emphasis has been given to the most time critical systems especially those involving spacecraft operations and those that would impact large numbers of CXC users.

4.2.1 Backups

Backups are obviously of critical importance to any IT operation. The CXC is on its third backup system since the start of the mission. We have had the misfortune of having our two previous backup vendors get bought by larger software providers. Although we have found excellent products with great service and support, buying from smaller vendors has not proven to be a good decision. In the software market, size matters. Large vendors can simply buy out small companies and eliminate them as competition. The CXC has now gone with Backup Express from Syncsort, an industry leader in Linux and NDMP backup.

We began backup operations with two DLT4000 tape libraries that held just 7 tapes each. This has grown to 5 tape libraries with a mix of DLT7000, DLT8000 and now LTO tape support. This has greatly improved our backup capabilities. The new LTO-3 tapes can hold 400GB (uncompressed) each and write at 68MB/sec.

4.2.2 Security

There are many security issues that must be considered when designing a network. I'll be discussing only the issues involving system availability. Data security, which is primarily concerned with making sure the proper people have access to the data, is a separate issue. Security measures can both enhance and detract from a highly available network design. Security that is cumbersome, difficult to use or difficult to maintain can cause almost has many issues as the attacks that it is designed to prevent. The CXC has implemented security in many layers; perimeter, subnet level, node level, and user level. This allows for redundancy and reduced vulnerability to attack. It also allows for flexibility within the network so users don't need to perform multiple types of authentication and data sharing is not restricted more than is required to prevent system compromise.

The perimeter security is provided by a Cisco firewall. This allows for flexible access restrictions. As with most networks outbound traffic is generally unrestricted. We have, however, blocked certain protocols outbound to enforce business policy. For instance only designated email servers are allowed outbound SMTP connections. This prevents users on the network from setting up independent mail servers. Misconfigured email servers have been exploited as SMTP rely servers by spammers in the past. This can create unnecessary traffic on the network. Inbound traffic is generally blocked except for certain protocols on certain machines. We allow users to access their desktops remotely using ssh (secure shell). Many of our users require access to their systems when at home or on travel. Other protocols, such as HTTP, SMTP, DNS, and FTP are allowed only to designated servers.

Within SAO restrictions are placed between subnets so that less secure subnets, such as the wireless, DHCP, and self managed subnets are restricted from making certain connections to more secure subnets. Only trusted subnets are allowed full access to NFS (disk sharing) and other more sensitive protocols.

Node level security is provided on two fronts. First, we use programs such as TCP Wrappers to restrict access to nodes appropriately. Second, we make sure to keep patches and OS releases up to date so the number of system vulnerabilities is reduced. The CXC runs through a monthly cycle of patching for all Sun desktops. The Linux, Mac and Windows desktops use built in updater's to pull new patches from the Internet as needed.

User level security is dictated by policies set down by the Smithsonian and includes regular password changes, user security education and signing a computer use policy. User accounts are scanned periodically so that unused accounts can be disabled and passwords are tested to ensure that they are well chosen.

In addition to this all email is filtered to screen out potential viruses. Since Windows is not a primary desktop we are less susceptible than many organizations to viruses, but it is still in our best interest to avoid any such email reaching our users.

4.2.3 Monitoring

There are two types of monitoring systems in place within the CXC. Both were developed in house by the CXC systems staff. The first involves a series of checks to test the availability of certain systems and services. Failure for any of these tests triggers an email alert sent to the pagers/cell phones of key personnel. The tests range from simple pings to actual web site retrievals. The tests are run from systems both internal to and external to the CXC to ensure that a widespread failure will still produce a trigger.

The second type of monitoring in place is aimed at security and network stability. These tests show new hardware that has connected to the network or hardware that has changed its hardware address. They also report errors found in the log files of our systems that may indicate a failure and changes to important configuration files.

4.2.4 Resonate Cluster

Web servers have become critical to the operation of the CXC. The web servers now provide a forum for all phases of CXC operations. The CXC web sites host user information, proposal planning and submission, peer review support, mission planning, processing status, V&V, science operations including monitoring and trends analysis, archive access, calibration information and public outreach. To allow for highly available web access we selected Central Dispatch from Resonate. This software allows a cluster of web servers to support each distinct web site. If a single server within the cluster fails all subsequent connections will go to another node in the cluster. The software also provides load balancing so there is no drop in response time during peak loads.

4.3 Serviceability

Redundancy is the key to easy serviceability of our systems. The previous section outlined our web server configuration. Because of the fail over built into that system any component could be taken offline to be upgraded or replaced without the web sites themselves being effected. We have designed similar setups for our other critical components. Our Network Appliance file servers have hot swappable disks that allow them to be replaced without downtime. Our operations group can shut down processing on one group of servers and move to another group when maintenance is required. Only upgrades to the main file servers require actual downtime for any critical component.

5. LESSONS LEARNED

5.1 Everything is mission critical

No matter what plans are in place and how well defined your expected response time is, downtime for any reason at any time is disruptive. If components are selected that are high quality, and stable software is deployed, users can reasonably expect that each morning their computer will be ready and waiting for them to use. The best policy is to avoid downtime issues by following RAS principles and keep a well-trained staff to handle any unforeseen problems.

5.2 Know about problems before the users

The monitoring system in place at the CXC has saved many visits from irate staff members. On several occasions over the nearly nine years I've been with the project, I've been able to respond to an outage before most people knew that it had occurred. Our most common outages have come from power and air conditioning failures. Our building is old and does not have generator backup so these types of failures are beyond the control of my staff. Being able to bring systems back on line before the scientific staff arrives has brought us a great deal of good will.

5.3 Backups, Backups, Backups

Disks fail. People make mistakes. The files that are lost are most likely to be a person’s only copy of their thesis. We have had two cases of errant scripts running “rm -r *” from the root directory. Excellent backups are the key to surviving these types of situations. It takes about 2/3 of a full time employee to manage the backup system. This is an invaluable investment and not one that should be shortchanged.

5.4 Staffing

Nothing beats having a happy, well trained, knowledgeable staff when it comes to dealing with whatever curves are thrown your way. The newest member of the CXC systems team was hired six years ago. Within the IT community this low turnover rate is almost unheard of. I attribute this mainly to the respect shown to the system staff by the CXC staff and leadership. The IT staff members are invested in the organization and their input is valued. Most decisions on hardware, software, and networking are made by the people who work with the systems daily rather than a CIO who sits in a corner office removed from the day to day work of the group. The IT staff is also encouraged to learn as much about the mission as they can. This means that their decisions are based on a thorough understanding of the mission and it's objectives. As IT manager I can often anticipate problems and weaknesses in the system before they surface rather than reacting to a crisis.

6. THE FUTURE

6.1 Beowulf clusters

The CXC is just beginning to implement an additional level of processing designed to create a catalog of X-ray sources. This is a large, CPU intensive task. In order to provide enough processing power in a cost effective manner, a cluster of 15 nodes running Red Hat Linux has been set up. These dual CPU systems are managed by Rocks Clusters software and will be used in source detection and characterization. The use of clusters in astronomy is likely to grow in the future as more software is written that targets these platforms. Currently much of the software needs to be custom designed to make use of the clusters.

6.2 Helpdesk

The CXC systems group has just instituted a new commercial helpdesk system. This will replace a simple mailing list that has served as the helpdesk system for many years. The systems group uses a rotating on call schedule to handle user requests. Because of this there is occasionally a lack of follow through on issues that span many days. Moving to a formal helpdesk system has greatly reduced the number of issues that get dropped or forgotten. Our helpdesk closes around 300 tickets each month so having this software has greatly improved our overall service to the community.

6.3 Industry trends

MBBF (More Bigger Better Faster) – The computer industry continues to produce faster systems with more storage and larger network pipes. This is extremely fortunate since astronomers continue to invent instruments with greater data collection capabilities. Some of the more interesting innovations are outlined below.

6.3.1 iSCSI/SAN

iSCSI is a protocol that allows for SCSI connections over standard Ethernet networks. SAN stands for Storage Area Network. The combination of these two allows large file servers to appear as local disks on the client system. This is an excellent way to combine the high-throughput and RAID management capabilities of large file servers with the access speed of local disks.

6.3.2 Hierarchical Storage Management (HMS)

Because of the complexity of managing large volumes of data, various hardware and software solutions are becoming available that automate data management. HMS systems ensure that the most frequently accessed files are on the fastest file servers. Less frequently accessed files are moved to secondary storage and eventually to tape. All file are available to the users and appear to be in the current directory but these files may actually be links to less expensive storage media. This makes the most cost effective use of storage systems.

6.3.3 10Gbit networks

As demands on networks increase, 10Gbit networks are becoming more common. The network equipment to run at these speeds remains expensive, but like 1Gbit networking (which is now available on many desktops) it is likely to fall in price. The 10Gbit IEEE standard was only approved in June 2002(4). It will initially affect large data centers but is likely to migrate to smaller organizations in the future.

REFERENCES

(1) http://chandra.harvard.edu/about/axaf_mission.html

(2) http://cxc.harvard.edu/cdo/about_chandra/overview_cxo.html

(3)http://www.bitpipe.com

(4)http://standards.ieee.org/announcements/8023aeapp.html