

ATHENA TECHNICAL PLAN: Section G.1.1

Athena Changes to Berkeley 4.3 UNIX[1]

by Win Treese

1 Introduction

This document describes the changes made by Project Athena to the 4.3 Berkeley Software Distribution (4.3BSD) of the UNIX operating system. Athena began with 4.2BSD, served as a beta-test site for 4.3BSD, and is now converting to the released version of 4.3BSD. Many of the changes and bug fixes that Athena found necessary for 4.2BSD were incorporated into the 4.3 release. Hence, the number and magnitude of the changes described here are much smaller than before. Most of the changes now reflect the special environment Athena has developed, as we attempt to use UNIX in a distributed workstation environment.

This version of the document primarily addresses changes that have been made for use on the Athena timesharing systems (currently, approximately 40 VAX 11/750's). Later revisions will discuss more of the changes necessary to support workstations. Later revisions will also include details of specific changes as well as appendices describing which files have changed. It is likely that the configuration changes mentioned below will be described in a separate document instead of this one.

2 Kinds of Changes

Athena modifications to 4.3BSD fall into the following categories:

- Directory hierarchy
- Necessary system changes
- Configuration changes
- Bug fixes

1. UNIX is a registered trademark of AT&T Bell Laboratories

- New software
- Minor additional features

Athena has made several changes to the directory hierarchy. These are primarily to isolate software that Athena has added to the system. Necessary system changes are modifications to the software that are required by the Athena environment. For example, it was necessary to add the client device driver for the Remote Virtual Disk (RVD) system to the Berkeley kernel. Configuration changes are those made to various configuration files (e.g., /usr/lib/sendmail.cf). Most of these files are altered by most sites, but the Athena modifications are often more extensive, so some of them are described here. Bug fixes are exactly that; most of them are reported to Berkeley. Many of the Athena fixes to 4.2BSD were included in 4.3. New software includes both software developed by Project Athena and that acquired from third-party vendors. Athena-developed software tends to be system software necessary for the Athena environment (e.g., the Kerberos authentication system). Most of the third-party software packages are user applications, such as the gnuemacs text editor from the Free Software Foundation or the scribe text formatter from Unilogic, Inc. Minor additional features are changes made to make life easier, such as filename completion in the C shell.

In a sense, these categories also define the criteria by which proposed changes are judged. The Athena approach to software has been to buy it when possible, and write it only if it cannot be bought and is necessary for the operation of our system. For the most part, gratuitous changes are avoided, though there are several on the system. The success of this approach is demonstrated by the relatively small number of changes that were necessary for the released version of 4.3BSD.

3 Directory Hierarchy

Athena has added the following directories to the standard system: /usr/athena, /usr/athena/lib, /etc/athena, and /usr/unsupported. In addition, we will soon be adding a directory /bin/athena. These directories hold the expected classes of software: /usr/athena contains Athena user applications, /usr/athena/lib contains Athena-specific library files, and /etc/athena contains special system maintenance commands. /usr/unsupported is used for a few pieces of software that are not supported by Athena, but that we distribute for use. Generally, a program there is supported by a specific Athena staff member who wishes to make it available. The directory /bin/athena will be used for commands that need to be available on the root filesystem, but which are also accessed by normal users. User directories are placed in /mit, which is a mounted filesystem on the timesharing machines.

4 Necessary System Changes

Changes to the system software from Berkeley fall naturally into two categories: those made to the kernel and those made to user software. Although we do not relish changing the kernel, it has often been necessary. With the arrival of 4.3BSD, our changes have become more modular, since many of the bugs from 4.2 have been fixed. The three sets of changes necessary for 4.3 support the RVD system, the RTI Ingres database system, and the CHAOSnet, a network system developed at MIT.

The RVD system, in its current implementation, consists primarily of a device driver. Hence, the installation is much like that of any other device driver. It also requires, however, an addition to the network protocols table, since RVD uses its own protocol.

The RTI Ingres database, from Relational Technologies, Inc., uses a pseudo-device and an associated driver for database locking. This installation is a simple device driver change.

The MIT CHAOSnet system requires a somewhat more extensive set of changes. It is used, however, on only a few of our machines that serve as gateways between the MIT campus network and the existing CHAOSnet. We expect that this will not be a permanent feature of our system. The most common use of the software is for mail routing and occasional transfer of files between a CHAOSnet host and an Athena host.

Modifications to the user software, unfortunately, have not been so modular. They do not fall into any convenient groups, so they are simply listed here in no particular order.

The line printer system (`lpr`) was changed to give a remote printer specification in `/etc/printcap` a higher priority than a specification for a printer connected to a tty port. This allows us to construct an `/etc/printcap` file that can be used on all hosts, since the local printer is used only if the remote hostname matches the local hostname. In addition, we currently use the environment variable `APRINTER` instead of `PRINTER` to specify the desired default printer. (Note: I would like to back out of this for the January release.) The standard printer filters were also modified to support the LNO1 and LNO3 laser printers. In addition, the line printer system used the clustertab system discussed below for translation of printer names.

The utility `df` was extended to understand filesystems on RVD disk partitions.

The directory `/usr/athena/lib` was added to the list of library directories searched by `ld`.

For some reason, 4.3BSD did not include the network version of `write`. Because of its usefulness in our environment, we re-installed the 4.2 version.

Various manual pages were updated to reflect these changes.

5 Configuration Changes

4.3BSD contains many files that control the configuration of the system. Some of these, such as /etc/passwd and /etc/group, are changed by every site using 4.3. Others, such as /etc/syslog.conf are used by a specific program and are locally modified as necessary. There are also system data files such as /etc/services that are used by several programs. Finally, there are files such as /etc/rc that control the system configuration at boot time.

5.1 Normally Local Files

For timesharing systems, the files /etc/passwd and /etc/group for each machine are generated periodically from the Athena User Registration database. This database contains the necessary information for each user and group at Athena.

5.2 Program Data Files

Various programs, particularly daemons, use configuration files to control their behavior. These are described here, in no particular order:

The sendmail system is controlled by /usr/lib/sendmail.cf. This file has been extensively modified at Athena. All mail from an Athena machine is sent to a central mail server (currently the machine "athena"). This server then routes the mail to the appropriate machine, whether the destination is an Athena host or not. Mail to another Athena machine is tagged to indicate that it should not be forwarded to the mail server. On most machines, therefore, the file /usr/lib/aliases is not used, since no alias translation takes place locally. Instead, the mail server maintains a global aliases file, which contains the home mail machine of each Athena user as well as all mailing lists that Athena maintains. This file is generated from the user database, which also stores information about mailing lists. Workstations are not configured to receive mail at all; users of workstations can get their mail on a timesharing system (for now) or from a "post office" server.

The file /etc/disktab is a concatenation of the 4.3BSD version and the IBM ACIS 4.2A version. This allows us to maintain one master file, as well as create RVD partitions for IBM and DEC machines on the other vendor's physical disks.

/etc/shells, which lists shells acceptable to chsh, now includes the tcsh shell. tcsh is a shell that supports EMACS-style command line editing. It is not currently a supported part of the Athena system.

The syslog daemon routes logging messages based on the file /etc/syslog.conf. Athena uses the syslog facilities to log various messages, particularly from RVD. We plan to experiment with using syslog to log messages to a remote machine for workstations.

Athena uses the file /etc/named.ca to initialize the cache of the nameserver /etc/named. The standard file /etc/named.boot points to this file.

The program /etc/inetd watches for connections to ports as specified by /etc/inetd.conf. Athena added an entry for the write daemon, as well as for the Kerberos utilities (described below).

The files /etc/X?.hosts list hosts allowed to create windows on a display controlled by the X window system.

In addition to these files, Athena has added three configuration files: /etc/rvdtab lists information about RVD packs that a workstation may wish to access, /etc/athena/lib/clustertab contains the database for the clustertab system described below, and /etc/hosts.login lists hosts that may be accessed through the login server (described below).

5.3 System Data Files

Some configuration files are used by a number of programs. The file /etc/services contains a list of network services and port numbers for them. Athena has extended this file for several network services, including RVD and Kerberos.

The hosts database, /etc/hosts, is no longer maintained on every Athena host. Its function has been assumed by the nameserver, /etc/named, except for the initial setting of the host name and address. We plan to extend the program /bin/hostname to take the Internet address as an argument, so we can eliminate use of this file completely.

5.4 System Startup

Athena has extensively modified the files /etc/rc and /etc/rc.local to handle several configuration issues without having to carefully customize the files on each machine. This issue is discussed in "System Configuration at Boot Time" (Athena System Engineering Memo #?).

6 Bug Fixes

Fortunately, 4.3BSD has fewer bugs for us to deal with locally (at least, so far). In the beta test version, it was necessary to fix some problems with sendmail and the nameserver /etc/named. We have not fully tested the final versions at this time.

In addition, we have modified the gethostbyname() function in the C library to accept a "dotted quad" Internet address and return the appropriate host structure without accessing the nameserver. This enables several programs to accept an Internet address as an argument and perform their function without calling the nameserver. In particular, /etc/route can now establish a default gateway without first having to look up the name from a nameserver outside the local Ethernet.

7 New Software

There are two classes of software added to the system by Athena: those packages acquired from third-party vendors and those developed locally at Athena. Since third-party software packages are generally standalone applications, they required no changes to the system (with the exception of RTI Ingres, as discussed above).

Most of the system software developed at Athena is necessary to the Athena environment of distributed workstations. Other pieces were important for the programming or user environment, though not strictly necessary for the system. These packages are listed here in no particular order:

The "clustertab" (for cluster table) system is used primarily for translating printer names. In the timesharing environment, Athena had several "clusters" of VAXen. Each cluster might have its own laser printer, so the real name of a laser printer might differ from cluster to cluster. In order to simplify the system configuration, clustertab was developed to handle the name translation. In the workstation environment, the current implementation of clustertab is not sufficiently flexible. We expect to replace it with a generalized name translation service.

The X window system was developed to take advantage of high-resolution bitmapped displays. X is one of the most successful of the Athena-developed packages, as it was distributed with 4.3BSD and is being supported by DEC and, to some extent, IBM. Since Athena is still the focus of X development, the latest system available at Athena is usually installed on our systems, instead of that from 4.3

Implementation of the Remote Virtual Disk (RVD) system was originally done at the MIT Laboratory for Computer Science, though Athena has done much development since acquiring it from LCS. This system is described in detail elsewhere. To support it, we have added several commands for manipulating an RVD server and controlling client use of RVD packs.

The Kerberos authentication system, also discussed elsewhere, is being integrated into many applications, including some standard programs from 4.3BSD such as rlogin. For the most part, these changes are not visible to users of the system.

The CHAOSnet user software and servers are installed on those machines connected to CHAOS.

All Athena timesharing machines run a modified version of /etc/getty called a "login server". This version first prompts for a machine name, and then a username and password. If necessary, /etc/getty executes an rlogin to a remote host. This function enables users in a timesharing cluster to use a terminal connected to any other machine in the cluster without needing an account on each machine.

The Athena User Registration system consists of a database of all users, as well as several tools for managing users, groups, and mailing lists.

In the early days of Athena, the Penplot graphics package was adopted as an interim graphics library until a good implementation of the Graphics Kernel Standard (GKS) became available. Penplot was developed at the MIT Joint Computer Facility (JCF). Athena will be phasing out the use of Penplot, though we are now responsible for its maintenance since the JCF closed.

8 Minor Additional Features

Finally, there are some minor changes that Athena has made simply because we prefer that the software works in a certain way. The C shell, for example, has been modified to handle TOPS-20-style command filename completion.

9 Conclusion

The Athena changes to 4.3BSD primarily consist of a new "layer" of software on top of the standard Berkeley distribution. We modify many configuration files, but that simplifies identifying where changes need to be made. When possible, changes and bug fixes for software are shipped back to Berkeley for inclusion in later releases. This worked well for Athena with 4.2BSD and the beta-test version of 4.3, and we expect it to continue to be successful in the future.

10 Changes to this document

```
$Log: athena_changes.mss,v $  
Revision 1.2 86/10/31 15:34:22 treese  
Modified for inclusion in the Technical Plan.  
Added discussion of expected changes to this document.
```


ATHENA TECHNICAL PLAN: Section J
Athena Standards and Specifications

by J. H. Saltzer

This section provides citations for standards adopted by Project Athena, documents related to those standards, and references for third-party software adopted by Project Athena. The first two subsections are organized in parallel for easy cross-reference between standards and documents related to those standards.

The contents of this section are as follows:

H.1 Project Athena Standards

- H.1.1 UNIX standard*
- H.1.2 Print stream standard
- H.1.3 Network protocol standards

H.2 Related standards and publications

- H.2.1 UNIX standard
- H.2.2 Print stream standard
- H.2.3 Network protocol standards

H.3 Athena Standard Third party software

- H.3.1 Spreadsheet
- H.3.2 Numerical algorithm library
- H.3.3 Data analysis system
- H.3.4 Graphics toolkit
- H.3.5 Graphics library
- H.3.6 Graphics chart package
- H.3.7 Editor
- H.3.8 Text formatter
- H.3.9 Data base management system
- H.3.10 Algebraic manipulation system
- H.3.11 Mail management system
- H.3.12 Lisp interpreter and compiler

* UNIX is a trademark of AT&T

H.1. Project Athena Standards

This subsection provides citations for standards adopted by Project Athena. The next subsection, Related Standards and other Publications, is organized to be parallel to this section. In many cases, the related standards provide information that is useful in interpreting or understanding the adopted standards.

H.1.1. UNIX standard.

1. UNIX Programmer's Manual, Berkeley System Distribution 4.3, University of California, Berkeley, August, 1986. (Three volumes.)

Provides the fundamental specification for the version of UNIX adopted by Project Athena.

2. UNIX Commands, Functions, and User's Guide, Revision 1.1, M.I.T. Project Athena, March 15, 1984.

An out-of-date copy of item 1, with local modifications interleaved.

3. Academic Information Systems 4.2 for the IBM RT PC, second edition, IBM Corporation, Palo Alto, California, PRPQ #5799-CGZ, September 30, 1986.

Description of the differences between ACIS 4.2A and the Berkeley UNIX 4.2 systems, and installation instructions for the IBM RT PC.

4. Ultrix-32m, Digital Equipment Corporation, 1986.

A complete copy of the material in the Berkeley 4.2 UNIX manual, tailored for the Digital Equipment VAX environment.

H.1.2. Print stream standard

1. PostScript Language Reference Manual, Addison-Wesley, Reading, Massachusetts, 1985.

H.1.3. Network protocol standards

1. Internet Protocol Transition Workbook, SRI International, Menlo Park, California, March 1982.

Contains the protocol specifications for Internet Protocol (IP), User Datagram Protocol (UDP), Transmission Control Protocol (TCP), Remote Login Protocol (Telnet), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Trivial File Transfer Protocol (TFTP), all of which

are standard Athena Communication protocols.

2. Mockapetris, P., "Domain Names - Implementation and Specification," Network Working Group Request for Comments 883, SRI International, Menlo Park, California, November, 1983.
3. Mockapetris, P., "Domain System Changes and Observations," Network Working Group Request for Comments 973, SRI International, Menlo Park, California, January, 1986.

This and the previous item together specify the Athena name resolution protocol.

H.2. Related standards and publications

This subsection lists documents that are related to standards adopted by Project Athena. Some of these documents are related standards, while others provide explanatory, tutorial, or background discussions of the standards themselves. The organization of this section is parallel to that of section H.1, which cites only the adopted standards.

H.2.1. UNIX standard

1. System V Interface Definition, American Telephone and Telegraph Company, Indianapolis, Indiana, 1985.
2. The 1984 /usr/group Standard, /usr/group, Santa Clara, California, 1984.
3. Buck, David L., Reader's Guide to the 1984 /usr/group Standard, /usr/group, Santa Clara, California, December 27, 1984.
4. IEEE Trial-Use Standard Portable Operating System for Computer Environments, Draft American National Standard 1003.1, Institute of Electrical and Electronics Engineers, Inc., April 1986.

H.2.2. Print stream standard

1. PostScript Language Tutorial and Cookbook, Addison-Wesley, Reading, Massachusetts, 1985.

H.2.3. Network protocol standards

1. Mockapetris, P., "Domain Names, Concepts and Facilities," Network Working Group Request for Comments 882, SRI International, Menlo Park, California, November, 1983.

H.3. Athena Standard Third party software

This subsection lists the names and sources of third party software that Project Athena has adopted as a standard part of its software library.

H.3.1. Spreadsheet

20/20 Integrated Spreadsheet Modeling Program

Access Technology, Incorporated, Ashford, Kent, U.K.
and
South Natick, Massachusetts

H.3.2. Numerical algorithm library

NAG Fortran Library

Numerical Algorithms Group, Oxford, U.K.
and
Downers Grove, Illinois

H.3.3. Data analysis system

RS/1

Bolt Beranek and Newman Inc., Cambridge, Massachusetts.

H.3.4. Graphics toolkit

BLOX

Rubel Software, Cambridge, Massachusetts

H.3.5. Graphics library

Visual: GKS

Visual Engineering, San Jose, California

H.3.6. Graphics chart package

Visual: ProChart and Visual: C-Chart

Visual Engineering, San Jose, California

H.3.7. Editor

GNU emacs

Richard M. Stallman, Cambridge, Massachusetts

H.3.8. Text formatter

Scribe

Unilogic, Ltd., Pittsburgh, Pennsylvania

H.3.9. Data base management system

Ingres

Relational Technology Incorporated, Alameda, California

H.3.10. Algebraic manipulation System

MACSYMA

Symbolics Incorporated, Cambridge, Massachusetts

H.3.11. Mail management system

Rand MH Mail handling system (UCI version)

University of California at Irvine, Irvine, California

H.3.12. Lisp interpreter and compiler (Temporary standard, to be replaced with a Common Lisp implementation when available.)

Franz Lisp

Franz, Incorporated, Berkeley, California