Early Adopters
An Internet 2 Middleware Project

Jay Graham
University of Pittsburgh
Computing Services & Systems Development
717 Cathedral of Learning
Pittsburgh, PA 15260
(412) 624-5244
jwg@pitt.edu

Jeffrey Cepull
University of Pittsburgh
Computing Services & Systems Development
419 Bellefield Hall
Pittsburgh, PA 15260
(412) 648-7225
cepull@pitt.edu

1. INTRODUCTION

Early Adopters, the campus testbed phase of Early Harvest, is a group of eleven institutions of higher education working to provide a testbed for the deployment of middleware technologies. This will result in early implementation for campus networks, and in a roadmap for other universities and colleges to follow. Early Adopters is sponsored by Internet2 with funding from the National Science Foundation.

Goals of the Early Adopters Project

Primary Goals

- facilitate campus deployments of core middleware technologies
- identify reasonable technical and policy approaches, and design issues and factors, that influence institutional selection of a particular approach enrich the technical contents of Early Harvest
- inform the larger community — for example, higher education more generally, NIH, and NSF — of requirements for middleware deployment and interoperability

Secondary Goals

- development of the EDUPerson
- adherence to IMS standards
- explore medical middleware issues
- generic — how is this expressed in the core deployment?
- specific — what medical data structures need integration into campus environments?
- outreach to encourage other institutions
- research options for authorization services
- evaluate new tools and technologies

LEAVE THIS TEXT BOX IN PLACE
AND BLANK

Keywords
middleware, IMS, LDAP, interoperability and EDUPerson

2. EARLY ADOPTERS - AREAS OF ACTIVITY

The items included under the heading of middleware differ depending on who is making the list. Many interesting categorizations exist. These categorizations are all centered around sets of tools and data that help applications use networked resources and services. Some services, like authentication and directories, are in all categorizations. Others, such as co-scheduling of networked resources, secure multicast, and object brokering and messaging, are the major middleware interests of particular communities, but attract little interest outside of those particular communities. A popular definition of middleware that reflects this diversity of interests is "the intersection of the stuff that network engineers don't want to do with the stuff that applications developers don't want to do."

Middleware has emerged as a critical second level of the enterprise IT infrastructure. The need for middleware stems from growth in the number of applications, in the customizations within those applications and in the number of locations in our environments. These and other factors now require that a set of core data and services be moved from their multiple instances into a centralized institutional offering. This central provision of service eases application development, increases robustness, assists data management, and provides overall operating efficiencies. Interoperable middleware between organizations is a particular need of higher education. Researchers need to have their local middleware work with that operated by national scientific resources such as supercomputing centers, scholarly databases, and federal scientific facilities and labs. Advanced network applications will transform instructional processes, but

they will depend on middleware to function. The fact that higher education is fractal in structure will create markets that need interoperable standards and products.

3. TAXONOMY

Core middleware services are those that all other middleware services depend on. The challenges in providing these services are as much political as they are technical. Many of the hardest issues involve the ownership and management of data in the complex world of higher education.

Identifiers. An identifier is a character string that connects a real-world subject to a set of computerized data. Identifiers were simple when each person had exactly one. Now people generally have several identifiers, and identifiers apply not only to people, but also to group of people, or to objects (or groups of objects) such as printers and applications. Thus the relationships among a subject's identifiers, and policies associated with the assignment of identifiers, become important issues.

Authentication. Given the breadth of interactions that are now computer-assisted, establishing that a particular request is associated with a specific real-world subject becomes critical. The traditional approach of login and clear text password is far too insecure and inflexible for the variety of ways that clients need to authenticate to servers.

Directories. Much of the information about real-world subjects needs to be contained in a general-purpose, high-performance server that can respond to application requests for information. There are substantial technical and political issues in the development and operation of a directory service. Technically, determination of the elements of the directory (the schema), the ways of addressing the elements (the namespace), and operational issues such as replication and partitioning need to be addressed. Applications must be reengineered to use the directory. Policy issues include ownership of data, feeds into and out of the directory, and setting permissions to read and write data.

Authorization. An important subset of the information about a real world subject is what it is permitted to do. Authorization can range from allowing access to refined controls of a remote electron microscope to permissions to place purchase orders below a specified level on an institutional account. Defining these rules, including means to delegate or reassign authority on a temporary basis, as well as delivering this information to applications, are some of the challenges in this newly emergent area.

Certificates and PKI. Below the core middleware services, at the boundary of the network layer, lie a number of services that can be classified as middleware-based networking or networking-oriented middleware. These services include:
Secure multicast. This is multicast extended to permit, at the network layer, secure access to join a multicast session. Bandwidth brokering. This is a service that securely allocates quality of service (QoS) to various applications and users within an institution or organization.

Typically these services require core middleware services, such as identifiers, authentication and directories, in order to operate. Above the core middleware services are a number of types of application-oriented middleware, or upperware. A rough grouping of such middleware would include:

Services for ubiquitous computing. Higher education needs a variety of open protocols and implementations that allow students to access their bookmarks and aliases from any location, as well as institutional and multiorganizational file systems to enable sharing and support collaboration tools.

Support for research computing. Efforts are underway to transform scattered national computational resources into a coherent grid, providing researchers consistent access across a variety of architectures, permitting co-scheduling of resources, coupling data, networking and computing together.

Support for administrative computing. The new generations of business systems have loosely-coupled components that depend on a common applications infrastructure, which provides services such as object brokering for component requests, message handling between components, and monitoring of transactions.

Again, these services depend on core middleware components in order to operate. In turn, as these areas continue to evolve rapidly over the next few years, new utilities may be developed within the core to support them.

4. EARLY ADOPTERS FAQ

What is the purpose of the campus testbed phase of Early Harvest?

Early Adopters serves three purposes: to help some campuses advance the state of their core middleware infrastructure; to generate additional best practices, particularly in process management, for inclusion in the knowledge base; and to advise the NSF on the issues and challenges in deploying middleware within higher education and research.

What is the structure of Early Adopters?

Early Adopters will begin with a two-day workshop (a limited amount of travel reimbursement is available) for participants, to be held in early December. The workshop will include extensive discussions of the results of the Early Harvest technical workshop (held in late September), implementation options, campus process discussions, and creation of planning materials.

During the next several months, campuses should move forward on their middleware design and deployment initiatives. To assist them, there will be a number of information-sharing mechanisms to support the campuses in their work. These will include biweekly conference calls, technical briefings, and access to consulting assistance.
In early spring, the participants will be reconvened in a second workshop, to discuss successes, challenges, and next steps, and to gather additional material for the Early Harvest best-practices guide.

What are the obligations of a participating campus?

The campus must commit to pursuing the design and deployment of core middleware services (coherent identifier management, authentication, and directory services). The campus also agrees to participate in the sharing of technology and process best practices, including involvement in the two workshops, biweekly conference calls, and the aggregation of best practices at periodic intervals.

What are the benefits to the campus?

The campus will get an early opportunity to implement an infrastructure that can be highly leveraged. The campus will be part of a well-supported initiative in this regard, with discussion and consulting opportunities. The leadership and contributions of the participants will be widely acknowledged.

What are the costs for a campus?

In building enterprise middleware services, there will be considerable requirements for time commitments from senior management within a number of units around campus. There will be the usual project costs of acquiring equipment and software. As is the case with any leadership activity, there is the risk of misdirection, limited economies of scale, and a harder overall effort than for those who follow later.

What is the role of Internet2 in Early Adopters?

Internet2 will provide expertise, coordination, and some limited funding support. The expertise will include materials gleaned from the Early Harvest technical workshop, national experts, and ongoing developments in core technologies. Coordination will include operation of the meetings and biweekly conference calls, brokering information needs among participants, and culling materials for the best-practices knowledge base. Funding will include the costs of the workshops and conference calls, and some limited travel reimbursements for participants.

Who needs to be on the campus team?

The campus team should include technology developers and technology support people, applications developers (including administrative systems, instructional applications such as web course systems, and basic services such as email and printing services), policy makers (including university legal staff and senior management) and key data providers (including the Registrar and Faculty and Staff Personnel). The campus team is intended to provide overall project oversight, obtain institutional commitments, and involve key constituencies.

How does a campus apply for participation in Early Adopters?

A campus needs to submit a brief (2-3 page) application letter to earlyadopter@internet2.edu by November 12, 1999. The letter should address the following issues:
Technical resources available to the effort, including central IT staff and campus applications developers. Involvement and commitment from the major data owners on campus, including the Registrar, Faculty and Staff Personnel, and other key institutional informational resource providers. Involvement and commitment from the major institutional policy makers, including senior management and university legal offices. Existing technical infrastructure, including unified campus name space, authentication deployments, and central directories or integrated data warehouses. Existing policy infrastructure, including specific guidelines for who has electronic access for major campus IT resources (e.g. network, accounts, email, libraries).

How will participants be selected?

Proposals will be evaluated primarily on 1) the strength of the campus commitment to pursue deployment of a core middleware infrastructure, and 2) its readiness in related technical and policy areas. In addition, selections will be made to maximize the diversity of institutions participating, in order to obtain the broadest possible gathering of best practices in deploying middleware in higher education.

Where can I find additional information?

Background material is available from the Internet2 middleware and Early Harvest pages. In particular, a copy of the original NSF proposal is available, as is the Middleware 101 presentation given at the I2 fall member meeting, which describes some of the specific issues that a campus may encounter in deploying core middleware.

5. CORE MIDDLEWARE

These five services are central to middleware as a whole.

5.1. Identifiers

A set of computer-readable codes that uniquely specify a subject.

An identifier is a function that maps real-world subjects into name or character strings, so that distinct subjects have distinct strings. A real-world subject may be a person, an object (for example, a printer or a file), a group, or a department. A real-world subject can have multiple identifiers. For example, a person may have a Social Security number, an email address, userids on several systems, a network ID, and others.

Identifiers have always been part of the campus IT environment, but until recently their use was relatively narrow and limited. As the number of computing and networked resources has proliferated, so too have identifiers. With the growing importance of these resources, issues of rights and responsibilities associated with each identifier become critical.

The key issues are assigning identifiers (How are they formed? Who hands them out? How long are they good for? Can they be reused? What resources are they valid for?) and relating identifiers (Are some dependent on others? Can an effective mapping be made among a real-world subject's set of identifiers?).

5.2. Authentication

The process of a subject electronically establishing that it is, in fact, the subject associated with a particular identity.

Authentication is the process of establishing whether or not a real-world subject is who or what its identifier says it is. Identity can be proven by:

- Something you know, like a password
- Something you have, as with smartcards, challenge-response mechanisms, or public-key certificates
- Something you are, as with positive photo identification, fingerprints, and biometrics

Authentication should be secure. It is the atomic service that enables all activities in the networked world. Authentication should be accessible to any application that wants to use the service. Implementing single sign-on, to whatever extent possible, has real benefits to both users and the overall IT environment. Authentication should be efficient; it should not tax the resources of either the system or the user. Authentication should be effective. Applications should not have to be customized to use alternative authentication schemes.

5.3. Directories

Central repositories that hold information and data associated with identities. These repositories are accessed by people and by applications to, for example, get information, customize generic environments to individual preferences, and route mail and documents.

Directories are the operational linchpin of almost all middleware services. They can contain critical customization information for people, processes, resources and groups. By placing such information in a common storage area, diverse applications from diverse locations can access a consistent and comprehensive source for current values of key data. In future information technology environments, directories will be among the most critical services offered.

Directories are databases that are optimized for reads, and that contain key institutional and personal data for use by a wide variety of applications. Directories need ways to describe the sequence of fields in the database (a schema), the names of the fields (a namespace) and the contents of the fields (attribute values). Directories also need indices into the database (identifiers).

Examples of fields in a directory include institutional status, bookmarks, email aliases, personal photos, permissions, private keys and calendars. Identifiers to access directory mail include social security number, public key certificates, unique ID, and email address.

5.4. Authorization

Those permissions and workflow engines that drive transaction handling, administrative applications and automation of business processes.

Of the current components of core middleware, the least developed and most amorphous is authorization services. It is definitely a service rather than a server — authorization functionality will be provided coherently through several means of delivery, including authentication, directory servers and certificates.

Examples are legion, which is what makes this area so important. Authorization will be the basis of workflow. It will drive permissions for accessing networked resources, allow us to control and delegate electronic responsibilities, and serve as the basis for future administrative applications. It will allow us to convert our complex legal policies into automated systems in a easily scalable fashion.

At its simplest, authorization is the next generation of ACLs — the read/write/execute controls that are embedded in file systems. Typically, authorization indicates what an identifier, properly authenticated, is permitted to do with a networked object or resource.

There are many challenges associated with authorization, including :

- Where to store the authorization characteristics
- How to transport those characteristics to applications
- How to ensure consistent meaning and validity to values associated with those characteristics
- How to effectively express the sophisticated and diverse characteristics implicit in policies in an processable list of attributes

There are several places to store authorization characteristics. Most often, they are kept in directories, either system-specific or as part of a campus-wide infrastructure. Alternatively, they can be stored within a file system, as a separate data system, or on an external device (such as a smartcard).

Transporting the characteristics to the application can be done in several ways as well. Applications can be periodically updated from a standalone authorization server or request authorization dynamically from the server via an RPC. Alternatively, the user can present authorizations to the application as part of the authentication process. For example, the authorizations can be carried within the Kerberos ticket or as part of a certificate. In order to assist consistent assignments of values within authorizations, a number of technological tools are useful. For example, default settings and inherited values help reduce the discretion of the authorizer. Similarly, providing easy ways of delegating permissions to authorize is an important feature.

The need to translate complex policies into automated combinations of more basic attributes has led to research into policy models and policy description languages. These tools are receiving some attention within IETF as they have significance for network layer controls as well.

5.5. Certificates and public-key infrastructures (PKI)

Certificates and PKI are related to the previous four core middleware services in several important ways.

There is considerable interest in the use of X.509 certificates to address a number of network computing needs in higher education. The technology itself is powerful and elegant, but there are several major challenges to the widespread successful use of certificates. This page discusses some of these issues.

The software, protocols and legal agreements that are necessary to effectively use certificates combine to form a Public Key Infrastructure (PKI). A PKI has several components.

- A Certificate Authority (CA), that manages and signs certificates for an institution
- Registration Authorities, operating under the auspices of the CA, that validate users as having been issued certificates
- PKI management tools, including software to manage revocations, validations and renewals
- Directories to store certificates, public keys, and certificate management information
- Databases and key-management software to store escrowed and archived keys
- Applications that can make use of certificates and can seek validation of others' certificates
- Trust models that extend the realm of secure communications beyond the original CA
- Policies that identify how an institution manages certificates, including legal liabilities and limitations, standards on
- contents of certificates, and actual campus practices

Among the potential uses for certificates are individual authentication, email encryption, digital signatures, and access controls. Each of these uses can place different requirements on the PKI components. For example, private keys for encryption may be escrowed, while private keys for signatures may not be.

6. PARTICIPATING INSTITUTIONS

- Dartmouth College
- Johns Hopkins University
- Michigan Tech University
- Tufts University
- University of Hawaii
- University of Maryland. Baltimore County
- University of Memphis
- University of Michigan
- University of Pittsburgh
- University of Southern California
- University of Tennessee, Memphis

The EA partners have been holding regular conference calls and meeting to accomplish several goals. Those goals include establishing a uniform approach and format for a scoping document that would be available to any institution trying to develop a strategy for integrating middleware. Best practices, especially those related to process and project management, will also be identified The EA partners will also be advising the NSF on issues, obstacles or opportunities for deploying middleware at higher educational sites.

An overview of each partner's stage of development, deployment or implementation will be presented. Common issues, barriers and successes will be summarized.

7. MIDDLEWARE AT THE UNIVERSITY OF PITTSBURGH

A detailed explanation of the impact of middleware on the University of Pittsburgh campus will be presented. Particular attention will be devoted to the role of middleware and central directories in the development and strategy for the University of Pittsburgh's new Accounts Management System. The presenters will report in detail on the current state of the system and the planned set of features that will be made available to the university community over the next 12-18 months, many of which have middleware at their core.

8. MIDDLEWARE FAQ

What is middleware?

The term middleware is used to describe a broad array of tools and data that help applications use networked resources and services. Some tools, such as authentication and directories, are in all categorizations. Other services, such as co-scheduling of networked resources, secure multicast, and object brokering and messaging, are the major middleware interests of particular communities, such as scientific researchers or business systems vendors. One definition that reflects this breadth of meaning is "Middleware is the intersection of the stuff that network engineers don't want to do with the stuff that applications developers don't want to do."

Why is middleware important?

Middleware has emerged as a critical second level of the enterprise IT infrastructure, between the network and application levels. The need for middleware stems from the increasing growth in the number of applications, in the customizations within

those applications, and in the number of locations in our environments. These and other factors now require that a set of core data and services be moved from their multiple instances into a centralized institutional offering. This central provision of service eases application development, increases robustness, assists data management, and provides overall operating efficiencies.

Why is middleware urgent?

There are several drivers bringing middleware to campus. Advanced scientific computing environments such as PACI are placing requirements on campus researchers for middleware services such as authentication and directories. Library projects such as the UCOP/Columbia certificate project will be extending across a broader higher education community. The Federal government is preparing requirements for digital signatures for student loan forms. New versions of software, such as Windows 2000, come with the tools to build ad hoc middleware components. It is urgent that campuses build a coherent infrastructure to respond to these drivers.

What makes the higher education and research communities distinctive in their need for middleware?

Many companies and other communities of interest are coming to understand the importance of middleware to their missions, and are proceeding with development. Higher education faces unique technical and policy issues in its deployment. Technical issues include the mobility of students, the diversity of equipment, and the requirements of advanced applications. Policy issues include ownership of data, FERPA and other public records issues, and extended collaborative relationships. Together these considerations make middleware deployment within higher education significantly harder than deployment outside of it.

When middleware becomes part of the IT environment, how critical will a robust infrastructure be?

The middleware components of the future IT environment will be every bit as critical as the underlying network infrastructure, requiring 24x7 service, high performance, and appropriate redundancy. Directory services will receive millions of hits per day; identifiers will have explicit control mechanisms; attribute services will be invoked by almost every application on campus. In addition, lawyers will place strict operational constraints on security services.

Is middleware a centralized or a distributed issue on campus?

It is both. Like network services on campus, there is a need for a consistent infrastructure across campus that is best provisioned centrally. At the same time, many parts of the contents of this infrastructure are best maintained by the individuals themselves, and by their departments. The trick is to create a centrally coordinated service that provides tools and authority for distributed management of the contents.

Aren't we going to get middleware from the commercial marketplace?

It is certainly the case that many basic middleware products that higher education will deploy will be commercial products. These products will come both from diversified software companies such as Microsoft and Novell, and from providers of more specific products, such as Netscape, HP, and ATT. At the same time, a number of distinctive characteristics of the higher education community create design considerations that require complex implementations. In addition, the research side of the academic enterprise needs additional discipline-specific middleware that will probably not attract much commercial interest. Finally, the collaborative nature of higher education will raise interoperability issues that must be addressed within the community.

What kind of investments will campuses need to make?

Like networking, middleware will require considerable commitments of time and money. However, the types of costs are different.
Networking has required large sums of capital (for fiber, routers, switches, etc.) and considerable operating costs (for external access, maintenance, etc.) Personnel costs have been relatively modest. For middleware, the hardware costs (servers, readers, etc.) are likely to be relatively low. Software costs are unclear now, but there are clearly considerable expenses in building bridges to legacy systems and to evolving middleware-enabled applications.

Unlike networking, middleware has a second major cost component: process time. A campus must develop consensus and support for the deployment of middleware, clarify data ownership and management issues, specify relationships among individuals, groups and information technology objects, establish legal agreements, and change the way that information is managed on the campus.

How does the Internet2 Middleware Initiative intend to address these needs?

Efforts will focus on advancing the level of middleware within higher education. A set of related activities will include fostering technical standards, aggregating and disseminating technical design and implementation strategies, fostering opportunities for vendors and Internet2 members to shape and deploy products, and integrating efforts with specific scientific and research communities.

What should campuses be doing now?

It is not too early for campuses to begin the processes that address the policy side of the challenge, building awareness about the need for middleware, identifying key constituencies that will be involved in the process, and taking basic inventories of the data and management relationships on campus. At the same time, experimentation in the core technologies, most notably in directory services, should be undertaken.

9. REFERENCES

Early Adopters Middleware Project by Kenneth J. Klingenstein, Project Director, Internet2 Middleware Initiative and Chief Technologist, University of Colorado at Boulder. Copyright 1999.

Identifier references:

http://www.stanford.edu/group/itssccs/project/registry/person_registry/attributes/

http://consult.stanford.edu/pub/internet/netinfo/i-d/draft-ietf-ldapext-authmeth-02.txt

http://www.stanford.edu/group/networking/directory/PubliclyUniqIdentV1_0.html

http://www.stanford.edu/group/itssccs/project/sunetid/sunetid.design/sunetid.requirements

http://www.arl.org/newsltr/194/identifier.html.

Authentication references:

Kerberos: A Network Authentication System by Brian Tung (The Addison-Wesley Networking Basics Series).

Authentication Systems for Secure Networks by Rolf Oppliger (Artech House Computer Science Library).

Directory references:

Understanding and Deploying LDAP Directory Services, by Howes, Smith and Good (Macmillan).