This document was originally prepared using the SCRIBE document formatting program. To produce this version the SCRIBE manuscript has been ported to pages.app for the Macintosh. The format thus resembles, but is not quite identical to, the original.

Section K.3

Server Deployment Plan

by J. H. Saltzer

Scope and Status

This document describes a set of plans for deployment of the several network-based services that are needed to support Athena workstations: Storage service (RVD and NFS), database service, authentication service (Kerberos), name service (Hesiod), user location service (Zephyr), mail forwarding service, mail holding (post office) service, dial-in service, and time service. It does not cover network hostname resolution service or gateways, both of which are provided by the M.I.T. telecommunications office for the campus as a whole. In addition, the details of deployment of printer service have not yet been explored.

The status of this document is *DRAFT*; it has received only a small amount of discussion and debate. *However, time is short, and in the absence of comment, this plan will become committed by default. If you have comments, please get them to me quickly*.

Readers who have seen an earlier draft of this plan will find that the change bars in the margin identify the most significant additions and changes.

Structure

Section 1 of this document describes storage service plans—resources available, allocation, server arrangement, performance, and the transition from time-sharing to a client-server model. That (long) section is followed by several much shorter sections each laying out deployment plans for one other service. The next to last section identifies which servers, under this plan, should be assigned high-priority maintenance service. Finally, the last section reviews deployment of the entire range of services in off-campus independent living groups.

1. Storage Services

Project Athena currently operates 52 VAX 11/750's and 1 VAX 11/785, mostly as time-sharing service computers. Over the coming 9 months workstation use will replace time-sharing use, and the time-sharing service computers will convert to various kinds of storage service.

1.1. Potential storage server resources currently deployed

This note classifies the VAX systems with a label of the form (x-y) where "x" is the number of RA-80 disks and "y" is the number of RA-81 disks.

	systems	RA-80's	RA-81's
Machines assigned to users:			
(1-1) machines "client"	30	30	30
(0-3) machines "backup"	5	0	15
(0-2) machines	4	0	8
Machines assigned to staff:			
(1-1) machines "client"	9	9	9
(1-2) machines	2	2	4
(1-3) machines	1	1	3
(0-2) machines	1	0	2
(0-3) machines "backup"	1	0	3
Total in use	53	42	74
Not installed	0	2	4
Total in inventory	53	44	78

space available for allocation after leaving room for operating system:

RA-80	124 MB; 61 MB available on 42 drives		2,562 MB
RA-81	456 MB; 456 MB available on 67 drives		30,552 MB
	390 MB available on 7 drives	(no RA-80)	2,730 MB
grand total:			35,844 MB

1.2. Allocation assumptions

1. All numbers are in gross megabytes; the UNIX file system format makes approximately 80% of the gross space available.
2. Allocations for private lockers of 2 Mbytes each for registered undergraduate students. Extension to graduate student population should be done by purchase of additional disk storage space rather than by reducing the undergraduate per-student allocation.
3. Budget space for 5000 allocations, to handle 4000 students and up to 1000 miscellaneous special cases:
   - fifth-year undergraduates
   - students not officially registered
   - graduate teaching assistants working with a subject

- etc.

4. Each subject actively using Athena for teaching has an average of 40 Mbytes allocated for use in exported libraries. Assume 100 such subjects in the immediate future. (In the longer run, as development activities decline, some staff development storage should be reallocable to teaching activities.)

5. Each faculty development project has an average of 25 MBytes allocated. Assume 100 such projects. Faculty not running development projects may also have private directories. Since many of them may not actually use their potential allocation, assume 1 Mbyte average for the remaining 900 faculty members.

6. Staff members receive an average 7.5 MByte private allocation. Each staff member's allocation is on some server in the headquarters cluster; staff will no longer have a private directory on every server owned by Athena. Assume 80 staff members. Student staff have standard student lockers; those participating actively in development will have have additional home directories allocated from staff resources.

7. Up to 100 staff-initiated development activities have an average of 100 MB storage allocations, which may be subdivided to provide storage for specific development projects and individual developers with additional storage. (These numbers seem big, but their product is the actual amount of storage currently in use by the staff.)

8. Note that for purposes of starting discussion with a specific proposal, a major assumption is that the general division of resources between staff and users does not change. This assumption should probably be revisited--the overall level of staff usage of storage seems unjustifiably high.

1.3. Allocations

```
     Student lockers:  5000 allocations of 2 MB each      10,000 MB
     System Libraries: 13 copies of 516 MB each            6,700 MB
     Class libraries: 100 at 40MB each                     4,000 MB
     Class development: 100 projects at 25 MB each         2,500 MB
     non-developing faculty: 900 at 1 MB each                900 MB
     Staff : 80 allocations at 7.5 MBytes each               600 MB
     Staff development storage spaces: 100 at 100 MB each 10,000 MB

total allocated under this plan:                          34,700 MB

Unallocated (overallocated)                                1,144 MB
```

1.4. Server Arrangement

The general plan is that the five user-cluster "backup" configuration machines, each with 3 RA-81's, become file servers for faculty development and class libraries, and personal faculty files, using NFS. Ten "client" VAX's become system library servers, using RVD. The twenty remaining "client" VAX's are converted to file storage service for student lockers, again using NFS (temporarily, two of these 20 are held in time-sharing service). In the future, as dormitory networks are deployed, the machines used for student locker service by on-campus residents will move to one of two centers, one on East Campus (building 66) and one on West Campus (either W20 or W91), along with some of the system library servers. The locker machines in building 4 (Hydrovaxes) will be assigned for use by off-campus fraternity and other users, so they will remain indefinitely in their present places. The file servers for class development and class libraries could either stay in central campus locations or move to the locker centers to consolidate and free up space.

The next subsections expand on this general plan. Following that is a discussion of server and network performance, transition plans, and some longer-term considerations. Figure 1 provides an overview of the server/workstation configuration.

1.4.1. Faculty Development and Class Library Servers

All personal directories of faculty, faculty project development areas, and directories used to export class libraries are placed on five (0-3) systems each with 1300 MB, for a total of 6500 MB. (Zeus, Apollo, Clio, Ringworld, Trillian) (For this purpose, Trillian needs to be moved from E40 to the building 37 machine room.)

1.4.2. System Library Servers>

Placed on 13 (1-1) RVD library servers each with one RA-80 and one RA-81. Following is the allocation of space on a single rvd library server:

```
          152 MB vsusr          user library
          152 MB rtusr
          152 MB spusr
           10 MB vssys          root file system
           10 MB rtsys
           10 MB spsys
           10 MB vsinstall      for installation kit
           10 MB rtinstall
           10 MB spinstall
total     516 MB  (One RA-80 + one RA-81)
```

The thirteen copies of this library server are deployed as follows:

```
copies(13):    HQ development (Pollux)
               HQ prerelease (Agamemnon)
               HQ release (Andromache)
```

```
Building W91 (Hactar)
Building W20 server 1  (Zarquon)
Building W20 server 2 (Slartibartfast)
Building 66 server 1 (tbd--Mnemosyne?)
Building 66 server 2 (tbd--Polyhymnia?)
Building 1 (Persephone)
Building 4 (4-035-S)
Building 11 (Gaea)
Building 37 server 1 (Socrates)
Building 37 server 2 (Prak)
```

1.4.3. User Locker Servers

Placed on 20 (1-1) servers, each with 517 MBytes of space. Each server has 250 students assigned, for a grand total of 5000 users and little more than 10 Gigabytes. (During the first semester of use, two of these machines, Poseidon and Demeter, are temporarily assigned to continue time-sharing use.)

Nessus	Hades	Urania	Louiswu	4-035-X
Demeter	Artemis	Erato	Speaker	4-035-W
Ares	Aphrodite	Euterpe	Teela	4-035-T
Poseidon	Atlas	Calliope	Prill	4-035-U

(Note--two of these servers are (0-2) servers, rather than (1-1) servers, so there are 660 spare MBytes available.)

1.4.4. Database Service

Experience in the use of the Ingres database system suggests that the most important mode of deployment of that system is in a form that permits several members of a class to share a single data base. So that data base access coordination can be accomplished, it is necessary that all requests to a given database be handled by a single processor. To handle this requirement, two VAX 11/750 time-sharing systems (Hera and Helen) will be designated as database servers. Initially, users needing access to a database will be given a time-sharing account on this machine. However, all access will be by remote login from workstations rather than by terminals directly connected to the database servers. At some future time it is expected that client front ends will become available that are properly integrated with other network services. At that time it will be possible to use the database servers exclusively as backend machines rather than as time-sharing systems. Note that allocation of space for database storage was included in the class library category.

1.4.5. Headquarters/Staff

```
Staff cluster:

     System libraries:    3 (1-1) servers (Pollux, Agamemnon, Andromache)
     Project storage:     6 (1-1) servers (Theseus, Heracles, Orpheus,
```

```
                                        Euridyce, Priam, Hector)
                          2 (1-2) servers (Menelaus, Odysseus)
                          1 (1-3) server  (Paris)
                          1 (0-3) server  (Jason)
                          1 (0-2) server  (Achilles)
    total cluster size  14 VAX 11/750's
```

1.4.6. Loose Ends

The building 1, 11, and W91 subnets have only one RVD server each.

The following user machines are not yet assigned any role in this plan:
Oath (operations suggests use for additional RVD library service)
M4-035-V (operations suggests use for additional RVD library service)]

1.4.7. Summary of Configuration

Total facilities covered by this plan:

	VAX 11/750	RA-80	RA-81	MB
User locker servers*	20	18	22	11,130
User Library servers	10	10	10	5,170
Staff system library servers	3	3	3	1,551
Class Development and libraries	5	0	15	6,510
Database servers	2	2	2	1,034
Staff files and projects	11	9	18	8,625
spares	2	0	4	1,824
total	53	42	74	35,844

* 2 temporarily assigned to continue time-sharing service.

1.5. Performance

In addition to allocation of storage capacity, there is a need to allocate machine cycles. According to the plan above, users will have access to 10 library servers running RVD, 20 locker servers running NFS, and five class library servers running NFS. With 1500 workstations deployed, of which about 100 are in the hands of staff, 100 in the hands of faculty developers, and 200 in off-campus locations, there will be a maximum of about 1100 student workstations potentially attempting to use those storage facilities. For purposes of analysis we assume that at the busiest hour of the evening, all 1100 workstations are simultaneously in use. (We also assume that during the day, when the faculty are using workstations, there will be less student activity.)

1.5.1. RVD library servers

Every workstation that is in use will have spun-up libraries. With 1100 workstations and 10 servers, there will be about 110 workstations/server. Since cluster assignments

can be centrally managed, the average and the maximum are the same. Current performance observations on RVD suggest that 110 workstations will push the limit, but probably not overload the server.

1.5.2. NFS locker servers

Most workstations that are in use will have spun-up lockers, but the intensity of use of those lockers will be very variable. With 20 user locker servers, there will be an average of 55 users attached to each server during the busiest periods. Because there is no control on just which students are using which workstations, the number actually attached can be substantially different from the average, ranging from zero to 250 (that is the number of lockers on each server.) One would expect to see between 25 and 100 users on each server most of the time. This load is quite large for NFS, and experience will be needed to see whether the resulting response time is intolerable.

1.5.3. NFS class library servers

We assume that the average workstation attaches a class locker only 50% of the time it is active; the rest of the time the student is reading mail or doing word processing for classes that do not provide libraries of programs and data. On that basis, the five class library servers will be supplying connections to about 550 workstations, or 110 workstations per server. Again, this load is quite large for NFS, and experience will be needed to see whether the resulting response time is intolerable. Since the style of use of some class libraries is likely to be quite intermittent, one would expect that a large number of connections may not present so large a load.

1.5.4. Directory structure

Because UNIX searches directories with a linear scan, some attention must be paid to keeping frequently used directories from growing to unbounded sizes. In addition, a workstation's table of mounted (whether local or imported) file systems is searched linearly whenever a directory inode is traversed, so the total number of mounted file systems must be kept under control. The following pattern applies:

1. Each student locker server will export a single file system within which all lockers are located. Each student will have a directory in the top-level directory of that file system. This directory will thus have typically about 250 entries. When allocations are made to new students, it is important to keep the number on any given server from drifting too far from the average. (Contingency plan: if a one-level directory in locker servers proves to cost too much in performance, then a two-level directory system, with 26 directories named "a" through "z" can be interposed. There seems to be no reason to go to a three-level system with 26X26 directories, as is in the present time-sharing systems.) Note that the Hesiod name service locates locker directories by user name, so assignment of students to specific locker servers can (and will) have no connection with the user's name. Assignments can be made on the basis of most likely work location, and then on the basis of which server near that work location has the fewest lockers.

2. Class libraries on a library server will all be allocated on one exported file system, to maximize the chance that a student who imports two at the same time actually imports only one file system. Faculty development project directories will be allocated on a different exported file system of the library server.

3. The workstation will mount all imported file systems on temporary nodes in the local directory named "/mit". The typical workstation will import two system libraries, the student's private locker directory, one to four class libraries, and perhaps one or two other items of interest such as the SIPB library, the personal directories of a few friends, or the common working area of a student activity. Under this plan, the number of imported file systems will typically be ten or less.

1.5.5. Network traffic

The final performance consideration is the amount of long-distance network traffic generated by users whose lockers are in a server on a different network from the one attached to the workstation being used. The current workstation and server deployment plan leads to the following configuration at the end of the summer, 1987:

network building	workstations	rvd servers	class servers	locker servers
W20	59	2	0	5
66	24	2	1	3
4/16	87	1	1	4
9/11	67	1	1	2
1/7	64	1	1	3
37	44	2	0	0
W91	7	1	0	0

If all 352 of the above workstations are in use at once, and lockers are scattered at random, on would expect about 60 (17%) of the workstations to be using a locker on their own network; the remaining 83% would be working through gateways.

Fortunately, the initial configuration has only about 1/4 the ultimately planned number of workstations, so the total traffic will be about 1/4 the later expected load. Thus it should be possible to observe the traffic in this configuration and use the results of those observations to help develop network configuration changes needed as more workstations are deployed. There are four reconfiguration strategies that may be usable to help keep network traffic from overloading the gateways:

1. Place some locker or class servers on more than one network.

2. Combine some networks using fiber repeaters.

3. Move some locker or class servers into different machine rooms closer to their clients.

4. Discover patterns of use and move user lockers to servers near the workstations they usually patronize.

The section below, on longer-term plans, makes some specific proposals along these lines.

1.6. Transition from Time-Sharing

The present time-sharing systems are home to a large number of student and faculty files, many of which should end up in new locations after the conversion to the new server arrangement. To ease this transition, several interim measures are planned:

1. The transition will be accomplished during the summer, 1987, a three-month period when usage is relatively light. One reason to complete the transition during this period is to assure that it is not necessary to run both a locker storage system and a large fleet of time-sharing systems at the same time.

2. Initial student locker allocations will be 1 Mbyte, rather than 2. Each locker machine will thus be allocated only to 50% of capacity. Students requiring a larger allocation can apply for it and will receive it if there is space available on their server. (And see the next item.)

3. Two of the 20 machines allocated for student locker use will temporarily be assigned to continue to run time-sharing, for those classes that cannot cope with the transition in time for the start of the Fall semester. (Tentative designation for this purpose: Demeter and Poseidon.)

4. Machines that undergo the transition from time-sharing to NFS service will be updated to the latest system release (the first to contain NFS service facilities) and reconfigured for dedicated NFS service. The dedicated NFS reconfiguration is essentially identical to the dedicated RVD configuration: a minimal set of commands and libraries, and RA-81 disks with a single file system occupying a full-disk partition.

5. An SMS-related tool is needed that allocates student lockers to storage servers and triggers directory creation on the server. Normally that tool would allocate space for incoming freshmen; this summer it must be used to create all allocations.

6. All present student center files (on Prill, Teela, Speaker, Louiswu, and Nessus) will be dumped to tape and reloaded into five separate file systems on Ringworld, where they will be exported as read-only NFS file systems (with the name of the file system being the name of the machine from which the file system came). Students will be expected to copy their files from Ringworld to their lockers during the fall semester. (If time and energy are available, it may be possible to perform much of that copying by automatic scripts in August.) (Read-only export is accomplished by mounting the file systems read-only on the server machine. If that method causes NFS to fail, then RVD can be used to accomplish read-only export.)

7. The remaining time-sharing systems will be cleared of student accounts as usual at the end of the Spring semester. Files remaining in faculty development accounts will be dumped to tape and restored to the backup server machine in their cluster, in a separate directory for each originating machine. Again, these file systems will be exported as read-only NFS file systems, or as read-only RVD file systems if necessary.

8. Both the old student center file systems and the old faculty development file systems will remain on-line until December, 1987, when their storage will be reclaimed.

9. It will probably be necessary to continue to run Helen as a class library server exporting RVD lockers through the fall semester, until all present users have converted over to the newer NFS model. At that time it can assume its new role as a database server.

1.7. Longer term plans

As the number of workstations located in dormitory rooms increases, the user locker servers should be physically moved to machine rooms closer to the dormitories, forming East Campus and West Campus locker centers. (The West Campus locker center, being larger, may actually be split between buildings W20 and W91.) Under this plan, students would normally have their lockers assigned in the locker center nearest their dormitory, and to a locker server that serves primarily their dormitory.

Dormitory networks would have a tail that extends to the locker center, so that a locker server there can be on the same physical network without going through a gateway. (N. B. a limit of 256 IP addresses on one physical network will start to be a nuisance here; it will probably be necessary to begin using the IP "rest" address byte.)

When working in a main-building workstation cluster, use of lockers will be via gateways.

Physical movement of locker servers may not actually be required, depending on the timing of expansion of Athena to handle graduate students. It may be more appropriate to install new servers in the locker centers and reassign space on the older central-campus servers to the gradute student population.

Class libraries will be in main-building machine rooms rather than in the East Campus or West Campus locker centers. If traffic gets too great, copies of libraries for popular classes may be placed in the locker centers.

At the locker center, each system library server would have multiple network interfaces, and thus serve several dormitory networks, one as primary and the others as backup. Locker servers are attached to the local network of the dormitory that they serve. The gateway from the dormitory to the campus network is located at the locker center.

1.8. General observations

1. Present VAX 11/750's can handle a maximum of 4 disk drives without purchase of additional interface hardware.

2. Every private locker is named with the owner's Athena name.

2. Service Management System

The Athenareg system will move from its current location on a VAX 11/750 to a large-configuration VS-II workstation located in E40 (taking the name Hector with it), where it will manage only the remaining VAX 11/750 time-sharing systems. At the time of the demise of the last time-sharing system, that workstation will be recycled into a different service. The newly-deployed Service Management System will also operate on a VS-II workstation located in E40.

3. Authentication Service (Kerberos) and Time Service

Kerberos is organized so that one host acts as a master server and any number of other machines can act as slave servers, receiving regular updates of their database from the master server. The deployment of slaves is a compromise between security and availability. Both the master and each slave represent a place at which all passwords can potentially be compromised, so their physical security must be very high. To accomplish this high security, both the master and all slaves will be dedicated MicroVAX-II processors. The master will be located in a locked office in building E40, and three slaves will be located in separately secured screen rooms inside locked Athena machine rooms in buildings 4, 11, and W-91. This number of slaves is sufficient to assure high availability of at least one at all times that spine network continuity is available.

Because Kerberos tickets contain time stamps and Kerberos authenticators have relatively short timeouts, it is helpful that workstations use the same time reference as the Kerberos servers. For this reason, the Kerberos servers also operate primary time servers. In addition, the master Kerberos site runs a program that polls the slave time servers once every five minutes to verify their accuracy. If any slave drifts by more than one minute from the master's value the master forwards a trouble reporting message to the staff responsible for Kerberos operations.

4. Name Service (Hesiod)

There will be five Hesiod name servers, running initially on Zeus, Apollo, Clio, Ringworld, and Trillian. As load grows, it will probably be necessary to move these name servers to less heavily loaded hosts such as the mail holding servers or to a set of MicroVAXes II or RT's dedicated to name service. A dedicated MicroVAX II or RT should be able to handle about 40 name requests per second, so three can handle about 6,000/minute. If a user makes an average of one name request per command typed, and types an average of two commands per minute, 1,000 active users would use about 33%

of that capacity. The number of dedicated name servers is thus determined more by an availability requirement than a capacity requirement.

5. User Location Service (Zephyr)

(Not yet designed; deployment plan yet to be determined.)

6. Mail Forwarding and Holding Services

6.1. Capacity estimate

Assuming each user receives 6 messages /day, 4,000 users would receive about 24,000 messages per day. If receiving an average message and then forwarding it or delivering it to the workstation requires 4 seconds of attention of the server, a single server could handle up to 20,000 messages/day.

6.2. Proposed Configuration

The mail forwarding service presently operated on the dedicated MicroVAX II named Athena.MIT.EDU seems to have more than enough capacity to handle the rate of messages currently flowing within Athena. When demand grows to exceed its 20,000 messages/day forwarding capacity, one additional, identical forwarder will be installed. Mail originating within Athena will be split randomly between the two forwarders; outside mail targeted for Athena users will continue to come to the one that advertises its name as Athena.MIT.EDU.

Mail storage service is handled by a set of post office servers, each being a dedicated large-configuration (BA-123) MicroVAX II with 210 Mbytes of disk, of which 170 Mbytes are available for mail storage. Allowing for an average of 100 Kbytes (25 screen-pages of messages) of not-yet-picked-up mail per user for 4000 users requires about 400 Mbytes of storage, so three such dedicated systems will be required. (Alternatively, three RT PC's configured with 200 or 400 Mbyte disk drives could be used.) Although the processor capacity of these three systems is about 2.5 times the expected load, one expects to see peaks in mail pickup patterns, so this capacity may be marginal. Careful monitoring of mail pickup delays during peak periods will be required to determine whether or not more servers should be deployed. In addition, it will be important to monitor the rate of accumulation of non-yet-picked-up mail to see if the space allocation suggested here is realistic.

6.3. Printer Service

(Plan yet to be developed, in conjunction with workstation deployment plan.)

7. Dial-in Service

Even though many workstation applications take advantage of large displays, graphics, and other features of X, there will probably be some residual demand for

ability to dial in from an ASCII terminal located off-campus, to do word processing, mail reading and some program development. To meet this requirement, a dial-in server will be configured. A dial-in server is a VS-2 or RT workstation configured for time-sharing service with 8 RS-232 ports. Workstation login will work as usual, requiring Kerberos authentication and attaching the identified user's NFS locker. In this configuration the toehold activation and deactivation system will activate and deactivate the workstation on first login and last logout.

Three such dial-in servers are required, located in the places where dial-in lines currently are attached to Zeus, Ringworld, and Jason.

To provide both RT and VS-2 service, both kinds of workstations will be needed as dial-in servers. A future development, still under discussion, would consist of deploying rlogin servers of both architectures and modifying the dial-in server to ask the user what kind of machine is needed and perform a relayed rlogin. In this mode a single dialin server could probably handle 16 or 24 lines. The rlogin servers would also be usable for rlogin access to Athena resources from other network locations.

8. Priority Maintenance

Under this plan, a small number of servers are "critical", which means that they should be subject to 24-hour hardware maintenance contracts, because outage would disable all (or a major part) of Athena's service. The servers that have this property are the five designated to provide NFS export class libraries (Clio, Ringworld, Apollo, Zeus, and Trillian.) Since the RVD servers are deployed in identical pairs, the frequency of service loss should be low enough to allow those machines to operate with normal maintenance. Any one locke server handles the files of only 5% of the user population, so those servers do not require special attention. Finally, name, post office, and Kerberos servers are backed up with a small number of "hot spares" that can be quickly reconfigured to deliver any of those services, 24-hour hardware maintenance is not required on them either.

9. Off-Campus Independent Living Groups

There are several off-campus independent living groups that have a configuration consisting of a few workstations, one or more large-configuration workstations intended for use as servers, and a low-speed (9600 baud) telephone link to the campus network. Both because the speed is low and because the reliability of a telephone link is problematical, the configuration of the living group is designed to allow operation independent of the campus network and administration.

9.1. Storage service

In the off-campus sites, individual workstations have 70 Mbytes of disk space, and it is assumed that users are assigned home workstations on which to store their files. In addition, each workstation can export its user file system with NFS. Thus lockers for private storage are unnecessary. Class and system libraries are still required.

Assume that off-campus independent living groups have only one variety of workstation, so only one copy of the system libraries (/usr, etc.) is needed in a local server. In this case, a large-configuration workstation with 210 Mbytes of disk storage allows room for 26 Mbytes of operating system, 152 Mbytes of system library, and 32 Mbytes for class libraries.

Especially in small living groups, it may turn out to be appropriate to trim the vsusr/rtusr library of facilities that are less useful to the local community, in order to make room for more class libraries. In addition, in a group with a small number of workstations, library export can be done using NFS rather than RVD, thereby minimizing loss of space from breakage in allocations among multiple file systems.

9.2. Other services

In addition to providing class and system library storage service, the large-configuration workstation used as a server machine operates several other services.

1. A locally-administered copy of Kerberos that contains passwords only for members of the living group.
2. A time server.
3. A name service providing authoritative data for a separate Hesiod domain for the independent living group.
4. A post office server, holding mailboxes for each user of the independent living group.
5. A mail forwarding server, which provides a link between the local post office and the on-campus mail forwarding system.

Until a design is available for Zephyr, the user location service, it is not yet apparent whether or not it is appropriate for a copy of it to be run in the independent living group environment.

Figure 1--Fall, 1987, Network Layout of Athena Non-Staff Workstations and Storage Servers

	bldg W91	bldg 10	bldg 37	bldg e40
rvd	Hactar		Prak	
			Socrates	
class			Trillian	
DB				Helen
work stations	Next(7)	10-0xx(2)	37-3xx(40) 35-225(4)	
ws count	7	2	44	

	bldg 66	bldg 16/4	bldg 11/9	bldg 1/7	bldg w20
rvd	Mnemosyne Polyhymnia	M4-035-S M4-035-V	Gaea	Persephone Oath	Slarty Zarquon
class	Clio		Apollo	Zeus	Ringworld
locker	Erato Euterpe Urania Calliope	M4-035-T M4-035-U M4-035-W M4-035-X	Atlas Aphrodite Artemis	Hades Ares	Teela Nessus Speaker Louiswu Prill
DB				Hera	
TSS				Poseidon Demeter	
work stations	66-080(24)	4-035(23) 4-117(1) 16-034(20) 2-225(8) 2-023(3) 2-167(8) TDC (24)	11-1xx(56) 9-5xx(11)	1-2xx(9) 3-4xx(24) 7-3xx(7) 1-142(24)	w20-575 (59)
ws count	24	87	67	64	59

total workstation count: 354 (Does not include 130 in headquarters and 36 in off-net locations.)