MIT Enterprise Architecture Guide

Introduction

The MIT Enterprise Architecture Guide (EAG) documents MIT's architectural principles and goals, the current state of MIT's enterprise architecture, and a future state architectural vision. The EAG also includes information regarding the ITAG architecture review process. Since this document serves to inform developers about available enterprise tools and services, we expect the EAG will be useful to enterprise system developers across the institute. Because this is a constantly evolving document, community feedback will drive future ITAG agendas and thus influence content in future versions. ITAG expects to update the EAG on a quarterly basis.

Audience

The intended audience of the EAG includes project teams making enhancements to existing systems, project teams developing new systems, sponsors of initiative, ITAG Members, and DLC Leadership. Each group can benefit in a different way from the EAG as detailed below.

- **Project Teams:** Project teams can use the EAG to gain an understanding of the current architectural landscape, the future vision of the enterprise architecture, and the services available to development teams. By understanding the recommended technical standards and available MIT services, project teams can re-use existing services and create applications that fit into the long term architectural vision. Teams can also leverage the information to develop new enterprise-wide services. Finally, the EAG will assist project team members in identifying whom to contact to mitigate risks in different aspects of their project.
- **Sponsors:** Sponsors can benefit from the EAG by gaining an understanding of the technical direction of the Institute as well as the Architectural Governance Process. This knowledge can then be used to shape their decisions regarding IT investments.
- **ITAG:** ITAG members can use the EAG to gain a common understanding of the Enterprise Architecture at MIT. Additionally, the EAG will be used during the project review process to provide a consistent representation of the context and principles of both the current and future state. Both of these items will assist ITAG in making informed architectural decisions as well as identifying gaps in the Enterprise Architecture.
- **DLC Leadership:** DLC Leadership can use the EAG to gain a common understanding of the Enterprise Architecture at MIT, to proactively identify potential risks in projects, and to assist in identifying individuals to help mitigate those risks.

How to use the EA Guide

The EAG is divided into the following sections: Introduction, Context and Principles, Current State, Key Systems Overview, Future State, Project Review Process, and Moving Forward. As a user of the EAG, you can read it in its entirety or reference a particular section that pertains to your need.

Whom to Contact with Questions

While the Enterprise Architecture Guide provides a breath of knowledge on the Enterprise Architecture at MIT, it may not contain answers to the queries you have. If you have specific questions regarding items presented in this guide you can contact ITAG with them via email through itag@mit.edu. If your questions are in the various business or service areas discussed, you can contact the appropriate owner listed in the contacts section in the appendix.

Suggestions/Update to the EA Guide

If you have any suggestions or updates to make to this EA Guide, please contact ITAG or the appropriate owner. They will be able to incorporate the changes as appropriate. This Guide and the artifacts presented in it can be updated with appropriate access rights.

Enterprise Architecture Block Diagram

The Enterprise Architecture Block Diagram shown below displays the various artifacts necessary to outline and detail MIT's Current and Future Enterprise Architecture. The diagram outlines the relationships and flow between these artifacts and is meant to be provide context for users of this Guide. Boxes with a dotted line and Grey Italic Lettering indicate artifacts not produced during the MIT EAP Reap process. Boxes with a solid line and Black Bold lettering display those artifacts which were.

As shown, the EA Guide Block Diagram is broken into 4 Sections: Current State, Future State, Strategy Implementations, Timeless/Evolutionary.

- **Current State:** The Current State section displays the artifacts produced as part of the MIT EAP Reap process (with the exception of one) which outline the Current State of MIT's EA. The flows outlined displays the relationships between the documents and the information which was gathered and subsequently used to derive other artifacts.
- **Future State:** The Future State section displays the artifacts produced as part of the MIT EAP Reap process which outline the Future State of MIT's EA similar to the Current State section.
- **Strategy Implementation:** This Strategy section displays the artifacts which outline MIT's proposed Future Enterprise Architecture. These artifacts also contain information detailing the method to achieve this proposal. As displayed, these artifacts were not produced as part of this initiative.
- **Timeless/Evolutionary:** The Timeless/Evolutionary section displays those artifacts which will support MIT's EA through its various stages.

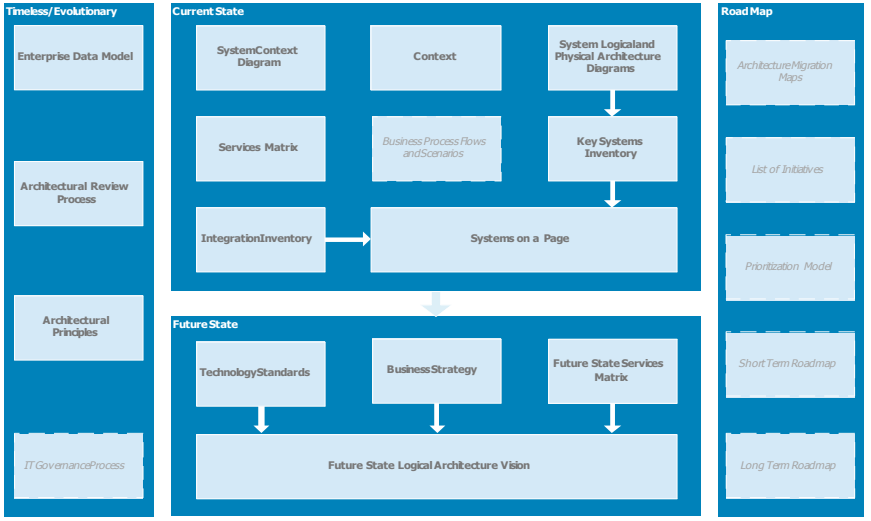

1. Executive Summary

Current State Summary

The project team used a variety of tools to document MIT's current state architecture. We conducted 22 interviews involving 25 people representing a cross section of the Institute and held four consensus building workshop sessions.

The team created a system context diagram, a systems on a page diagram describing how key enterprise systems relate to each other, 22 architecture diagrams describing key enterprise systems (12 logical architecture and 10 physical architecture diagrams), an enterprise data model, and a services matrix describing currently available enterprise services.

System context diagram, systems-on-a-page diagram and enterprise entity relationship diagram are shown below.

Key themes that emerged from the interviews conducted to document the current state are listed below:

- Current issues with latency of data updates can be reduced by moving to a more real-time integration model
- There is no single source of people information at MIT, leading to wide variety of problems
- There is no clear vision for how to manage information and security for people who belong to extended MIT community
- There is a significant amount of data shadowing across enterprise and departmental systems at MIT
- There is no clear policy around data ownership (custodians) at MIT
- An enterprise standard Software Development Lifecycle (SDLC) process is missing
- There is an opportunity for IS&T to clarify the process for engaging its services, as well as an opportunity to offer additional support services that DLCs are expecting
- While there is an enterprise solution for authentication and authorization services, these services are not uniformly adopted by enterprise systems

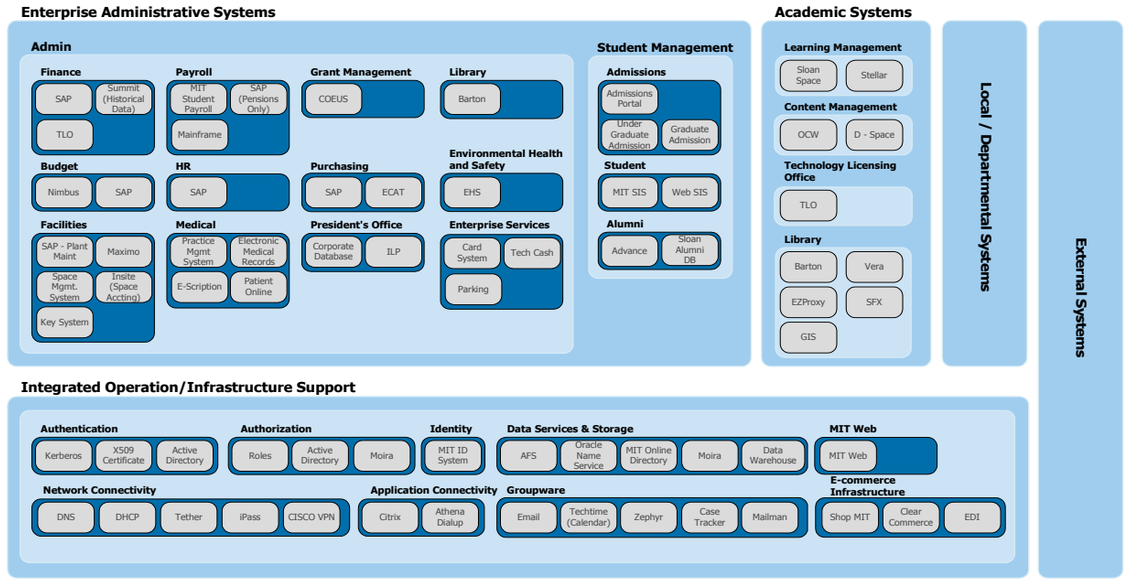

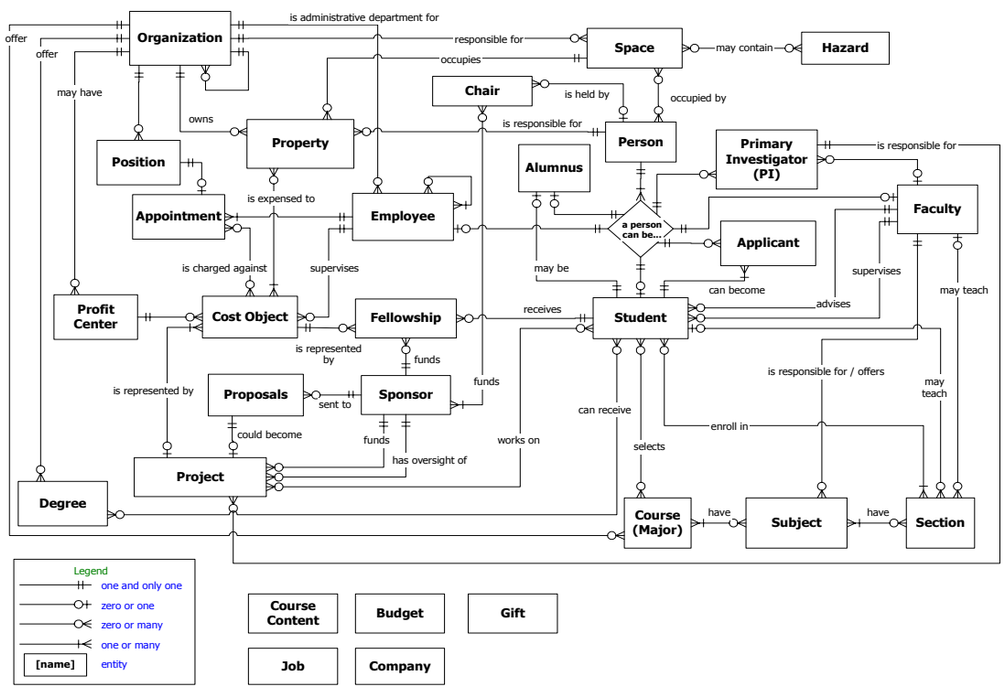

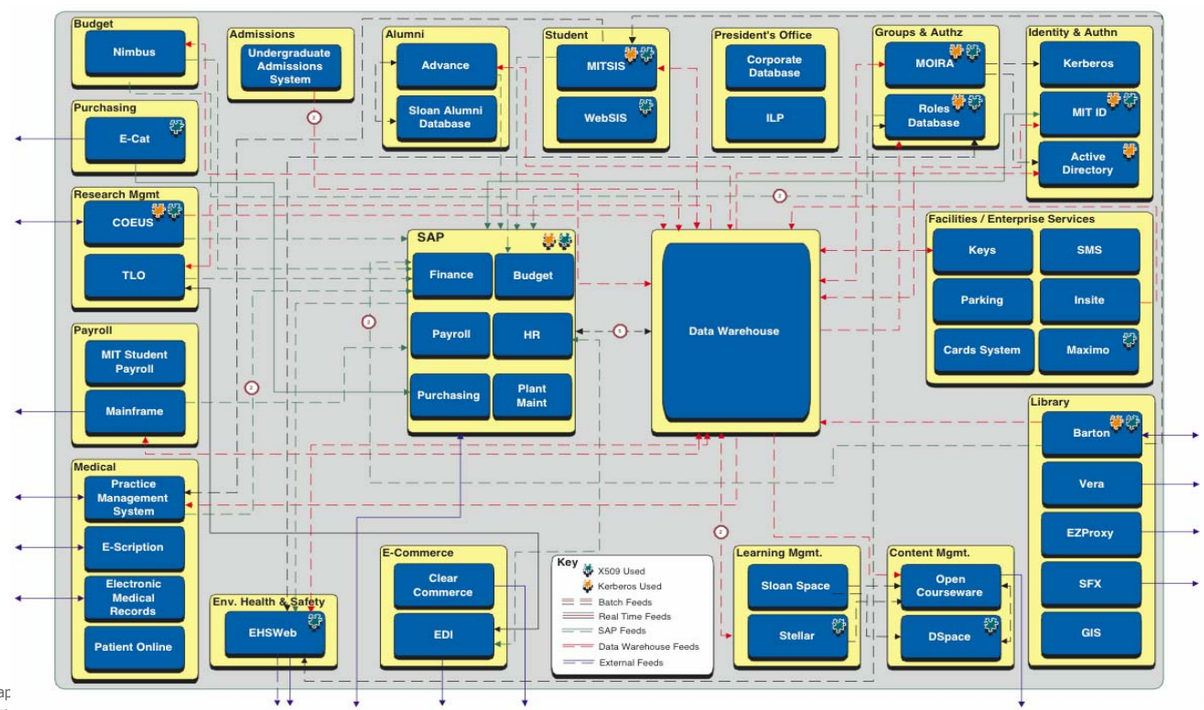

Future State Summary

The enterprise architecture future state vision should support MIT's future vision for the Institute while operating within the Institute's future state context. Furthermore, the architecture should be consistent with a set of principles defined during the future state workshops. The context and principles are summarized below:

Context:

- SAP will continue to be the primary ERP system; MIT may have other systems providing some ERP services.
- The MIT Data Warehouse will be the central repository for administrative data that is of interest to more than one DLC.
- Our user community will be based throughout the world, and will require 24x7 access to our systems; the definition of the MIT community will be amorphous, and will continue to evolve
- There will be increased integration between MIT and other universities; there will be increased need for collaboration between members of MIT community and external community (e.g. other universities, research labs, etc.)
- The MIT environment is heterogeneous
- The MIT network will evolve to support needs of the enterprise; we may have many research networks, we will have an IPv6 network and we will need differentiated services to better support user needs

Principles:

Security: applications should ensure data and access security

Ownership: clear and explicit ownership of enterprise data

Leverage assets: leverage existing services and capabilities

Accessibility: be aware of to needs of all users (location & disabilities)

Real-time: Minimize latency of data updates

Standards: promote consistency using standards

Logical Architecture:

In addition to the future state context and desired architecture principles, the project team also worked on identifying the services that should be part of any future state architecture vision for MIT and technology standards for commonly used components. Context, principles, technology standards and the services matrix were key input into developing the future state logical architecture diagram shown in the adjacent diagram.

The logical architecture diagram is the ideal state architecture which MIT should move towards. All technical decisions in the future should be accessed with this framework in mind. Each subsequent enterprise application should move the Institute closer to realizing this vision.

Moving Forward:

The nest steps for the institute are to refine the future state vision further to add detail, and then to develop a roadmap for implementing the vision.

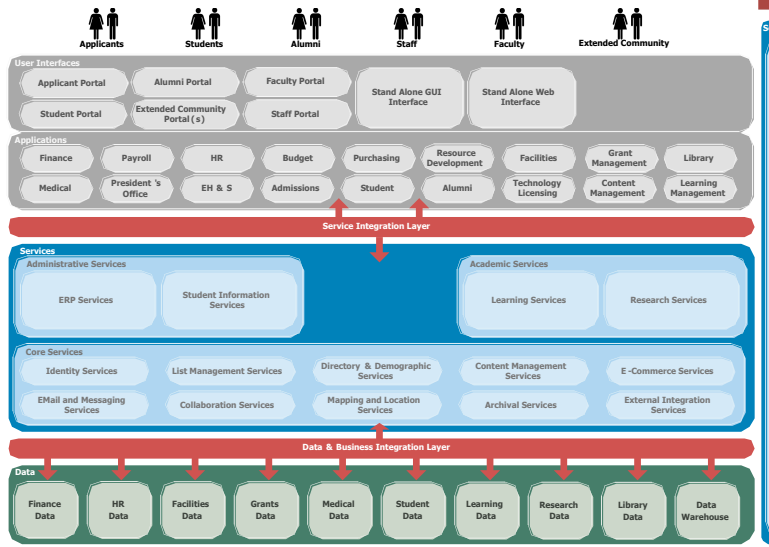

2. Context and Principles

Context and Principles

The following statements are intended to represent the business and technology landscape for the next three to five years. Thus, these items must be considered when discussing the future state vision.

Future State Context:

1. SAP will continue to be the primary ERP system; MIT may have other systems providing some ERP services
2. The MIT Data Warehouse will be the central repository for administrative data that is of interest to more than one DLC
3. Our user community will be based throughout the world, and will require 24x7 access to our systems; the definition of the MIT community will be amorphous, and will continue to evolve
4. There will be increased integration between MIT and other universities; there will be increased need for collaboration between members of MIT community and external community (e.g. other universities, research labs, etc.)
5. The MIT environment is heterogeneous
6. The MIT network will evolve to support needs of the enterprise; we may have many research networks, we will have an IPv6 network and we will need differentiated services to better support user needs

Goals

During the same discussions a number of goals were identified. These are listed below:

- Business rules and processes for accessing data will be well documented
- We will have a central repository (logical) for academic data
- We will eliminate "swivel chair" integration
- We will have a clear definition of what our community is but it may be complex with many parts

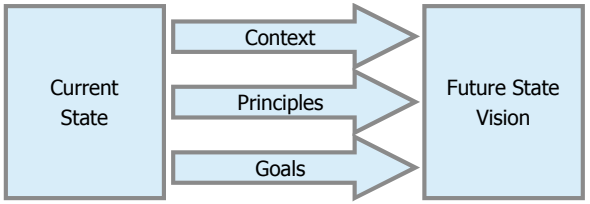

Principles

Principles are intended to be simple statements of concepts that can be easily remembered, and used to guide the development of enterprise applications to evolve and improve the enterprise architecture. The statements below were agreed upon by the group and are intended to be used by application architects and developers to understand how they can contribute towards realizing MIT's enterprise architecture vision.

Security: applications should ensure data and access security
- Sensitive data must be protected in storage and in transit
- People should have single identity to all enterprise applications (single sign-on)
- Usernames should be consistent across applications

Ownership: clear and explicit ownership of enterprise data
- All enterprise level data entities should have a single identified system of record
- Systems should fulfill their custodial obligations for data they are the system of record for

Leverage assets: leverage existing services and capabilities
- Leverage capabilities in our existing investments where appropriate (SAP, Data warehouse, roles, etc.)

Accessibility: be aware of needs of all users (location & disabilities)
- Enterprise applications should be accessible from anywhere
- Enterprise applications should support accessibility standards

Real-time: Minimize latency of data updates
- Minimize latency of data updates

Standards: promote consistency using standards
- All new enterprise applications should adhere to recommended technical standards
- Use of open source tools and specifications

3. Current State

The current state assessment is intended to provide a consolidated understanding of systems currently in use at MIT. In scope it covers Enterprise systems; an Enterprise system is one which is important to the operation of more than one Department/Lab/Center or is important to the operation of the institute in general.

The section includes:

- The Systems Context Diagram
- The Integration Inventory
- The Systems on a Page Diagram, showing interfaces and interactions between systems
- The Services Matrix
- The Enterprise Data Model
- The Enterprise Entity Ownership Matrix
- Current State Assessment Themes
- **A summary of themes from the assessment of the current state architecture**

Systems Context Diagram

The Systems Context Diagram shows the various enterprise class systems at MIT grouped into functional sets.

- **Enterprise Administrative Systems** are systems that are used to administer the operational aspects of MIT. This includes all aspects of Enterprise Resource Planning (ERP), Facilities Management, Health and Safety and other administrative functions. It also includes Student Management, which encompasses undergraduate and graduate admissions, registration, enrollment and alumni systems.

- **Academic Systems** are the sum total of those systems whose purpose is to advance MIT's academic capability. This includes systems for learning management (course administration, homework tracking etc.), content management (the management and publication of course materials and the results of research), the licensing and management of intellectual property resulting from research at MIT and the Library systems used to locate and research information.

- **Integrated Operation/Infrastructure Support** are all the services which enable systems to be deployed, managed and accessed at MIT.

- **Local and Departmental Systems** are systems that are used exclusively within one Department, Lab or Center (DLC). These may range from small, desktop based applications to large client/server or web applications. While these systems are not considered enterprise class they may be extremely mission-critical to the operation of a single department, and may represent large investments in IT.

- **External Systems** are systems that are not owned, leased or operated by MIT, but with which one or more systems at MIT interact. Examples of external systems are:
  - Meta Data providers for the Library Systems
  - IDX eCommerce Clearing House (for Medical Payment transactions)
  - Grants.gov for grant proposal and award information

For more detail on specific external interfaces please refer to the application architecture diagrams later in this section.

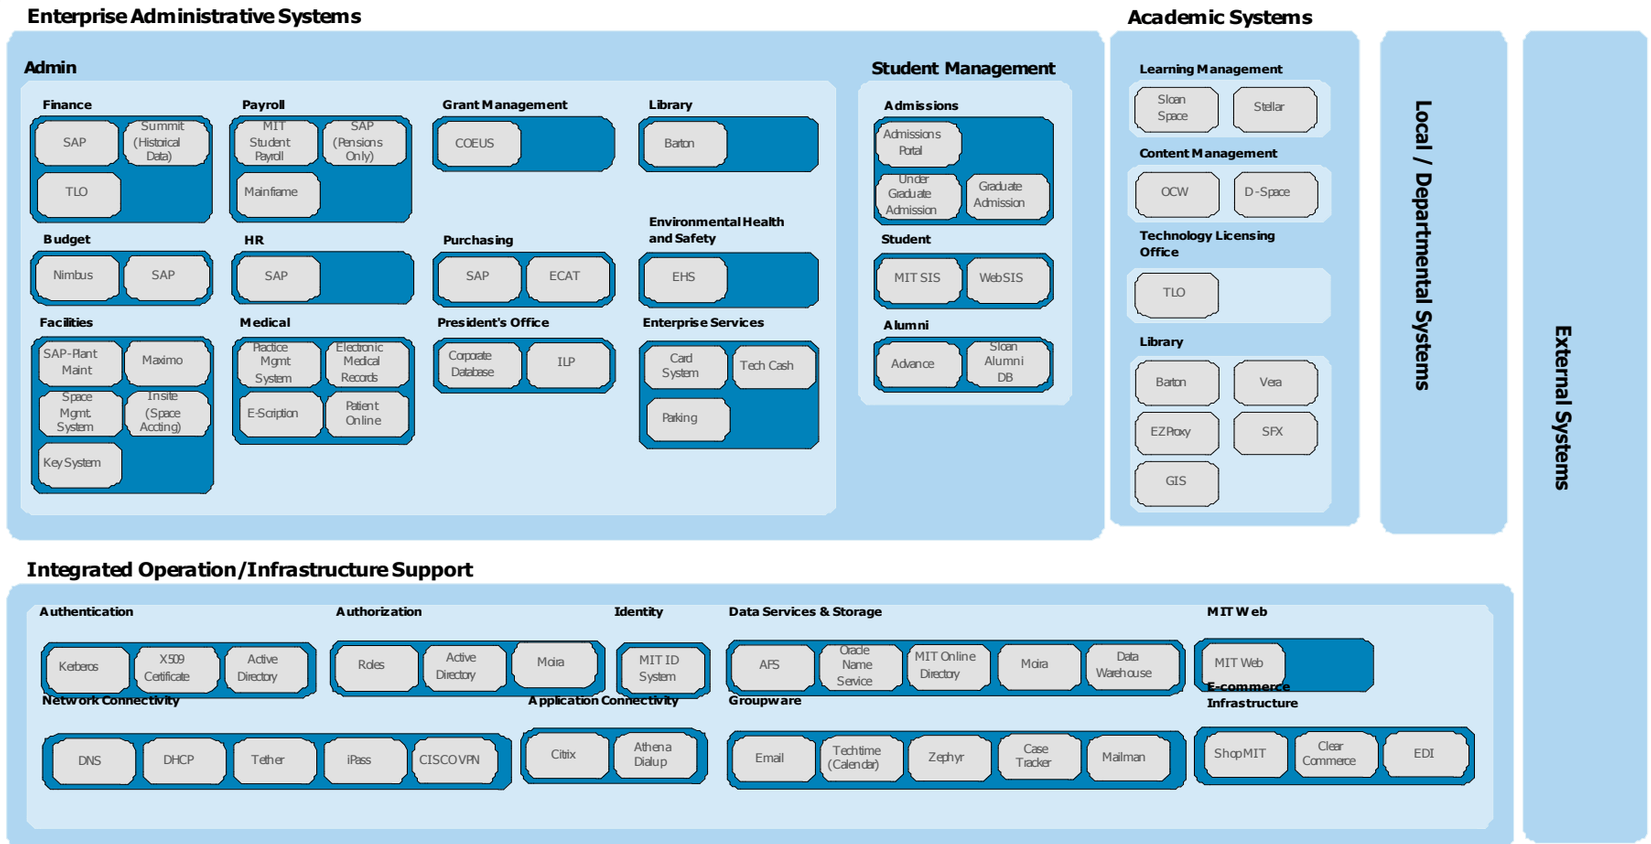

Instructions for Integration Inventory

Description	The Integration Inventory is the consolidated list of integrations that exist between enterprise systems. Given the nature of enterprise architecture this document will evolve and change as the architecture changes.
Version	0.1
Worksheet Definitions	
Integration Inventory	The Integration Inventory; see above
Column Definitions	
ID	A numeric ID that can be used to refer to the integration.
Source or Initiating System	In the case of a batch feed, this is the system from which data is sent. In the case of a real time integration, this is the system which is the client and initiates each transaction.
Target System	In the case of a batch feed, this is the system which receives and processes the data. In the case of a real-time integration, this is the system which acts upon the request of the initiating system
Data or Transaction Type	The type(s) of data being fed in a batch feed or the transaction type being requested by the initiating system.
Batch or Real-Time	Whether the integration is a batch feed, or a real-time integration between systems.
Frequency	In the case of batch feeds, how often is the feed executed.
Data Format	The class of data format that data or messages are sent in (e.g. XML, Flat Files etc.)
Owner	The person or group responsible for maintaining the integration.

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
001	Admissions	Data Warehouse	Admissions Applicant Load Admissions High School Admissions High School Statistics Admissions Lookup Data Admissions Master Data	Batch	Annually	Flat File	Data Warehouse Group
002	Admissions	Data Warehouse	Grad Admissions Applicant Graduate Admissions Application Graduate Admissions Application Graduate Admissions Degree Objective Grad Admissions Master Data Graduate Admissions Programs Grad Admissions School Attended Graduate Admissions Survey Graduate Admissions Test Score	Batch	Weekly	Flat File	Data Warehouse Group
003	Advance	Sloan Administrative DB	Alumni Profile Data	Batch	Hourly	XML	
004	Alumni	Data Warehouse	Alumni Biographic Data Alumni addresses Alumni degrees	Batch	Weekly	Flat File	Data Warehouse Group
005	Broad Institute	SAP	Purchase Orders, Goods Receipts	Batch	Multiple Times a Day	Flat File	
006	Campus Legacy Payroll	Lincoln SAP	Gross Pay	Batch	weekly	Flat File	Lincoln ITO/LFO
007	Campus Legacy Payroll	Lincoln SAP	Leave	Batch	weekly	Flat File	Lincoln ITO/LFO
008	COEUS	Data Warehouse	OSP Award Comments OSP Awards OSP Award Cost Sharing Data OSP Award Terms OSP Award Indirect Cost Data OSP Sponsors	Batch	Daily	Flat File	Data Warehouse Group
009	COEUS	EDI	Proposals				
010	COEUS	SAP		On Demand		Flat File	
011	Cyborg	Data Warehouse	HR data from Cyborg	Batch	Daily	Flat File	Data Warehouse Group
012	Data Warehouse	Active Directory	Contact Information				
013	Data Warehouse	Advance		Batch	Daily/Nightly		
014	Data Warehouse	Case Tracker	Directory Information	Batch	Daily		Jeff Schiller
015	Data Warehouse	COEUS	MIT ID, People Data, Appointment Data, HR Data	Batch	Nightly	Oracle SQL Connection	
016	Data Warehouse	Current Payroll System	Chart of Account	Batch	Daily	Flat File	
017	Data Warehouse	EHSWeb	List of all people Buildings and rooms List of all job titles	Batch	Daily/Nightly	Oracle SQL Connection	Jim Repa
018	Data Warehouse	LNS (Lab for Nuclear Science)		Batch	Daily/Nightly		
019	Data Warehouse	Medical: Practice Management System	People, Benefits, Eligibility	Batch	1 - 2 times a week - Nightly	Flat File	
020	Data Warehouse	MIT Directory	Directory Information: Name and Address	Batch	Daily/Nightly	Flat File	

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
021	Data Warehouse	MITID		Batch	Nightly		
022	Data Warehouse	MOIRA	Directory Information & MIT ID, Class of Person	Manual, a-periodic	A periodic	Flat File	
023	Data Warehouse	Nimbus	Master Data and Actuals	Batch	Daily/Nightly	Oracle SQL Connection	Loti Petrov, budget office
024	Data Warehouse	OCW	Employee Directory and CIP codes	Manual via Brio Query	Twice a Year		Cec
025	Data Warehouse	Practice Management	Payments made to Patients and Vendors	Batch	Nightly	Flat File	
026	Data Warehouse	Request Tracker	Directory Information	Batch	Daily		Jeff Schiller
027	Data Warehouse	RLE		Batch	Daily/Nightly		"Fletch" - Miron (First name)
028	Data Warehouse	Roles	All people at MIT who have kerb principles - Qualifiers - Org Units (Various types, mostly from the warehouse) - Financial Units - Cost Collectors	Batch	Daily/Nightly	Oracle SQL Connection	Jim Repa
029	Data Warehouse	SAP	MIT ID info, Buildings and Room Data	Batch	Daily/Nightly	Flat File	
030	Data Warehouse	SAP	Certificate Data	Batch	Daily/Nightly	Flat File	
031	Data Warehouse	SIS/MITSIS		Batch	Daily/Nightly	Oracle SQL Connection	Kent
032	Data Warehouse	Stellar	Courses & Enrollments	batch	daily	Oracle SQL Connection	Craig Counterman
033	Data Warehouse	Stellar	Kerberos users	batch	daily	Oracle SQL Connection	
034	Data Warehouse	TLO/Forrester	Financial	Batch	Daily/Nightly	???	TLO/Howard Bailey
035	ECAT	SAP	Purchase Orders	Real Time	On Demand	XML	Steve Landry
036	EDI	COEUS	Awards				Note: this will eventually go to Gov.
037	EHSWeb	Data Warehouse	EHS Room sets along with corresponding DLC, prim supervisors and School Area info EHS Access Control EHS Training for certification EHS training Access control	Batch	Daily	Flat File	Data Warehouse Group
038	EHSWeb	NetCaster (External)	People registered for training at MIT	Batch	Daily/Nightly	Flat File	Jim Repa
039	EHSWeb	NetCaster (External)	User's web based courses required	Real Time	On Demand	HTTP Hidden Form Variables	Jim Repa
040	Facilities: SAP Plant Maintenance	Online Directory	Contact Information	Real Time		Using GUI XT	Mike Sherman, Facilities
041	Insite	Data Warehouse	Space Data (Buildings and Rooms)	Batch	Daily	Flat File	Data Warehouse Group
042	Library: Barton (Aleph)	Data Warehouse	Library order arrival data Library Catalog -- incremental (current calendar year only) Library Circulation Data (Incr Load) Library Definitions from Aleph Library Transactions Library Invoices Library Item Detail FULL load Library Lookup Data Library Master Data Library Orders	Batch	Daily	Flat File	Data Warehouse Group
043	Library: Barton (Aleph)	MITSIS	Accounts receivable - students (Circulation bills for overdue and lost items)	Batch	Monthly	Flat File	Christine Moulen

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
044	Library: Barton (Aleph)	SAP	Accounts receivable - non-students (Circulation bills for overdue and lost items)	Batch	Monthly	Flat File	Christine Moulen
045	Library: Barton (Aleph)	SAP	Accounts payable (Acquisitions invoices for items purchased)	Batch	Daily (3-5 per week)	Flat File	Christine Moulen
046	Library: Document Services	E-Commerce Server (Omar)	Credit Card Processing				
047	Lincoln Lab SAP	EHSWeb	Completed Courses Registered Courses	Batch	Daily/Nightly	Flat File	Jim Repa
048	Lincoln Labs SAP	Data Warehouse	Training Records	Batch	Daily	Flat File	Data Warehouse Group
049	Lincoln SAP	SAP	HR	Batch	nightly	XML/ALE Idocs	Lincoln ITO/HRIS
050	Mainframe: Fleet Bank -> Magtape -> MITVMA	SAP	Bank Statements	batch	three tapes, once a month	Flat File	Administrative Computing
051	Mainframe: Payroll (on MITVMC?)	SAP	Journal vouchers	batch	monthly	Flat File	CAO (Payroll)
052	Mainframe: Payroll (on MITVMC?)	SAP	Payroll Disbursements	batch	once or twice a week	Flat File	CAO (Payroll)
053	Mainframe: Payroll -> MITVMA	SAP	Salaries & Salary Distributions	batch	monthly	Flat File	CAO (Payroll)
054	Medical: Practice Management System	SAP	Requests for Payment	Batch	Daily	Flat File	
055	Medical: Practice Management System	SAP	Non Students Accounts Received	Batch	Once a week at least, Can be more	Flat File	
056	Medical: Practice Management System	Students Accounting	Billing Information about Students	Batch	Once a week at least, Can be more	Flat File	
057	MIT Course Catalog	OCW	Course descriptions, master course numbers, faculty etc.	Manual via Spreadsheet	Once a Year	Excel Spreadsheet	
058	MIT ID	Data Warehouse	MIT ID	Batch	Daily	Flat File	Data Warehouse Group
059	MIT ID	SAP		On Demand		RFC - Encrypted	
060	MITSIS	Practice Management	Student Demographics and Enrollment	Batch	Once or twice a week	Flat File	
061	Moira	Active Directory	User and Group Info				
062	Moira	AFS	Athena Update Protocol		Several times a day		
063	Moira	Data Warehouse	Krb Mapping data from MOIRA	Batch	Daily	Flat File	Data Warehouse Group
064	Moira	DNS	Athena Update Protocol		Several times a day		
065	Moira	Hessiod	Athena Update Protocol		Several times a day		
066	Moira	Kerberos	Athena Update Protocol		Several times a day		
067	Moira	Mailman	Athena Update Protocol		Several times a day		
068	Moira	mit.edu (Mail Routing Table)	Athena Update Protocol		Several times a day		
069	Moira	Network Database	Athena Update Protocol		Several times a day		
070	Moira	Post Offices	Athena Update Protocol		Several times a day		
071	Moira	Print Servers	Athena Update Protocol		Several times a day		
072	Moira	Stellar	Groups	Batch			
073	NetCaster (External)	EHSWeb	Completed Courses	Batch	Daily/Nightly	Flat File	Jim Repa
074	Nimbus	Data Warehouse	Budget Groupings	Batch	Weekly	Flat File	Data Warehouse Group
075	Nimbus	SAP	Budget	Batch	Nightly	Flat File	OBFP/Lody Petriv
076	Nimbus (JV Feed)	SAP	Financial	Batch	Nightly	Flat File	OBFP/Lody Petriv
077	OCW	Dspace	Archived Course Content	Batch	On Demand		Will be in future, it does not exist right now
078	PPL Keys	Data Warehouse	Keys	Batch	Weekly	Flat File	Data Warehouse Group

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
079	Roles	Data Warehouse	XSIS Roles Roles SAP Authorization Information HR Roles	Batch	Daily	Flat File	Data Warehouse Group
080	Roles	EHSWeb	Authorization Data Master Department Codes and Names	Batch	Daily/Nightly	Oracle SQL Connection	Jim Repa
081	Roles Database	SAP	Authorizations	Batch	Daily	Flat File	
082	Roles Database	SAP	Hierarchy information	Batch	Daily	Flat File	
083	SAP	Broad Institute	Purchases Received	Batch	Multiple Times a Day	Flat File	
084	SAP	EDI	Invoices, Benefits, Procurement Card?				Check if procurement cards come into SAP through EDI.
085	SAP	EHSWeb	Completed Courses Registered Courses	Batch	Daily/Nightly	Flat File	Jim Repa
086	SAP	Lincoln Labs					Future
087	SAP	MIT ID		On Demand		RFC - Encrypted	
088	SAP Benefits	Data Warehouse	Benefits Enrollment Detail Benefits Lookup Benefits Master data	Batch	Daily	Flat File	Data Warehouse Group
089	SAP Finance	Data Warehouse	Balances, Incremental Institute Budget SAP Cost Element Hierarchy, Groups SAP Change Log Items Commitments History Financial Actuals Credit card tx from Financial Detail Fin Detail Clearing Data Financial Commitment Financial Overhead Costs Fund Descriptions from CAO's office Balance Sheet Accounts Balances Financial Processing Status LDS Person-to-Cost Collector assignment data One Time Vendor - Financial Detail Overhead calculation rules Payment information update (Bank tape file) SAP Payment (Check) Information SAP Profit Center Hierarchy SAP Profit Center Group s Pension Payroll PENSION Personal Data, Statuses, Addresses, Actions	Batch	Daily	Flat File	Data Warehouse Group

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
090	SAP HR	Data Warehouse	HR Academic Chair HR Appointments HR Appointment Letters HR lookup table. HR Master Extract Data HR Objects (jobs, positions) SAP HR Org Hierarchy SAP HR Org Hierarchy HR Person miscellaneous (new HR) HR Personal Data, Statuses, Addresses, Actions HR Training and Events Personal Data HR Object relationships	Batch	Daily	Flat File	Data Warehouse Group
091	SAP Other	Data Warehouse	SAP Project Wbs Hierarchy SAP Access Control SAP Lookup tables. Used in all SAP conversions SAP Master Tables	Batch	Daily	Flat File	Data Warehouse Group
092	SAP Purchasing	Data Warehouse	Purchasing Master Tables Purchasing Data - Raw Data Conversion Purchasing Requisitions Sales Orders Sales Contract Invoices Sales Contracts	Batch	Daily	Flat File	Data Warehouse Group
093	SIS	Data Warehouse	Student 5th week enrollment table Student Enrollment Y Report Financial Aid Time Dimensions	Batch	Term	Flat File	Data Warehouse Group
094	SIS	Data Warehouse	External Test Scores Financial Aid Applicant, Award, and Need Fin Aid Federal Work Study Financial Aid MIT Grant Packages and Grant Management Financial Aid Master Data Fin Aid Student Payroll Distribution Detail Financial Aid Requirement Trackings Graduate Award Term Detail Pre-Registration Data SIS Master Data Student Medical Insurance Student Biographic Student Subject Enrollment Student Term Enrollment Student Thesis Milestones	Batch	Daily	Flat File	Data Warehouse Group
095	SIS	Data Warehouse	Student Degree Information Student Tuition Assessment Student Degree Information	Batch	Weekly	Flat File	Data Warehouse Group
096	SIS	Data Warehouse	Sloan School Subject Type	Batch	On Demand	Flat File	Data Warehouse Group
097	Sloan Administrative DB	Advance	Alumni Profile Data	Batch	Hourly	XML SCP	

MIT Integration Inventory

ID	Source or Initiating System	Target System	Data or Transaction Type	Batch or Real Time	Frequency (if Batch)	Data Format	Owner (Group or Person)
098	Sloan Space	Dspace		On Demand		XML Web services	Will be in Dec. Does not exist right now. Add a feed from Dspace to Sloan as well.
099	Sloan Space	OCW	Course Content	Manual Download	Twice a Year		
100	Stellar	OCW	Course Content	Batch and Manual Download	Twice a Year		
101	Student Payroll	Data Warehouse	Student Payroll	Batch	Monthly	Flat File	Data Warehouse Group
102	TLO/Forrester	SAP	Financial	Batch	Monthly/Periodic	Flat File	TLO/Howard Bailey
103	Web Hits	Data Warehouse	Web Hit Detail	Batch	Weekly	Flat File	Data Warehouse Group

Systems on a Page

The Systems on a Page diagram shows the interactions between all of the non-infrastructure enterprise systems at MIT. It is a visual representation of the information contained within the Integration Inventory. Several key facts can be discerned from the diagram:

- The Data Warehouse is a central clearing house for a large number of feeds
- While Kerberos and X509 certificates are widely used they are not ubiquitous
- Many systems have external integrations, which are managed independently
- The Majority of the remaining integrations are with modules in SAP

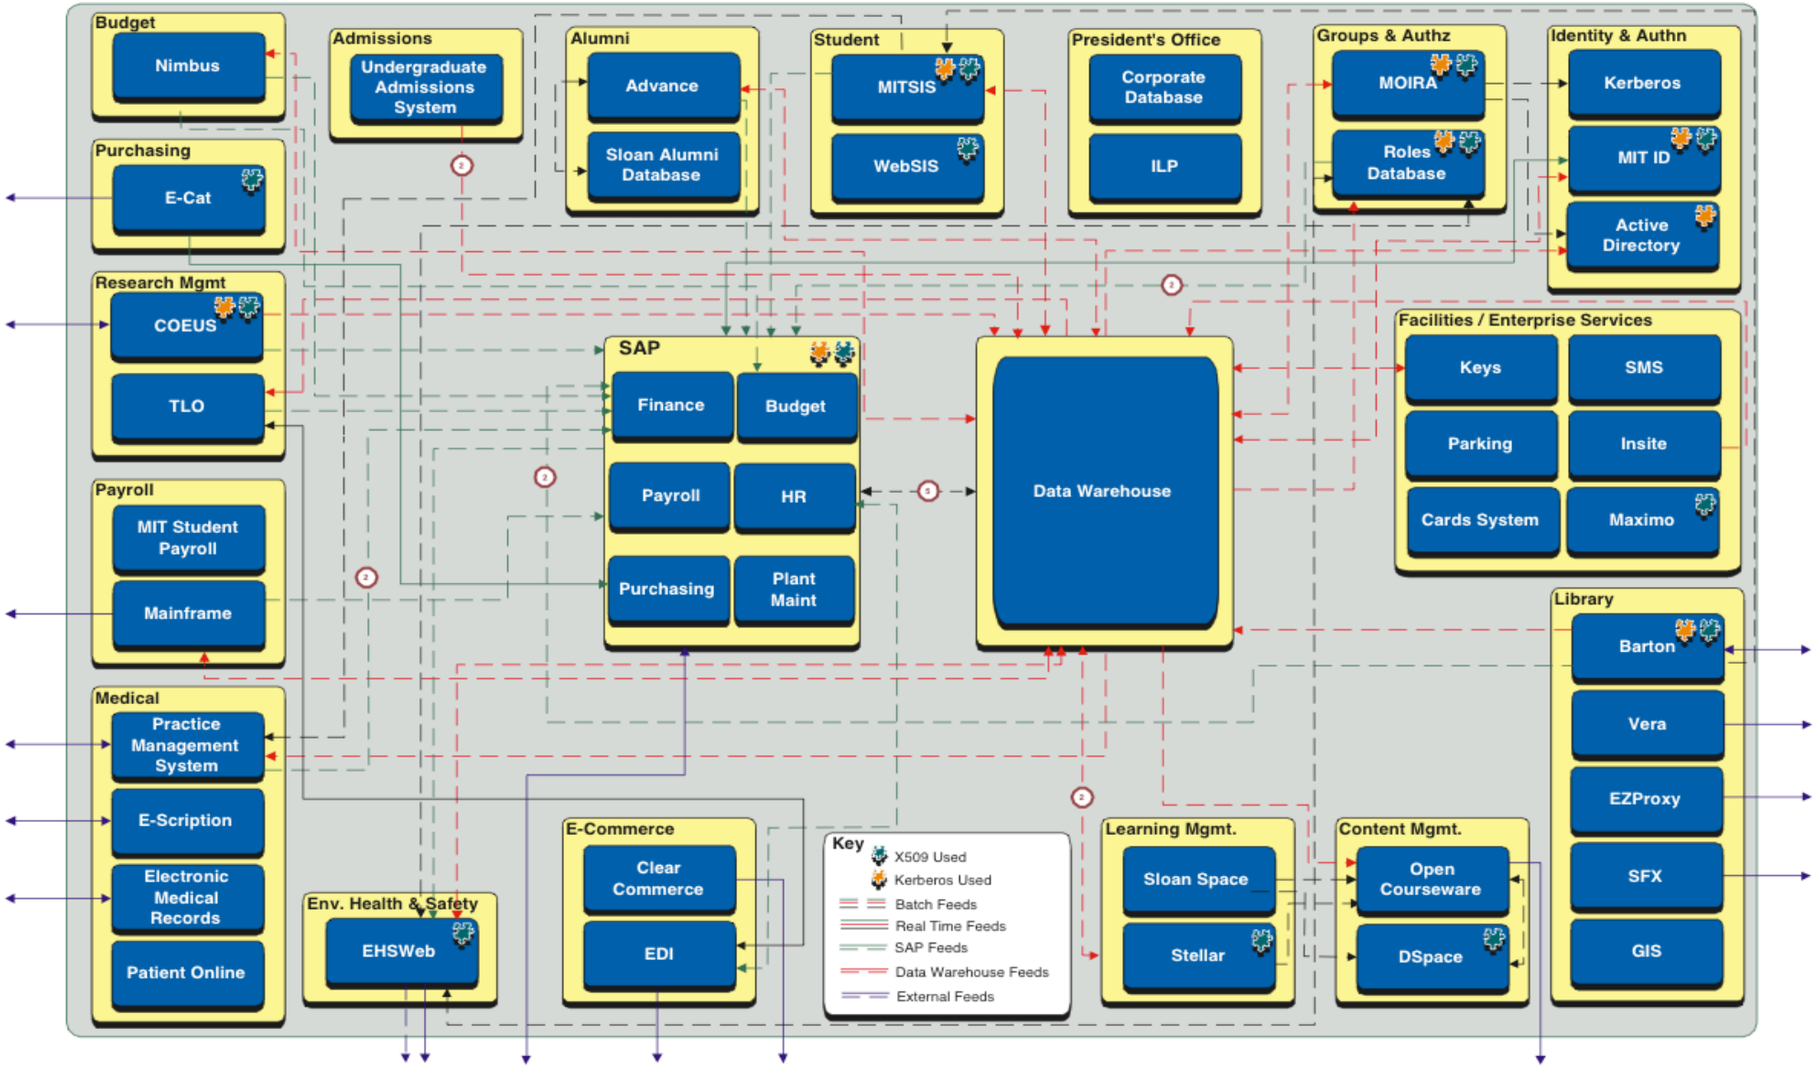

Instructions for Services Matrix

Description	The Services Matrix documents the services that are currently offered within the MIT Enterprise Architecture and MIT IT community.
Version	0.1
Types of Service	
Remote Service	A service is a piece of software functionality that is executing somewhere separate from your program that you can invoke in a defined manner.
Embeddable Service	SDKs & Libraries are reusable components that can be embedded into an application, but there is no central infrastructure for them.
Business/Operational Service	A Business/Operation Service involves interaction with a person or staff member.
End user Service	Services to end users and not applications are those which are used for functionality but have no programmatic access.
Worksheet Definitions	
Systematic Services	The catalog of services that are for use by applications, and are relevant to a services oriented enterprise architecture
Non-Systematic Services	The catalog of services that are for use by people and impact the enterprise architecture only peripherally
Column Definitions	
Service	The short name or description of the service.
Definition	A longer definition of what the service is.
Currently Used By	A representative sample of systems which currently make use of the service.
Current Service Type	Whether the service is currently a Remote Service, an Embeddable Service, a Business/Operational Service or an End User Service
Current Implementation	A description of the current implementation of the service, usually including the sytems which implements it.
Future Service Type	Whether, in future, the service should be a Remote Service, an Embeddable Service, a Business/Operational Service or an End User Service
Future Form	A description of how the service should be implemented in the future.
Areas for Investment	Areas in a current service, or anticipated future service that require investment in order to realize.
Contact	The name of the person who should be contacted for more information about the service, for example if you wished to used it from a new application.
Notes	Any further notes about the service or follow up items.

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
Infrastructure									
Security									
Authentication									
Authenticate a User	Allow an application to authenticate the user (i.e. assert that they own the identity supplied).	Barton, SAP, COEUS, MITSIS etc.	Remote Service	MIT Kerberos	Remote Service	MIT Kerberos	• Version 4 needs to be eliminated. • Keep up with ITEF Protocol.	Jeff Schiller	
Authenticate a User	Allow an application to authenticate the user (i.e. assert that they own the identity supplied).	COEUS, Stellar, Dspace, Barton, Ecat etc.	Remote Service	X509 Certificate	Remote Service	X509 Certificate		Jeff Schiller	
Password Reset	Allow a user to reset their password when locked out and attempting to access an application.		Business / Operational Service	User must present ID in person at Accounts Department in N42.	Remote Service	Unknown	Get statistics on man hours/money needed.		
Authorization									
Roles (direct access)	The Roles Database provides a consistent way to store and maintain access rules for other applications. Applications with an interface to the Roles Database interpret the access rules from the Roles Database and enforce them.	SAP, Data warehouse	Remote Service	Roles Database	Remote Service	Roles Database	Need a service interface that is higher level than the current SQL based access.	Jim Repa	*To Do: Break this line item up in to the actual services provided by Roles*
Identity									
Create MIT IDs	The MIT ID is a 9 digit number used to uniquely identify any member of the MIT community. An MIT ID can be created through the MIT ID Database web client.		MIT ID - Remote service			MIT ID - Remote service	Need to link or consolidate IDs, when someone is a student + alum + employee		
Retrieve MIT IDs	The MIT ID can also be retrieved through the MIT ID Database web client by supplying a person's first and last name.	Medical	MIT ID - Remote Service			MIT ID - Remote service			
Network									
DHCP	The DHCP (Dynamic Host Configuration Protocol) Service lets a user connect his/her computer to MITnet from a variety of sites on campus without reconfiguring his/her computer's network settings each time the computer is moved to a new location.		Remote service						
DNS	The Internet Domain Name Service (DNS) can translate host names into equivalent IP addresses and vice versa, as needed by various Internet programs.		Remote service						
Host ID Management			Remote service						
Messaging & Communication									
SMTP(S) Servers - Email Transmission	Outgoing mail servers are referred to as SMTP servers. The outgoing mail server at MIT is named outgoing.mit.edu.		Remote service						

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
IMAP/POP3 Servers	IMAP (Internet Message Access Protocol) is a standard set of rules for storing, accessing and working with e-mail on a post office server. One of the main advantages of IMAP is that it makes your e-mail easily accessible from multiple locations and computers. POP (Post Office Protocol) is a set of rules for storing and accessing your e-mail on a central server. When you access messages, they are downloaded to your local computer (or Athena home directory) and deleted from the server.		Remote service						
List Management (Mailman, Moira)	Mailman, Moira and Listserv are used to manage mailing lists at MIT.		Remote service						
MIT EDI Gateway	EDI is the electronic transfer of information between two trading partners' systems using a set of transactions that have been adopted as a national or international standard for the particular business function.		Remote service						
Data									
List Management									
Moira (direct access)	Moira is Project Athena's Service Management System. It controls the configuration of resources, including user accounts, remote file systems, printers, mailing lists, access control groups, and many other things.		Remote Service	Moira	Remote Service	Moira	List management needs to be integrated with Roles		*To Do: Break this line item up in to the actual services provided by Roles*
Content Management									
Directory									
Repository Services									
AFS - Remote File Service	AFS, the Andrew File system, is currently used by Athena as the file system for all user home directories and most of the other lockers. AFS is a distributed file system.		Remote service						
Information feeds from the Data Warehouse			Business/Operation service SQL - Remote service				XML Web Service for data retrieved		
Development Tools									
Change management									
Operations									
System Management									
Web Counter	The web counter service allows data gathering of the number of hits made to any web application.		Remote service				Question? Can this be used for sites outside Athena? Where can it be used?		
Admin									
eCommerce									

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
Clear Commerce Credit Card Processing	Clear Commerce is an enterprise software that sends transaction information to MIT's bank for verifying and processing payments on customers' credit cards.		Remote service						
Shopping Basket			Embeddable service						
Academic									
Learning management									
Educational Application									
Other									

Services Matrix - Non-Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
Infrastructure								
Security								
Authentication								
Acquiring a Server Certificate	Allow a system to use a certificate to identify and authenticate itself to a user.		Business / Operational Service	X509 Certificate	Business / Operational Service	X509 Certificate		Jeff Schiller
Password Reset	Allow a user to reset their password when locked out and attempting to access an application.		Business / Operational Service	User must present ID in person at Accounts Department in N42.	Remote Service	Remote service	Get statistics on man hours/money needed.	
End user certificate management		Stellar	End user service					
Authorization								
Encryption								
Identity								
Network								
Messaging & Communication								
Web Casts	Events at MIT can be cast via the Web by using a fee-based service offered by AMPS.		Business/Operation service					
Data								
Reporting Service								
Reporting through the Data Warehouse	Reporting is done in the Data Warehouse through Brio or web reports. In addition, some departments have built their own custom applications to access the Data Warehouse for reporting purposes.		End user service					
Content Management								
Research content archival			Business/Operation service					
Directory								
MIT Online Directory	The MIT Online Directory allows a user to search for and view information about people in the MIT Community.		End user service					
Repository Services								
Information feeds to the Data Warehouse	The Data Warehouse is a storage for any data in the Institute which needs to be accessed by multiple systems. The Data Warehouse is updated daily by systems of records. Other systems can then extract this information from the Data Warehouse as appropriate.		Business/Operation service					
Information feeds from the Data Warehouse			Business/Operation service SQL - Remote service				XML Web Service for data retrieved	

Services Matrix - Non-Systematic Services

	Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
	GIS	GIS (Geographic Information Systems) are computer tools for managing data about where features are (geographic coordinate data) and what they are like (attribute data), and for providing the ability to query, manipulate, and analyze those data.		End user service			Web Service/Remote Service		
Development Tools									
	Change management								
Operations									
	System Management								
	Co-Location Services	Co-Location Services allow MIT applications and servers to be located in a separate location to allow for backup and recovery in case of any failure.		Business/Operation service					
	Server Monitoring	Server Monitoring is a service provided by IS&T to monitor servers for various applications in W91.		Business/Operation service					
	Consulting Services	IS&T offers various consulting services to the MIT Community.		Business/Operation service					
	Sub Domain Management			Business/Operation service					
	Backup and Restore			End user service					
	Issue Tracker			End user service					
Admin									
	eCommerce								
Academic									
	Learning management								
	Educational Application								
	Matlab	Matlab is a technical computing environment for high-performance numeric computation and visualization, produced by The MathWorks Inc. It includes a number of subject specific toolboxes as well as a dynamic system simulation package, Simulink.		End user service					
	Calendaring	A Calendaring service allows a user his/her schedule, or agenda in Calendar parlance, and also coordinate easily with the schedules of other users of the same Calendaring service.		End user service			Remote service		
	Events.mit.edu	This web site displays the events at MIT for the current day. It also allows the user to view upcoming events in various categories.		End user service					

Services Matrix - Non-Systematic Services

	Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
	Campus Map	The campus map pinpoints where you are, and where you're going. The campus map uses geographic information systems (GIS) data from the official maps maintained by the Department of Facilities, resulting in a more accurate mapping system. Using XML, the map is also integrated with the lists of departments, labs, and centers on the MIT top-level pages		End user service					
Other									
	Video Production	Video production is a fee-based service offered by AMPS to the MIT Community.		Business/Operation service					
	Survey Service			Business/Operation service					

Enterprise Conceptual Data Model

The Enterprise Conceptual Data Model is an Entity-Relationship Diagram (ERD) that illustrates high-level data entities within the MIT information systems domain and their relationships to one another.

The ERD is a conceptual data model that captures the overall structure of data independent of any database management system or other implementation considerations. This view of the ERD captures the entities at their highest level.

The ERD as shown at the right is meant to be a communication vehicle to a developer but not useful for any other purpose. The ERD must be expanded to include the objects comprised in each entity to be of value to developers.

Definitions

An **entity** is a class of persons, places, objects, events, or concepts about which we need to capture and store data. Boxes on the ERD represent entities (e.g. Student, Sponsor).

A **relationship** in the ERD associates instances of two entities through a connecting line. A relationship indicates that there is a natural link between entity types. A relationship has two ends, each connected to an entity. A relationship end whose properties are being discussed is generally referred to as the target end whereas the opposite end is referred to as the source end.

The **cardinality** of each relationship is indicated on the diagram using symbols at the connection point of the relationship lines to the entity boxes (see Legend). The cardinality of a relationship on the diagram should be interpreted as "one instance of the source entity may have [*multiplicity*] of the target entity". For example the relationship between Subject and Section can be expressed in two ways:

- A Subject may have zero or more sections
- A Section must be associated to one or more subjects

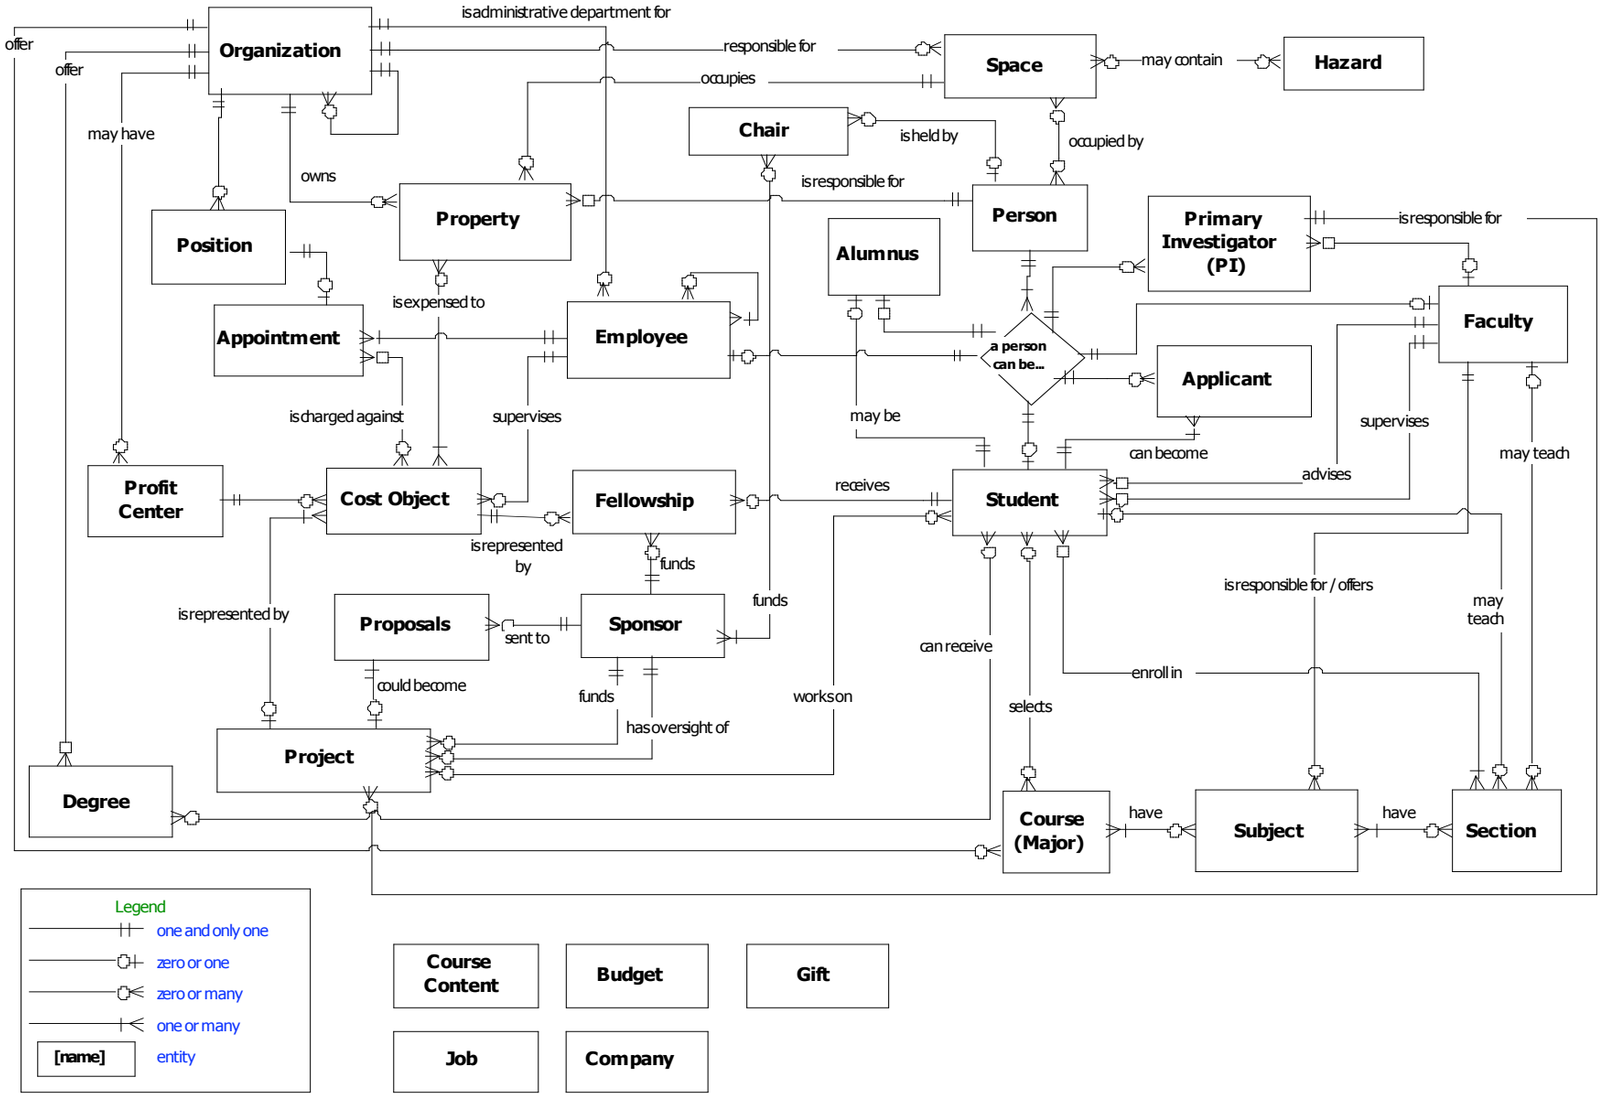

Instructions for MIT Entity Ownership Matrix

Description	The MIT Entity Ownership Matrix contains information about all of the entity types shown on the Enterprise Conceptual Data Model.
Version	0.1
Worksheet Definitions	
Contact Information	Entity Ownership Matrix; see above
Column Definitions	
Entity	Details the name of the Entity.
Class	Details the class for the entity. In some cases entities are broken down further in to classes, in order to show that different systems have primary responsibility for different segments of an entity.
Definition	Details the purpose of the entity.
Primary System	Details the the primary system responsible for maintaining the data related to that entity

MIT Entity Ownership Matrix

Entity	*Class*	*Definition*	*Primary System*
Administrative Department		In HR terms this is the single DLC responsible for performing the administrative functions such as parking and OPA. For faculty and non-faculty ranks this is the DLC of the appointment with the highest percent effort; if individual has two appointments with 50/50 effort, the DLC will be programmatically selected. If this logic does not produce the correct Administrative Department, the DLC administrator may request a DLC override.	SAP HR
Appointment	All	The intersection of a Person, Position and Organization. Currently not captured in the SAP system of record. There are 4 types of appointments: primary, dual, joint, and additional). See attached document.	Data Warehouse
Appointment Department/s		In HR terms this is the DLC that appoints an individual. There may be one or more appointment departments depending on whether an individual holds one or more appointments.	SAP HR
Budget	All		Nimbus (Moving to SAP)
Chair	All	A chair provides partial or full salary support for a faculty member. It may provide discretionary funds for the faculty member holding the chair. These funds may come from the Institute, a school, or an external donor. Chairs may be rotting term chairs (developmental for junior faculty) or non term chairs (as long as the faculty member is at MIT).	SAP HR
Company	All	In financial terms it shows which "company" a cost collector is attached. MIT uses separate "companies" in SAP to differentiate financial transactions for different legal entities with differing business rules or reporting requirements. There are currently three company codes, "CUR" for main campus, "TECR" for MIT's alumni magazine, Technology Review, and "LCP1" for Lincoln Laboratory.	SAP
Cost Object	All	There are three types of cost objects: Cost Center, Internal Order, and WBS (Work Breakdown Structure) Element. A Cost Center is a general or operating account. Cost Centers are budgeted on the fiscal year. Internal Orders are non-sponsored Fund Account (e.g., funding from the MIT Provost or gifts). They are not tied to the fiscal year and may or may not have budgets and/or receive interest income. A WBS Element is a sponsored account (carrying a 4-digit sponsor code, e.g., a corporate fellowship program, research grant, or contract). These cost objects are used to track expenses for a particular activity.	SAP
Course (Major)	All		SIS
Course Content	Offline Delivery		
Degree	All		SIS
Directory Department Override		In HR terms this is the DLC that will appear in the directory for the person. The directory department default is the DLC where the person holds a salaried appointment. If the person holds multiple salaried appointments, the directory department default is the administrative department. If the default DLC is not correct, an override DLC may be entered.	SAP HR
Fellowship	All		SIS
Gift	All		Advance
Hazard	All		EHSWeb

MIT Entity Ownership Matrix

Entity	*Class*	*Definition*	*Primary System*
Job	All	A job is the generic description or classification of a position. Many specific positions can link to a job. For example: There is one job code for an administrative officer (AO) but many specific positions (i.e., AO for Biology or AO for Ocean Engineering).	SAP HR
Organization	All	An Organizational Unit is a School, Department, Lab, Center or Division.	SAP HR
Person		An individual that has some affiliation with MIT.	
Position	All	A position is a specific (individual) employee assignment. A position has descriptive data attached to it, such as titles (official job title and position title), a specific description of responsibilities and skills (posting description), and a job code. Positions can be filled by a person; unoccupied (vacant); unoccupied and being recruited for; or cancelled. In SAP, a unique position is created for every "slot", occupied or not, within the Organizational Unit. Positions exist independent of the employee. Any position-related data are attached to the position, and individual employees who move into that position inherit those data. When they leave the position, they leave behind both the position and position-related data.	SAP HR
Primary Investigator	All	Primary Supervisor of a sponsored research project.	COEUS
Profit Center	All	A financial organizational unit in which cost objects are grouped.	SAP
Project	All	In financial terms Research Projects are WBS Elements. A WBS Element is a sponsored account (carrying a 4-digit sponsor code, e.g., a corporate fellowship program, research grant, or contract).	SAP
Property	All	Property is either physical, intellectual. It is owned, partially owned or leased by the Institute.	TLO, ILP, Property
Proposal	All	In terms of a research project it is a document submitted to a potential sponsor outlining proposed research to be done, and the personnel, materials, methods, space and cost for doing it.	COEUS
Section	All		
Space	All		
Sponsor	All	A state, local, national, or foreign government, a non profit foundation, or a for profit organization.	COEUS
	Online Delivery		
	Alumnus	MIT Graduate	Advance
	Applicant		SIS
	Employee	A person that has a paid or unpaid appointment.	SAP HR
	Faculty	There are 2 types of faculty members at MIT. They are tenured track and non tenured track. Tenure Track faculty have or are eligible for tenure and have academic responsibilities in academic departments. Examples of non tenure track faculty include professors or the practice and coaches and they are not eligible for tenure.	SAP HR
	Student	An individual that has registered or is eligible to register for classes.	SIS

Assessment Themes

Through the interviews and workshops conducted as part of the Enterprise Architecture Guide initiative, several themes emerged. They are documented here.

Integration

MIT has made significant progress in the last ten years in evolving from an architecture in which most integrations were accomplished as point-to-point integrations, to a model where the majority of integrations are performed in a similar manner through the Data Warehouse. Furthermore, the introduction of SAP as the ERP system for MIT, and the expansion of SAP's role has significantly reduced the number of point-to-point integrations as the functionality of more systems is encompassed in the SAP implementation.

The current model for integration is one in which nearly all integrations are batch feeds with a periodicity of twenty-four hours or greater. This introduced significant latency in to the architecture where in some cases a real-time integration would be more appropriate. Where as the Data Warehouse provides a de facto standard for performing batch integrations, there is as yet no standard or preferred way to perform real-time integrations.

People Information

There is no single source of information on people at MIT. This occurs for a number of reasons, but causes a wide variety of problems, not least of which is that people may end up with multiple identities (MIT IDs), and have duplicate information and fragmented information in systems. The main causes of this problem stated are:

- Different systems are interested in different communities. For example, HR may manage employee data, but the Medical Center may require information on spouses, dependents etc. that is not collected by HR or anyone else.
- There is no clear way of managing the movement of people between the categories of student, employee and alumni. Further, it is possible for a single person to be all three of these at separate times, or at the same time.
- There are no standard definitions of data types. What constitutes a person in one system may be different in another. Instead of specifying a common superset of attributes from which all systems draw a definition, systems simply define the entity according to their requirements.

Security Services

MIT has a world class and leading set of security services. MIT developed Kerberos and was one of the earliest adopters of X509 certificates for widespread client side authentication. Similarly the deployment of Moira and more recently the Roles Database shows significant foresight and effort around the aggregation and maintainability of authorization information.

Despite this fact, many systems at MIT still have their own separately maintained set of usernames and passwords. Several different reasons were stated:

- Off the shelf packages that do not readily support Kerberos or X509 certificates often cannot be customized to integrate with the MIT authentication systems
- There are no clear guides to integrating Kerberos or X509 with your application, and no documented process for requesting help
- X509 certificates are still perceived as problematic for the user
- It is unclear what the institute policy is on issuing Kerberos principles and X.509 certificates to member of the extended community, and thus how to authenticate members of the extended community is also unclear
- The MIT root certificate is not signed by one of the more well known root certificates, and this causes problems on some platforms

Extended Community

There is no clear vision for how to manage information and security for people who belong the extended community at MIT. There is also no clear definition of who forms part of the core MIT community and who is considered part of the extended community.

Data Shadowing and Ownership

Consistent with the integration model outlined above, there is a significant amount of data "shadowing" occurring between systems at MIT. The reasons for this are not simple or singular. In general, applications at MIT do not provide other applications with real-time access to their data in any way. Therefore if one system requires frequent access to information owned and maintained by another system, the only real option is to mirror data from the source system.

Software Development Lifecycle Process

There is no standard software development lifecycle process at MIT, and it will be challenging to create one in the future given the disparate nature of the groups that develop and maintain systems. Neither is there such a standard process for the development of the enterprise class systems at MIT. Due to the variety of SDLPs used, and the varying level of formality of them, there does not appear to be any agreement on a core set of standard milestones or activities that always form a part of software development. This makes it difficult to implement any standard process across projects because there is no way of determining when in a project something must occur. For example, it is difficult to implement an Architectural Review process because it is not possible to clearly state when in a technology project it should occur.

IS &T Support

A number of themes arose around the support that IS&T offers the rest of the IT community at MIT. The primary outcome was that most users of IS&T services were satisfied with the level of service that they received from IS&T, but that there were several areas in to which IS&T could expand their services to provide other useful offerings to the community. These included:

- Support for servers running one or more variants of the Windows operating system. Currently IS&T will support co-located servers running several variants of Unix and Linux, but do not offer support for servers running windows. This is primarily intended to discourage the use of Windows as a platform for enterprise solutions, because the combination of MIT's open network and Windows' security problems are thought to be too high risk. However, there are a number of cases where the optimal solution has included one or more components running under Windows, and the IS&T policy has forced DLCs to either host and manage the server at their own facility or procure external hosting for the server. There appear to be sufficient instances of this that IS&T may be able to provide support for windows servers more cost effectively than individual DLCs.
- A standardized way to engage and communicate with IS&T. There are no clearly documented and enforced policies for engaging IS&T or for requesting support. As a result a relationship-based culture has emerged where DLC personnel have in some cases learnt who is the "right person" to contact with certain questions. This creates a number of problems. Firstly, if the "right person" moves function or leaves the organization, the process for requesting support is unclear. Secondly it is difficult to train new individuals to provide support in IS&T as there is no clear point to integrate them in to the support process. Finally DLC personnel new to the Institute have a hard time engaging IS&T as they have no relationships and knowledge of who to contact.
- More cost effective solutions for hosting small servers at W91. IS&T provide excellent hosting options for large enterprise scale servers, but the pricing for hosting smaller servers is thought to be prohibitive by some DLCs.

4. Future State

The Future State section contains the first iteration of documentation on the future enterprise architecture. Included in this section are:

- A Future State Vision diagram
- Technology Guidelines
- A forward looking Services Matrix
- A summary of integration options

Logical Architecture Vision

The Logical Architecture Vision is a conceptual model evolving the enterprise architecture at MIT in the future. As such it raises as many questions as it answers, and provides a good forum for future discussions about the architecture.

- The **Data** layer illustrates the idea that enterprise data (not departmental data or departmental extensions to enterprise data) should be logically visible and consolidated by data domain. This does not imply that, for example, all student data must be ph ysically located in the same database, or managed by the same system, but that there are a coherent set of rules for locating, a unified view of, and a standard way to access student data.
- The **Data & Business Integration Layer** speaks to the fact that there should be a consistent way of interacting with enterprise data, and a coherent strategy for sharing data across systems in the case that shadowing data is still necessary. This layer is key to achieving the goals of the data layer, i.e. logical consolidation and access to data by domain

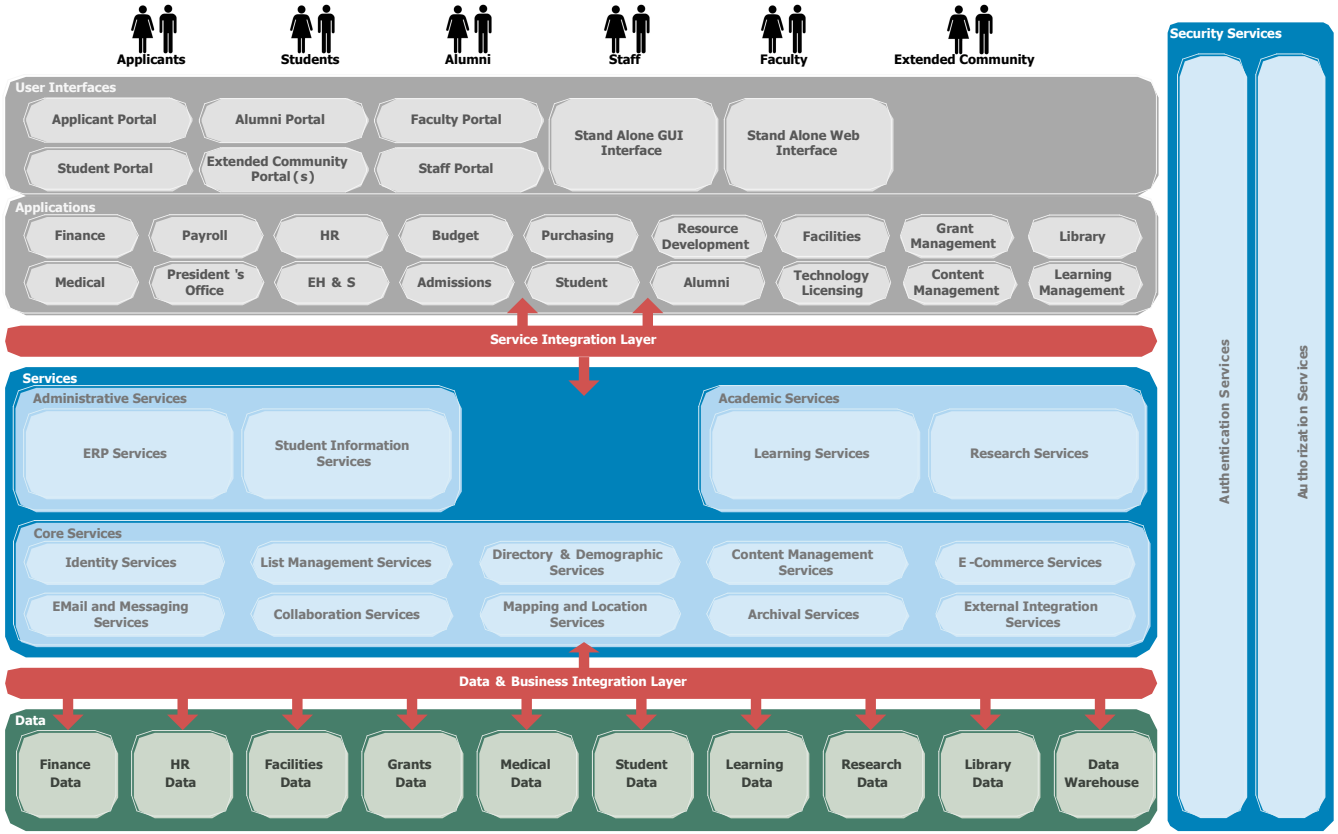

- The **Services** layer represents the separation of re-usable services from application logic. The service layer will thus consist of services with clearly defined contracts that can be used by any application. The services have initially been classified in to three major groupings: Core Services for technical and basic services, Administrative for services that provide access to administrative data or processes and Academic Services for services specific to the education and research domains.
- The **Service Integration** layer is responsible for exposing the services in the architecture in a consistent manner while enabling services to be implemented in a variety of technologies. Ideally it should also define the standard contract for a service type, therefore allowing substitution of service implementation without affecting clients of the service.
- The **Applications** layer shows groupings of applications that are built for specific purposes. By leveraging services available in the architecture, applications should generally be quicker to develop and easier to maintain.
- The **User Interfaces** layer conveys the idea that users should have a single point of access for related functions that they use. This might be implemented as one or more portals.
- The **Security Services** layer, while conceptually similar to other types of services has been shown separately because it has significant impacts at all levels within the architecture. It will be necessary to apply access and control security to data, to services, to applications and finally to user interfaces.

Instructions for Technology Guidelines

Description	The Technology Guidelines spreadsheet captures information about what technologies are in use today and are suggested for use in several scenarios in the future.
Version	0.1
Worksheet Definitions	
Technology Guidelines	The Technology Guidelines; see above.
Column Definitions	
Existing Standards	What is in wide use at MIT today and is therefore a de facto standard.
Recommendations for New Development of Enterprise Class Systems	Recommendations for technologies to use in developing enterprise systems in the future so that they complement and contribute to the enterprise architecture vision.
Recommendations for Department Mission Critical Systems	Recommendations for technologies to use in developing department mission critical systems so that they can be adequately supported in the enterprise environment.
Recommendations for Other Systems at MIT	General recommendations for non-enterprise and non-mission critical systems development.

MIT Technology Guidelines

	Existing Standards	*Recommendations for New Development of Enterprise Class Systems*	*Recommendations for Department Mission Critical Systems*	*Recommendations for Other Systems at MIT*
Server Hardware	Sun / Sparc Dell / Intel	*IF IS&T will maintain		
Server Operating System	Sun Solaris Red Hat Linux Windows (NT, 2000, 2003) Mac OS VM/CMS Novell Netware	Solaris Linux *(Need to develop criteria for deciding between recommended operating systems for projects)*		
Database Software	Oracle PostGres SQL Server 2000 My SQL Ingres FileMaker Adabas DB2	Oracle PostGres *(for small scale applications; need to determine transaction volume that qualifies as small scale)*		
Application Server Software	Tomcat / JRun / JBoss BEA WebLogic (5.0 - 8.1) Microsoft .NET	BEA WebLogic?		
Web Server	Apache IIS AOL Server	Apache		
Reporting Software	BRIO (Hyperion) Access Crystal Reports MS Query			
Desktop Operating System	Microsoft Windows (NT, XP, 2003) Mac OS Linux	Windows XP MAS OS 10		
Web Browser	IE (5.x, 6.x) on Windows Netscape Mozilla Safari Opera Lynx	IE Safari Mozilla Lynx *(in cases where meeting accessibility guidelines is important)*		
Programming Languages	Java C# PowerBuilder VB C Cobol etc……	Java -> primary C -> only if necessary		
Scripting		Perl		

Instructions for Services Matrix

Description	The Services Matrix documents the services that are currently offered within the MIT Enterprise Architecture and MIT IT community.
Version	0.1
Types of Service	
Remote Service	A service is a piece of software functionality that is executing somewhere separate from your program that you can invoke in a defined manner.
Embeddable Service	SDKs & Libraries are reusable components that can be embedded into an application, but there is no central infrastructure for them.
Business/Operational Service	A Business/Operation Service involves interaction with a person or staff member.
End user Service	Services to end users and not applications are those which are used for functionality but have no programmatic access.
Worksheet Definitions	
Systematic Services	The catalog of services that are for use by applications, and are relevant to a services oriented enterprise architecture
Non-Systematic Services	The catalog of services that are for use by people and impact the enterprise architecture only peripherally
Column Definitions	
Service	The short name or description of the service.
Definition	A longer definition of what the service is.
Currently Used By	A representative sample of systems which currently make use of the service.
Current Service Type	Whether the service is currently a Remote Service, an Embeddable Service, a Business/Operational Service or an End User Service
Current Implementation	A description of the current implementation of the service, usually including the sytems which implements it.
Future Service Type	Whether, in future, the service should be a Remote Service, an Embeddable Service, a Business/Operational Service or an End User Service
Future Form	A description of how the service should be implemented in the future.
Areas for Investment	Areas in a current service, or anticipated future service that require investment in order to realize.
Contact	The name of the person who should be contacted for more information about the service, for example if you wished to used it from a new application.
Notes	Any further notes about the service or follow up items.

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
Infrastructure									
Security									
Authentication									
Authenticate a User	Allow an application to authenticate the user (i.e. assert that they own the identity supplied).	Barton, SAP, COEUS, MITSIS etc.	Remote Service	MIT Kerberos	Remote Service	MIT Kerberos	• Version 4 needs to be eliminated. • Keep up with ITEF Protocol.	Jeff Schiller	
Authenticate a User	Allow an application to authenticate the user (i.e. assert that they own the identity supplied).	COEUS, Stellar, Dspace, Barton, Ecat etc.	Remote Service	X509 Certificate	Remote Service	X509 Certificate		Jeff Schiller	
Password Reset	Allow a user to reset their password when locked out and attempting to access an application.		Business / Operational Service	User must present ID in person at Accounts Department in N42.	Remote Service	Unknown	Get statistics on man hours/money needed.		
Authorization									
Roles (direct access)	The Roles Database provides a consistent way to store and maintain access rules for other applications. Applications with an interface to the Roles Database interpret the access rules from the Roles Database and enforce them.	SAP, Data warehouse	Remote Service	Roles Database	Remote Service	Roles Database	Need a service interface that is higher level than the current SQL based access.	Jim Repa	*To Do: Break this line item up in to the actual services provided by Roles*
Identity									
Create MIT IDs	The MIT ID is a 9 digit number used to uniquely identify any member of the MIT community. An MIT ID can be created through the MIT ID Database web client.		MIT ID - Remote service			MIT ID - Remote service	Need to link or consolidate IDs, when someone is a student + alum + employee		
Retrieve MIT IDs	The MIT ID can also be retrieved through the MIT ID Database web client by supplying a person's first and last name.	Medical	MIT ID - Remote Service			MIT ID - Remote service			
Network									
DHCP	The DHCP (Dynamic Host Configuration Protocol) Service lets a user connect his/her computer to MITnet from a variety of sites on campus without reconfiguring his/her computer's network settings each time the computer is moved to a new location.		Remote service						
DNS	The Internet Domain Name Service (DNS) can translate host names into equivalent IP addresses and vice versa, as needed by various Internet programs.		Remote service						
Host ID Management			Remote service						
Messaging & Communication									
SMTP(S) Servers - Email Transmission	Outgoing mail servers are referred to as SMTP servers. The outgoing mail server at MIT is named outgoing.mit.edu.		Remote service						

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
IMAP/POP3 Servers	IMAP (Internet Message Access Protocol) is a standard set of rules for storing, accessing and working with e-mail on a post office server. One of the main advantages of IMAP is that it makes your e-mail easily accessible from multiple locations and computers. POP (Post Office Protocol) is a set of rules for storing and accessing your e-mail on a central server. When you access messages, they are downloaded to your local computer (or Athena home directory) and deleted from the server.		Remote service						
List Management (Mailman, Moira)	Mailman, Moira and Listserv are used to manage mailing lists at MIT.		Remote service						
MIT EDI Gateway	EDI is the electronic transfer of information between two trading partners' systems using a set of transactions that have been adopted as a national or international standard for the particular business function.		Remote service						
Data									
List Management									
Moira (direct access)	Moira is Project Athena's Service Management System. It controls the configuration of resources, including user accounts, remote file systems, printers, mailing lists, access control groups, and many other things.		Remote Service	Moira	Remote Service	Moira	List management needs to be integrated with Roles		*To Do: Break this line item up in to the actual services provided by Roles*
Content Management									
Directory									
Repository Services									
AFS - Remote File Service	AFS, the Andrew File system, is currently used by Athena as the file system for all user home directories and most of the other lockers. AFS is a distributed file system.		Remote service						
Information feeds from the Data Warehouse			Business/Operation service SQL - Remote service				XML Web Service for data retrieved		
Development Tools									
Change management									
Operations									
System Management									
Web Counter	The web counter service allows data gathering of the number of hits made to any web application.		Remote service				Question? Can this be used for sites outside Athena? Where can it be used?		
Admin									
eCommerce									

Services Matrix - Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact	Notes
Clear Commerce Credit Card Processing	Clear Commerce is an enterprise software that sends transaction information to MIT's bank for verifying and processing payments on customers' credit cards.		Remote service						
Shopping Basket			Embeddable service						
Academic									
Learning management									
Educational Application									
Other									

Services Matrix - Non-Systematic Services

Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
Infrastructure								
Security								
Authentication								
Acquiring a Server Certificate	Allow a system to use a certificate to identify and authenticate itself to a user.		Business / Operational Service	X509 Certificate	Business / Operational Service	X509 Certificate		Jeff Schiller
Password Reset	Allow a user to reset their password when locked out and attempting to access an application.		Business / Operational Service	User must present ID in person at Accounts Department in N42.	Remote Service	Remote service	Get statistics on man hours/money needed.	
End user certificate management		Stellar	End user service					
Authorization								
Encryption								
Identity								
Network								
Messaging & Communication								
Web Casts	Events at MIT can be cast via the Web by using a fee-based service offered by AMPS.		Business/Operation service					
Data								
Reporting Service								
Reporting through the Data Warehouse	Reporting is done in the Data Warehouse through Brio or web reports. In addition, some departments have built their own custom applications to access the Data Warehouse for reporting purposes.		End user service					
Content Management								
Research content archival			Business/Operation service					
Directory								
MIT Online Directory	The MIT Online Directory allows a user to search for and view information about people in the MIT Community.		End user service					
Repository Services								
Information feeds to the Data Warehouse	The Data Warehouse is a storage for any data in the Institute which needs to be accessed by multiple systems. The Data Warehouse is updated daily by systems of records. Other systems can then extract this information from the Data Warehouse as appropriate.		Business/Operation service					
Information feeds from the Data Warehouse			Business/Operation service SQL - Remote service				XML Web Service for data retrieved	

Services Matrix - Non-Systematic Services

	Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
	GIS	GIS (Geographic Information Systems) are computer tools for managing data about where features are (geographic coordinate data) and what they are like (attribute data), and for providing the ability to query, manipulate, and analyze those data.		End user service			Web Service/Remote Service		
Development Tools									
	Change management								
Operations									
	System Management								
	Co-Location Services	Co-Location Services allow MIT applications and servers to be located in a separate location to allow for backup and recovery in case of any failure.		Business/Operation service					
	Server Monitoring	Server Monitoring is a service provided by IS&T to monitor servers for various applications in W91.		Business/Operation service					
	Consulting Services	IS&T offers various consulting services to the MIT Community.		Business/Operation service					
	Sub Domain Management			Business/Operation service					
	Backup and Restore			End user service					
	Issue Tracker			End user service					
Admin									
	eCommerce								
Academic									
	Learning management								
	Educational Application								
	Matlab	Matlab is a technical computing environment for high-performance numeric computation and visualization, produced by The MathWorks Inc. It includes a number of subject specific toolboxes as well as a dynamic system simulation package, Simulink.		End user service					
	Calendaring	A Calendaring service allows a user his/her schedule, or agenda in Calendar parlance, and also coordinate easily with the schedules of other users of the same Calendaring service.		End user service			Remote service		
	Events.mit.edu	This web site displays the events at MIT for the current day. It also allows the user to view upcoming events in various categories.		End user service					

Services Matrix - Non-Systematic Services

	Service	Definition	Currently Used by	Current Service Type	Current Implementation	Future Service Type	Future Form	Areas For Investment	Contact
	Campus Map	The campus map pinpoints where you are, and where you're going. The campus map uses geographic information systems (GIS) data from the official maps maintained by the Department of Facilities, resulting in a more accurate mapping system. Using XML, the map is also integrated with the lists of departments, labs, and centers on the MIT top-level pages.		End user service					
Other									
	Video Production	Video production is a fee-based service offered by AMPS to the MIT Community.		Business/Operation service					
	Survey Service			Business/Operation service					

Enterprise Integration Options

Several options exist for integrating applications within the enterprise. This section gives a brief outline and comparison of the major options.

Option 1: Custom Point-to-Point Integrations

Definition: A direct point-to-point link is created between applications for each business function

Attributes:

- Designed and implemented solely for the purpose of directly integrating two specific systems
- Custom code for
  - Data extraction
  - Business rule processing
  - Data loading
- Custom data format
- Inconsistent transmission protocol
- Used for both real time and batch integrations

Advantages:

- No need to invest in expensive tools up front
- No need for developers to learn new skills and packages
- No extended time frame for developing and deploying enterprise integration strategy

Option 2: Messaging or EAI (Enterprise Application Integration) Tools

Definition: Source systems "publish" enterprise messages to a common bus; application "subscribe" to relevant messages and act on them

Attributes:

- "Wraps" each application
- Acts as a broker between applications
- Provides near real-time, guaranteed, once-only delivery
- Stores and forwards messages
- Provides an environment in which to define rules

Advantages:

- Systems are integrated but not coupled
- Business rules are centralized in the message broker and transformation engine
- Allows for near real-time integrations which reduced latency
- Solves the n^2 problem; as the number of systems increases, the integration effort expands linearly

Option 3: Web Services

Definition: Functionality to be integrated is exposed via XML on an open protocol such as SOAP. Other systems can consume this service if needed. Inputs and outputs to the web services are XML

Attributes:

- Common language of communication across heterogeneous systems
- Based on standard Internet technologies
- Self describing and advertising
- Supports dynamic discovery & integration
- Services fit within an overall architectural model
- Widely supported by major vendors

Advantages:

- Solves problems similar to those EAI solves, as well as
- Need for expensive integration tools
- Use of proprietary integration platforms

Option 4: ETL (Extract, Transform, Load) Tools

Definition: A standard set of tools and processes used to extract, transform and load large volumes of data between systems. Very useful in populating a data warehouse

Attributes:

- Provides tools for data cleansing; correcting misspellings, resolving conflicts (city & zip code incompatibilities), missing elements, parsing elements
- Can combine data sources: Matching on key values, fuzzy matches on non-key attributes, textual comparisons to reference tables
- De-duplicate processing: Identifying and eliminating duplicates
- Can create surrogate keys: Operational systems and the data warehouse have different assumptions and data requirements thus the data warehouse requires its own set of primary keys
- Create aggregates to boost performance of common queries in data warehouses and data marts
- Loading and indexing: For large data warehouses specialized bulk loading processes are required

Advantages:

- Extremely efficient for moving large volumes of data in short timeframes
- Applies consistent transformations
- Can provide or integrate with meta-data for the enterprise data model

Comparison of Options

	EAI	Point to point	Web Services	ETL
Concept	• Publish/Subscribe mechanism • Most suitable for real time data needs • Loosely coupled	• Custom code for each integration need • Suitable for complex integration needs • Tightly coupled	• Standards based integration • Most suitable for inter-organization integration • Loosely coupled	• Suitable for large volumes of data • Generally used to move data between two or more databases
Strengths	• Reliability (guaranteed delivery) • Enables real-time business decisions • Out of box adapters for many enterprise systems	• Familiar technologies and processes • Many point to point integrations already exist • No major up front investment required	• Standards based integration • High degree of reuse • Wide tool support including open source • Low up front investment	• Metadata driven approach • GUI tools for most tasks (little coding) • Extremely efficient for large data volumes
Weaknesses	• High upfront cost • Relatively complex design patterns	• Costly over time • Tight coupling • Scalability issues • Opportunities for reuse are slim	• Lack of transaction support • Not a publishing model • Less established technology	• High upfront costs • Complexity of tool • Batch oriented
When to Use	• Real time data is important • High volume, low footprint data exchange • Many consumers of the same data	• Should be rarely used • When defined enterprise strategy cannot work • Proto typing	• Integration model is request/reply • Real time requirements • High volume, moderate data	• In conjunction with a data warehouse

Open Knowledge Initiative (OKI)

MIT is major contributor to an open source project called the Open Knowledge Initiative (OKI). OKI develops specifications with the goal of creating technology neutral, standard client interfaces for a variety of services. The focus is on specifications for educational software, but this extends to core technology services that are used in educational environments. The vision for OKI is that a system which leverages OKI interfaces for interaction with enterprise services should be able to be moved from one education environment to another and integrated with the local service offerings with simple configuration changes.

The goals of OKI are:

- Allow enterprise systems to exchange & synchronize information
- Allow different organizations to exchange & synchronize information
- Allow systems to use enterprise services
- Allow for modular software which plugs into a known framework
- Single system responsible for information

An OKI OSID (Open Service Interface Definition) defines the client contract for a specific service type. The current collection of OKI OSIDs are:

- **Common Services**
  - Authentication
  - Authorization
  - SQL
  - Logging
  - Shared
  - Filing
  - Dictionary
  - Hierarchy
  - Agent
  - ID
- **Extended Services**
  - User Messaging
  - Scheduling
  - Workflow
- **Educational Services**
  - Course Management
  - Digital Repository
  - Assessment
  - Grading

For more information on OKI please visit:

- http://www.okiproject.org
- http://sourceforge.net/projects/okiproject

5. Architecture Project Review Process

Process for enterprise projects

Architecture Project Reviews

What is the Purpose of the Review?

The ITAG Review is a review of project architecture during the design and development phases of an application system to help ensure a successful project implementation. This review also ensures that the proposed system fits into the existing MIT environment as well as the future architectural vision. There are multiple benefits of these reviews:

Project Success

- Architecture is reviewed by a group of senior architects across the institute
- Assistance in leveraging the existing architecture
- Proactively identify risks to the project
- Provide Institute-wide context to project team

Maintain the integrity of MIT's IT environment and expand the user community's access to MIT resources

- Does the project present risk to the IT environment (e.g., infrastructure, other applications, users, state and federal regulations, or institute policy)?
- Allows ITAG to proactively recognize when modifications to the architecture are required
- Allows the project team to provide input to the extension of the MIT architecture
- Does the project leverage the existing common services where applicable?
  - Provides cost effectiveness across the institute, or
  - Provides integration benefit (e.g., single sign-on)

How can ITAG assist you?

- Facilitate access to data from system of record
- Identify opportunities for leveraging institutional resources

What is the Review Process?

1. Project announcement forwarded to <itag.mit.edu> at project initiation
2. Project Review Board determines whether a review is necessary (if no, send email waiver)
3. Project Review Board assigns an ITAG Project Review Coach
4. ITAG Review Coach works with project to gather information and recommend a review timeframe
5. Project submits review material (at least one week prior to review)
6. ITAG performs review (normally 30 – 60 minutes in length)
7. ITAG Review Coach distributes review summary and recommendations (no later than 1 week after the review)
8. Project needs to re-engage ITAG Review Coach if any of the following occur:
   - Proposed architecture changes
   - Issues resolving the action items

What projects need to be reviewed by ITAG?

Applications that are community or enterprise wide (e.g., non departmental applications)

Applications that are likely to expand beyond the initial departmental scope

All IS&T projects (reviews could be a subset ITAG, although everybody would be welcome)

All projects in an ITAG member's area

Applications that utilizes sensitive data

Applications that create and store user IDs and passwords

What projects do not need to be reviewed?

Departmental applications (e.g., Totally self contained with a DLC)

- No intention or likelihood of expansion outside the department

Projects within a DLC whose data does not affect an Institute data entity.

- No intention or likelihood of data exchange

When is the Review?

Project has articulated the business goals

Project has a conceptual idea of hardware and software requirements

Hardware and Software decisions have not have been finalized

Designs have not been finalized (drafts should/may be completed)

Budgets have not been finalized

Reviews should also be conducted BEFORE a contract is signed

What preparation should be done by the project team?

During the review, please be prepared to answer questions such as:

What impacts does your project have on the community or other systems?

- What users are impacted by your project?
- What other systems will your project affect?
  - Data feeds

What IT services are required?

What security concerns exist?

- Will the system be handling sensitive data now, or in the future?
- Is a separate security review required?

Required Materials

Scope

Timeline

Context Diagram

- Major functions and connections to other systems
- Layered diagram showing major functions and connections to functions of the infrastructure (e.g., Roles, Single Sign-on, etc.)

Conceptual Data Model

- Major entities, their definition, and relations and cardinalities

ITAG architectural questionnaire

6. Moving Forward

Moving Forward

The following areas will need to included in the next version of this guide:

- Alumni
- Housing
- Human Resources
- Resource Development
- Student Services

Topics to be considered for inclusion into future versions of this guide:

- Encryption Strategy
- Kerberos ID Policy
- Data Ownership and Polices
- OKI Implementations
- Role of Open source packages at MIT
- Role of Filemaker at MIT
- Planning for applications leveraging SAP 3-5 years out

7. Key Systems Overview

Instructions for Key Systems Inventory

Description	The Key Systems Inventory details the Enterprise Systems at MIT.
Version	0.1
Worksheet Definitions	
Key Systems Inventory	Key Systems Inventory; see above
Column Definitions	
Summary	A Summary of the system functionality
Maintained By	The Department in charge of the system
Contact	The full name of the contact person of the system
Accessed Off Campus	Details if the system can be accessed Off Campus or Not
Intended Availability	Details the hour availability of the system
Charge for Service	Details any charge for using the system
Primary User Base	Details the primary users of the system
Secondary User Base	Details the secondary users of the system, if any
Hardware Platform	Details the Hardware Platform the system is built on
Operating System	Details the Operating System of the system
Database	Details the Database used in the system
Services used/Available	Details any services the system uses or provides
Other Key Technology	Details any other relevant key technology of the system
Hosting Location	Details where the systems is hosted
Future Plans	Details any plans for the system in the near future

Key Systems Inventory

	Summary	Maintained By	Contact	Accessed Off Campus	Intended Availability	Charge for Service?	Primary User Base	Secondary User Base	Hardware Platform	Operating System	Database	Services used /Available	Other Key Technology	Hosting Location	Future Plans
SAP	MIT's Primary ERP Solution. Two Installations exist. One for Main Campus and on for Lincoln Lab. (This row refers to details of the installation on Campus)	Administrative Computing: Technical Services	Kevin Lyons, Ron Parker, Steve Landry	Yes, with a web certificate or kerberos principle	24 X 7 with down time weekly: Sat 1:00 am - Sun 6:00 am	No	Employees and Faculty	Anyone with an MIT ID and authentication principle	HP ES-20s & ES-40s	Tru-64, v5.1.a transitioning to v5.1.b soon	Oracle 8i transitioning to 9i soon	Roles, X509 Certificates, Kerberos		W91	SAP-BUD, migration of Nimbus functionality to SAP by Dec 2004.EHS SAP project: inspection and audit phase by Feb 2005, needs assessment by June 2005. Employee payroll into SAP by Jan 2006. Hardware renewal for the SAP platform in 2005. May also be an upgrade of the SAP version around the same time.
Summit															
TLO															
MIT Student Payroll															
Mainframe Payroll															
COEUS	MIT's Grant Management System. The main modules in Coeus are: Proposal Development, Proposals, Awards, Subcontracts, Negotiations, Person, Conflict of Interest and Report Tracking.	Office of Sponsored Programs	Steve Dowdy	Yes	24 x 7	No	Research Administrators, Controller's Office, All other administrative officers	PIs (Principal Investigators	One Solaris box for Web and Application Tiers (2 processor USparc III, 2GB Memory), One Solaris box for the Database Server (2 processor USparc III, 2GB Memory	Sun Solaris	Oracle 9i	X509, Kerberos,	Power Builder Interface (will become obsolete in the future), JSP/Struts, Java Swing Client	W91	Real time integration with SAP, Evolve Human Subjects to create a PI portal view of the world, Possibly Animal Care functionality

Key Systems Inventory

	Summary	Maintained By	Contact	Accessed Off Campus	Intended Availability	Charge for Service?	Primary User Base	Secondary User Base	Hardware Platform	Operating System	Database	Services used /Available	Other Key Technology	Hosting Location	Future Plans
Nimbus	NIMBUS is the system of record for general and auxiliary Institute budgets. It is integrated with SAP and maintains and feeds budget information to SAP and the MIT Data Warehouse.														
Ecat	ECAT is MIT's Web-based system for purchasing directly from selected MIT partner vendors. ECAT is integrated into SAP web for quick creation of requisitions.														
Maximo	Facilities' legacy system, for the processing of repair and preventative maintenance orders. (SAP Plant Maintenance replaces Maximo)	Facilities	Mike Sherman	Yes	24 x 7	No	Facilities, Housing, Electric Plant	Students							
Insite	Space Accounting System.	Facilities	Mike Sherman	No	24 X 7	No	MIT Students/Faculty/Staff		Windows 2000	Windows 2000	Oracle		IP Screening to restrict to on campus access	W 91	

Key Systems Inventory

	Summary	Maintained By	Contact	Accessed Off Campus	Intended Availability	Charge for Service?	Primary User Base	Secondary User Base	Hardware Platform	Operating System	Database	Services used /Available	Other Key Technology	Hosting Location	Future Plans
SMS	Space Management System. Allows people to model scenarios.	Facilities	Mike Sherman		24 X 7	No	MIT Students/Faculty/Staff				Oracle		Campus Map Service	W 91	
PPL Keys	Manages the inventory of all keys on Campus.	Facilities	Mike Sherman						VAX		Ingress			W91	
Practice Management System	This system is responsible for the operational aspects of the Medical Center.	Medical Dept.	Alison Knott, Kin Lie	Yes	24 X 7 with a short down time daily at 2am	No	Medical department and MIT Health employees		Dell	Windows 2000	Cache	Roles	IDX Application plus custom modules	E-23	Upgrading the to FlowCast, the latest generation of the IDX PMS system, with a web interface. In the process of making everything totally redundant within the health center.
Electronic Medical Records	An implementation of Allscripts' Touchworks product and is the central repository for all information pertaining to a patients medical record	Medical Dept.	Alison Knott, Kin Lie	Yes	24 X 7	No	Medical department and MIT Health employees		Dell	Windows 2000	SQL Server	Roles	TouchWorks by AllScripts	E-23	In the process of making everything totally redundant within the health center.
Patient Online	Patient On-Line is an online patient portal that is, at time of writing, limited in functionality.	Medical Dept.	Alison Knott, Kin Lie		24 X 7	No	Patients		Dell						Patient Online is a major ongoing initiative.
E-Scription	This system	Medical Dept.	Alison Knott, Kin Lie		24 X 7	No	Medical department and MIT Health employees		Dell						

Key Systems Inventory

	Summary	Maintained By	Contact	Accessed Off Campus	Intended Availability	Charge for Service?	Primary User Base	Secondary User Base	Hardware Platform	Operating System	Database	Services used /Available	Other Key Technology	Hosting Location	Future Plans
Corporate Database															
ILP															
Barton	Barton is a Library Management System and is essential to nearly all aspects of operation of the library at MIT.	Libraries	MacKenzie Smith	Yes	24 X 7	No	Barton Front End: All of MIT; Barton Business Back End: MIT Staff	Barton Front End: All of the general public	Sun E450	Sun Solaris	Oracle	Kerberos, IS&T Gold Contract for Maintenance		14S	Enhancements to the web UI of Barton.
EHSWeb	Tracks employee training needs and trainings attended for health and safety	EHS	Jim Repa	Yes	24 X 7 with a 1/2 hr nightly backup	No	Anyone who needs to take training: Researchers/Faculty/Students, EHS Office, Admin/PI in DLC	Emergency Response People	Sun	Sun Solaris	Oracle	X509, Roles		W91	Migrate PI/SARA to SAP
Card System															
Tech Cash															
Parking System															
Admissions Portal															
Undergraduate Admissions															
Graduate Admissions															
MITSIS	the MIT Student Information System, is an administrative system that serves all of the student services at MIT.	Student Services	Kent Dorsey	Yes	24 x 7 with daily downtime from 3am to 6am	No	MIT Students/Faculty/Staff		HP Alpha	OpenVMS	Oracle 7.3.4		SQL Forms 3.0 and SQR, Oracle Forms applications and batch programs run here, ProC, C, COBOL programs also run here	W91	
WebSIS	WebSIS, MIT's web-based Student Information System, provides students with full access to their financial and academic records.	Student Services	Kent Dorsey	Yes	24 x 7 with daily downtime from 3am to 6am	No	MIT Students, Faculty, Staff		HP Alpha	Tru64 Unix	Oracle 7.3.4		Netscape Web Server, SQR	W91	

Key Systems Inventory

	Summary	Maintained By	Contact	Accessed Off Campus	Intended Availability	Charge for Service?	Primary User Base	Secondary User Base	Hardware Platform	Operating System	Database	Services used /Available	Other Key Technology	Hosting Location	Future Plans
Advance															
Sloan Admin DB	Holds the Alumni data for Sloan School Alumni	Sloan School of Management	Al Essa			No									
Sloan Space	SloanSpace is a portal for the entire extended Sloan community.	Sloan School of Management	Al Essa	Yes	24 x 7	No	MIT Students, Faculty, Alumni, Staff	Industry Sponsors, Researchers	Sun B 440	Sun Solaris	Oracle		AOL Web Server, Oracle Database Server	W91	Reach a point where DOT LEARN is a stand alone project so if Sloan left the project would still go on.
Stellar	Stellar is a Learning Management System developed at MIT for managing the authoring and delivery of courses.	AMPS	Craig Counterman	Yes	24 x 7	No	MIT Students, Faculty, Staff		Sun	Sun Solaris	Oracle	Kerberos	Apache Web Server, Tomcat Application Server	W 91	
OCW	Makes the course materials that are used in the teaching of almost all MIT's undergraduate and graduate subjects available on the Web	OCW	Cec d'Oliveira	Yes	24 x 7	No	Anyone		Windows 2000		SQL Server 2000			Sapient	
DSpace	Provide long term storage options of digital products to MIT faculty and Researchers.	Libraries	MacKenzie Smith	Yes	24 x 7	No	MIT faculty and researchers		HP	Redhat Linux	PostgreSQL	X509		14S	
Vera	Allows access to MIT Library resources from off campus	Libraries	MacKenzie Smith	Yes but only for the MIT Community	24 x 7	No	MIT Students/Faculty/Staff		Sun	Sun Solaris	FileMaker			14S	

Key Systems Inventory

	Summary	*Maintained By*	*Contact*	*Accessed Off Campus*	*Intended Availability*	*Charge for Service?*	*Primary User Base*	*Secondary User Base*	*Hardware Platform*	*Operating System*	*Database*	*Services used /Available*	*Other Key Technology*	*Hosting Location*	*Future Plans*
EZProxy	Provides a way for libraries to extend web-based licensed databases to their remote users.	Libraries	MacKenzie Smith	Yes	24 x 7	No	MIT Students/Faculty/Staff							14S	
SFX	SFX is a tool to help you link from an article citation in a database to the full text of that article.	Libraries	MacKenzie Smith	Yes	24 x 7	No	MIT Students/Faculty/Staff								

Athena Environment

Athena is a collection of systems that are developed and maintained by MIT to support the core IT infrastructure of the Institute. They cover a number of service areas including authentication, authorization, network access, list management, remote file systems and more. The main components of the Athena solution are:

- **Kerberos** is a network authentication protocol initially developed by MIT to provide strong authentication for network applications using secret-key cryptography. The specification is now maintained by an industry body, but MIT continues to maintain and distribute an implementation of the protocol and use it internally as the primary authentication mechanism for enterprise applications. It provides a single sign on capability through the issuing of tickets to an agent running on a user's desktop; these tickets are then used to authenticate to a "Kerberized" application. Kerberos is designed for use with traditional client/server applications. MIT has extended authentication in to the web domain using a Certificate Authority, described later.

- **Moira** is primarily a list maintenance application. It is the repository for several types of data used by other systems within the Athena environment as well as enterprise systems at MIT. Moira is the authority for mapping MIT IDs (the primary form of electronic as well as physical identity token at MIT) to Kerberos principles that can be used for authentication. It also manages a set of lists that are used to provide enterprise authorization information to applications. Much of the data contained within the Moira database is fed periodically to the Data Warehouse for reporting and distribution.

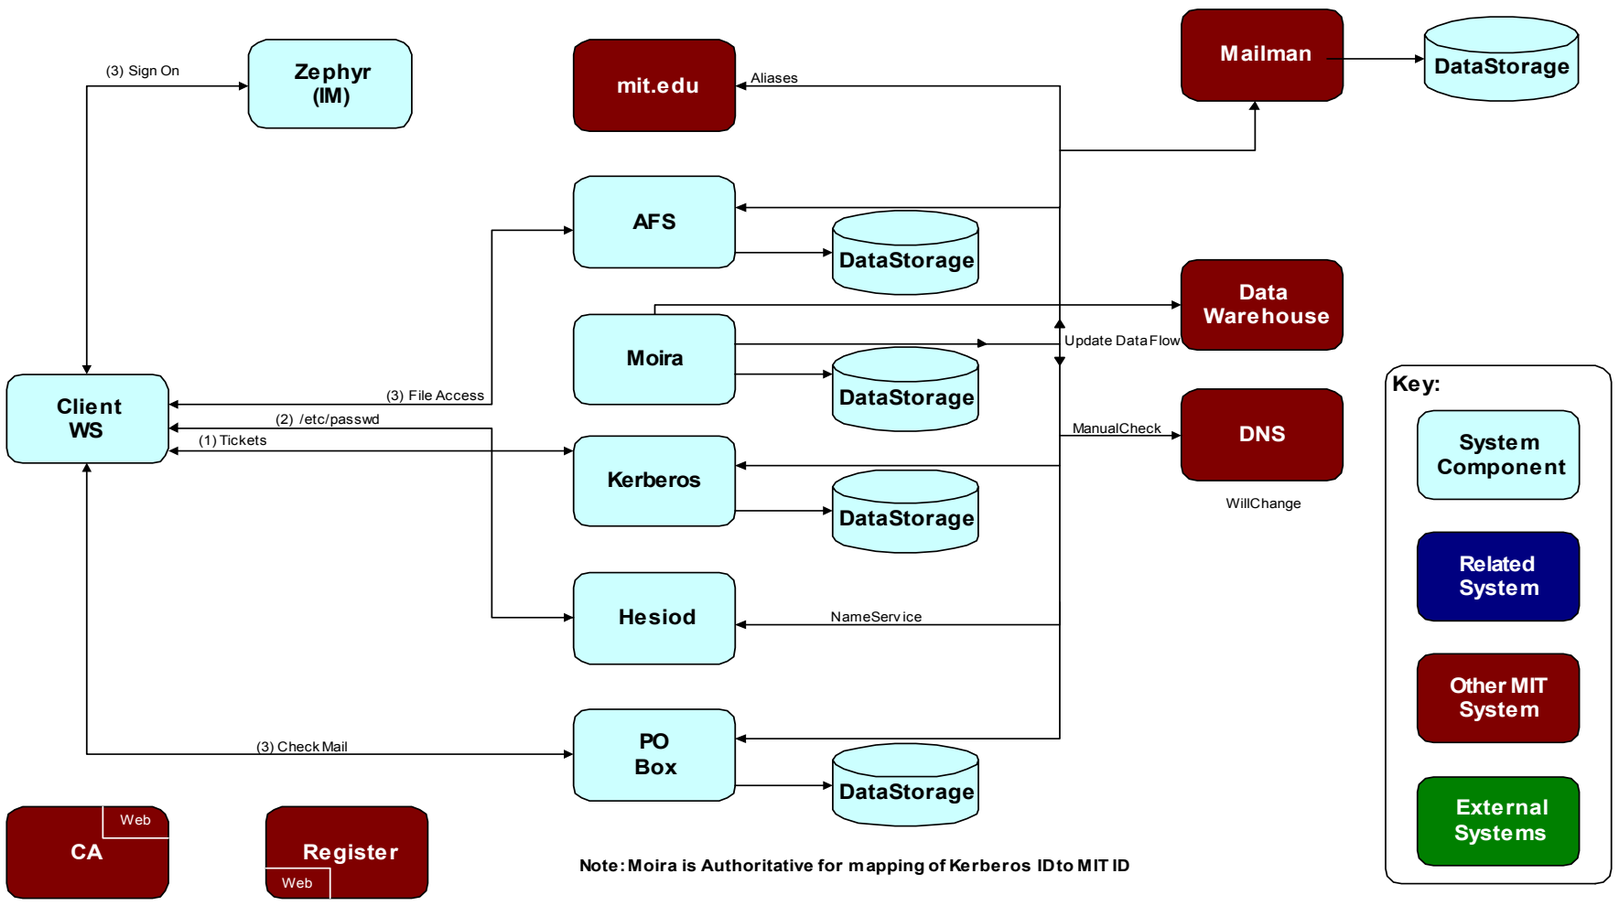

- **Hesiod** provides a mechanism for discovering various types of information within the Athena environment. It keeps track of post office servers, printers and machines. For example Hesiod provides the ability for a user to determine which post office server currently has their email; clients capable of interacting with Hesiod can therefore seamlessly manage the movement of a user's account between post office servers.

- **AFS** is an implementation of the Andrews File System. This provides a remote file system capability for use by computers within the Athena environment. It is used to store the home directories of users with Athena accounts and is the store for the web.mit.edu website. It is Kerberos aware, and hence can apply permissions based upon the Kerberos identity of a given user.

- **Zephyr** is a presence and messaging service supported within the Athena environment. When a user is logged on at an Athena workstation the system tracks this information centrally, and can route text messages between users logged on to the system.

- The **Certificate Authority** is a system for issuing X509 certificates. It can create certificates for identifying both web clients and servers. The Certificate Authority allows the Kerberos authentication principle to be extended for use with web applications; a user can identify themselves to the Certificate Authority using their Kerberos principal. Having authenticated themselves a Kerberos user can then create any number of certificates with validity periods of up to a year, for the purpose of authenticating themselves to web applications that use X509 certificates for authentication. This has the advantage of maintaining a single identity for systems using both Kerberos and X509, and therefore allows authorization and other types of information to be leveraged.

- The **Domain Name Service** takes information propagated from the Athena environment on a regular basis. Athena manages the assignment of domain names to computers on campus, both within and beyond the Athena environment. This information is fed to the DNS servers on a regular basis in order that they convey accurate information to clients.

Barton

Barton is a Library Management System and is essential to nearly all aspects of operation of the library at MIT. Its public web interface is used by the majority of students and faculty at MIT. The general public have access to the MIT library, and as a result are also users of the web interface to Barton. The public interface allows users to accomplish tasks such as searching the catalog, checking on availability etc.

Barton also has a GUI user interface to manage the internal functions of the library including catalogue management, purchasing and other administrative tasks. The primary user base of the GUI application are the library staff, although it can also be used by visitors to the library to perform self service checkout. The GUI application supports Kerberos as the mechanism to authenticate users.

The Barton application is an implementation of the Aleph 500 Integrated Library System. More information on this product can be found at: http://www.exlibrisgroup.com/aleph.htm

Barton Logical Architecture Diagram

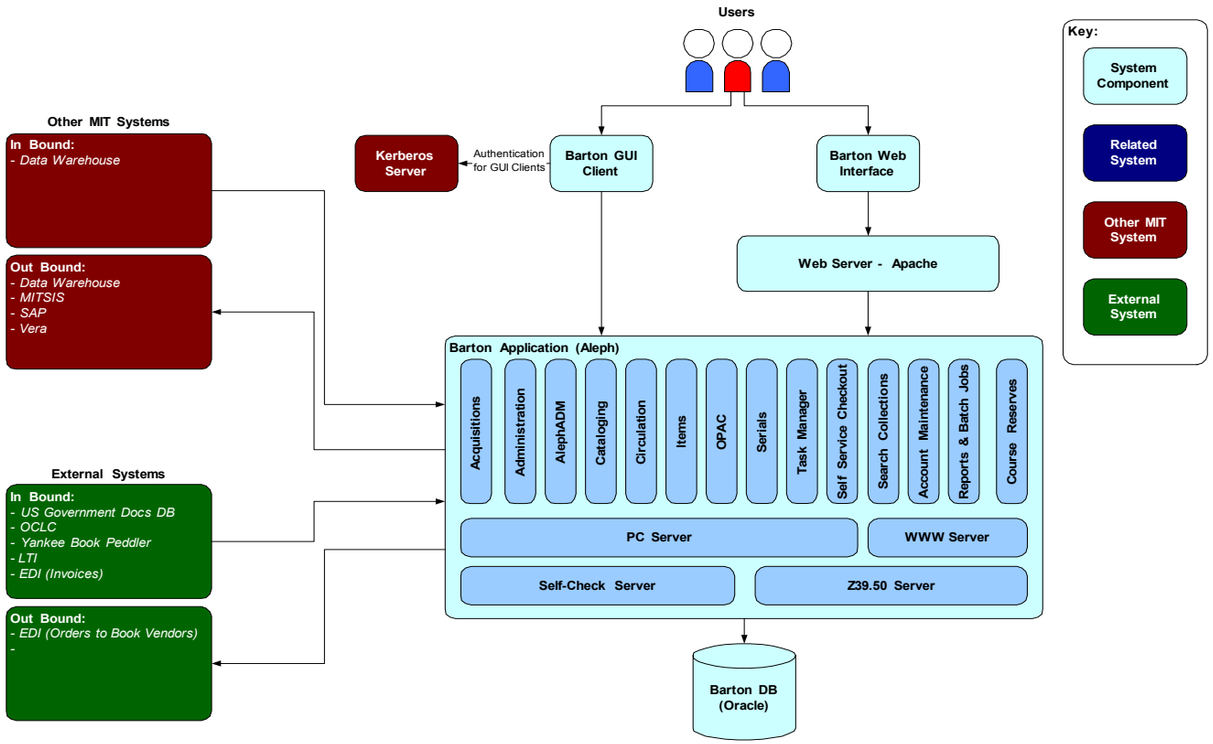

Barton Physical Architecture Diagram

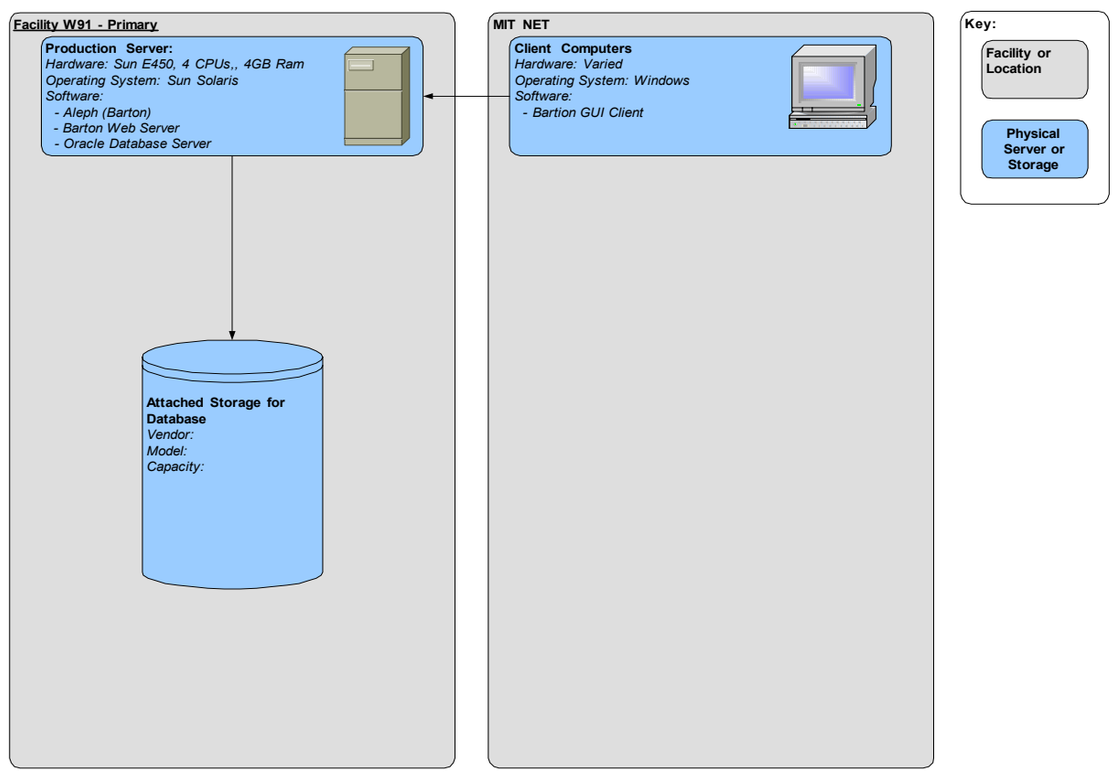

COEUS

COEUS is an application developed internally at MIT to manage grants. There are several aspects to this process that are managed by COEUS

- Proposal Development and Submission
- Awards Tracking
- Grant Accounting

COEUS was until recently exclusively accessed through a PowerBuilder client, which accessed an Oracle database. The result of this initial architecture is that significant amounts of business logic reside in PL/SQL stored procedures in the Oracle. This has proven valuable as it enabled a recent re-write of the application from PowerBuilder to a three tier Java architecture. The current COEUS architecture makes use of a set of application components residing in an Apache Tomcat instance. The components are invoked by the JSP/Struts web interface to create a web based interface to some portions of COEUS; the web interface is a "light" interface that does not offer the full functionality of the COEUS application to the user.

The complete user interface for COEUS is a Java Swing GUI client. This enables the user to accomplish any task that COEUS is capable of. The Swing client invokes the same application components as the web interface. The Swing client is a direct replacement for the PowerBuilder client, which will probably be retired in the near future.

The Swing application uses Kerberos for authentication, though not in the standard way. Problems with the Java Runtime Environment on windows mean that a Swing application cannot easily access the Kerberos ticket cache located on the user's machine. Therefore the Swing application must re-activate the Kerberos login process for the user and create a separate ticket cache for the application. This results in the application using the secure and well tested Kerberos infrastructure and eliminates the need for a separate set of credentials for the user. It is expected that in future, when the Java Runtime Environment compatibility problems have been solved, that COEUS will become a fully Kerberized application.

COEUS currently uses the MIT EDI service to communicate with a government agency for both sending proposals and receiving award information. It is expected that in the near future this will be migrated to a direct integration to the grants.gov infrastructure using Web Services.

COEUS Logical Architecture Diagram

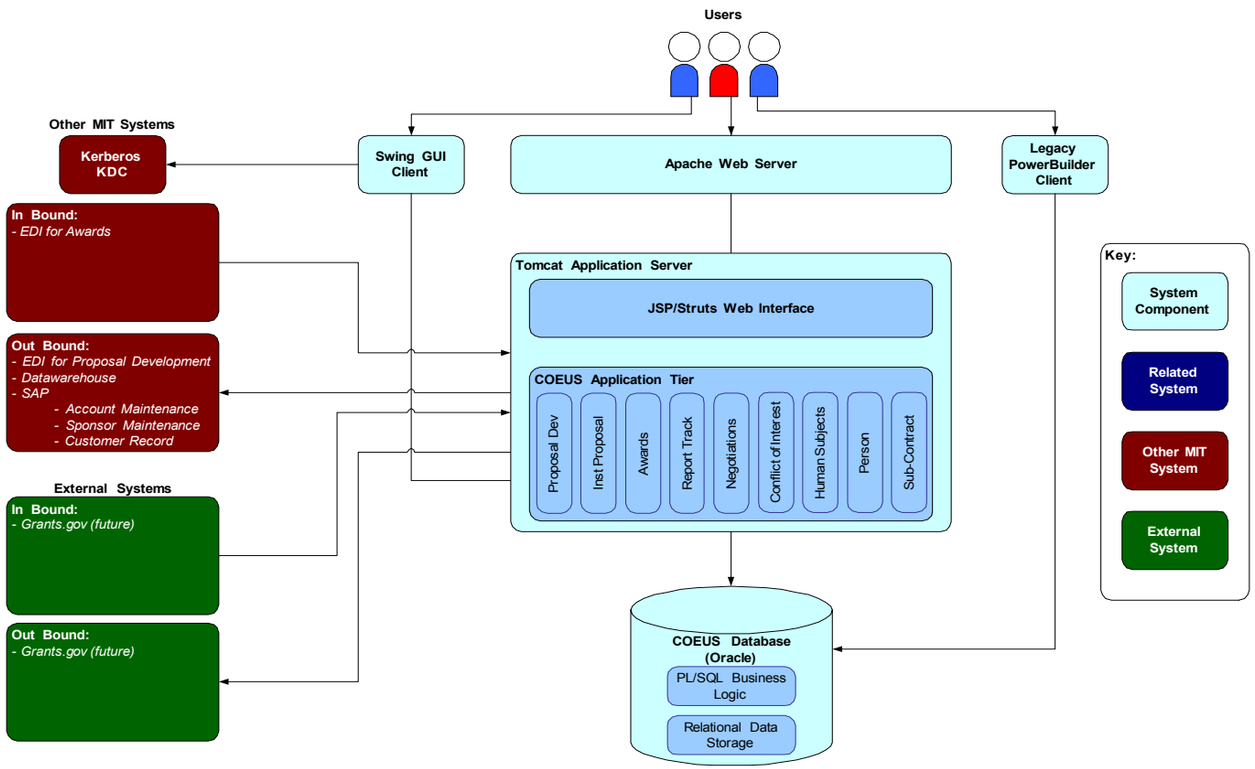

COEUS Physical Architecture Diagram

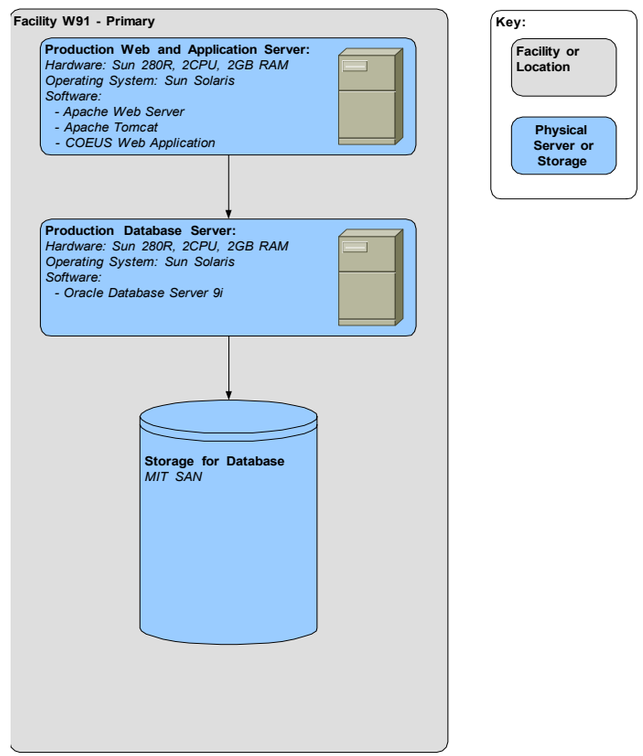

Data Warehouse

The Data Warehouse provides the MIT community with integrated data from various administrative systems stored in one location. It is a read only database and therefore, information represented in it is maintained by other systems, called "systems of record". The main purpose of the Data Warehouse is to serve as a central reporting and data distribution environment for Departments, Labs and Centers at MIT. The Data Warehouse acts as a hub, to facilitate the exchange of information between systems and therefore serves as the enterprise information infrastructure at MIT.

Reports can be created from the Data Warehouse using any SQL query tool. The tool most often used at MIT is BrioQuery.The Data Warehouse website contains standard reporting templates which can be downloaded and run by DLC users. Users can also create their own adhoc reports as needed. Reports and/or query results can be exported into other software programs, such as Microsoft Excel. DLC users can also join local data with data from the Data Warehouse to produce custom report of interest to their particular DLC.

There are various advantages of using the warehouse as a reporting tool. The Data Warehouse tables are specifically designed for end user adhoc reporting. Since the Warehouse is a central repository for MIT Data, a user can obtain all the information he/she needs from various systems at one place instead of having to obtain it from each system separately.

The operational hours of the Data Ware house are as follows. The Data Warehouse is available seven days a week, except on Saturday afternoon from 6:00 -10:00 P.M. This time is used for weekly backups. The Normal User hours of the Warehouse are between 8:00 A.M. and 8:00 P.M daily. During these hours all data is stable and is not altered by any incoming feeds. Data is loaded daily into the Warehouse between 8:00 P.M. and 6:00 A.M. Although the Warehouse can be accessed during these house, the data maybe in an inconsistent state since it is being refreshed by the system of records. Data is extracted from the Warehouse between 6:00 A.M. and 8:00 A.M. This time is set aside to allow systems refresh their data extracts from the Data Warehouse without affecting end user response time. Users may also access the Warehouse during this time, but the response time may be slow.

Data Warehouse Logical Architecture Diagram

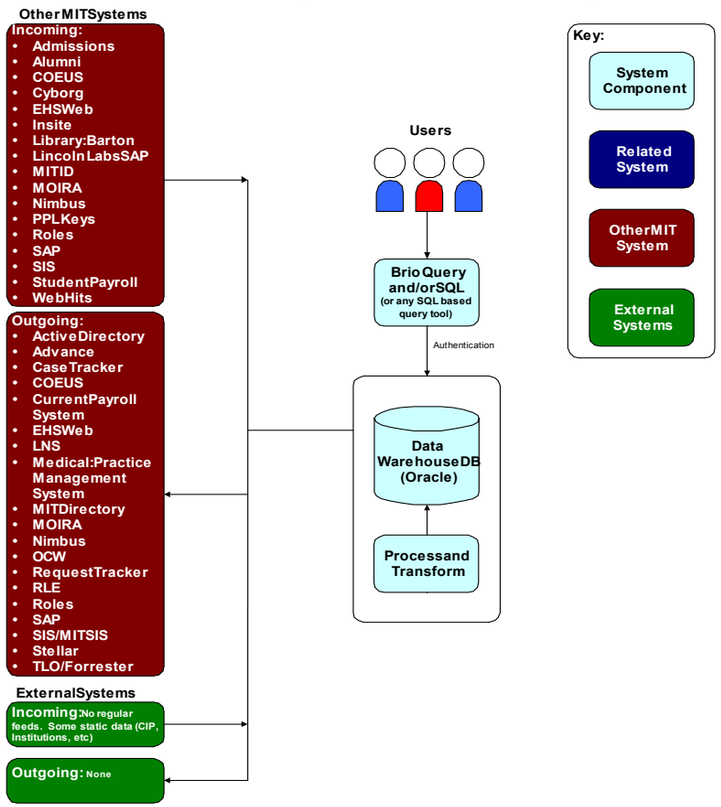

Data Warehouse Physical Architecture Diagram

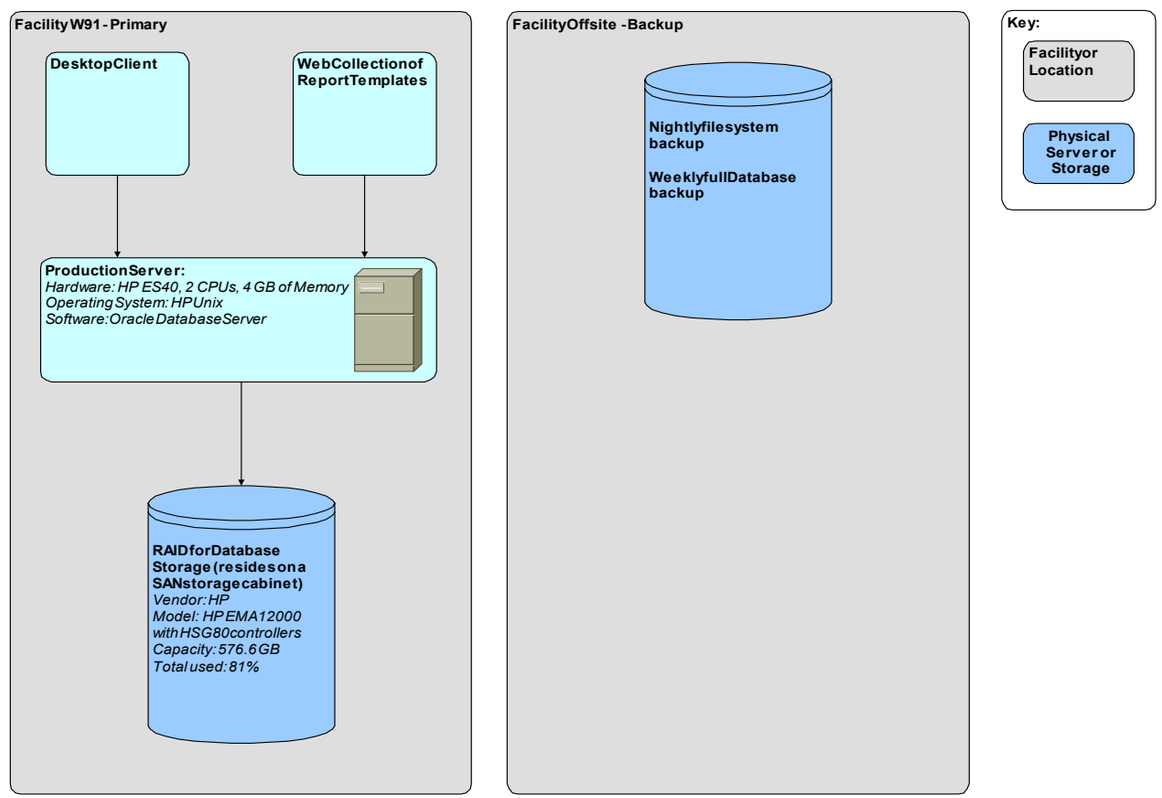

DSpace

DSpace was developed as a result of a joint project between MIT Libraries and Hewlett-Packard. Its purpose is to provide long term storage options of digital products to MIT faculty and Researchers. Users of DSpace can easily search, access, and read DSpace items through the Internet. While contributors can digitally distribute and preserve various formats of content including text, audio, video, images, datasets, and more. Digital work can be stored in various collections that are maintained by different MIT Communities.

At MIT, DSpace provides its users access to all research material at the institute through one interface. However, understanding that access restrictions are a need, policies can be customized to follow closely the need of any individual community at MIT. In addition the repository is organized in a manner which takes into account the various workflow issues which can be present at a multi-disciplinary institute such as MIT.

DSpace is an open source product and is available under the BSD open source license to other research institutions. These institutions may chose to run the project in its current form or modify it to their specific needs. It can be downloaded from SourceForge where users can access installation instructions to help new users install and run the application. (DSpace will run on any UNIX or LINUX operating system.)

As stated, DSpace is content is produced and managed by various communities at MIT. A DSpace community represents any academic unit at MIT which produces research. Each community should designate a coordinator to work with the DSpace staff. All individuals wishing to submit work through DSpace must belong to an existing community in DSpace.

Shown below is the logical and physical diagrams for DSpace. For more information on DSpace technology and the DSpace project, visit: http://www.dspace.org

DSpace Logical Architecture Diagram

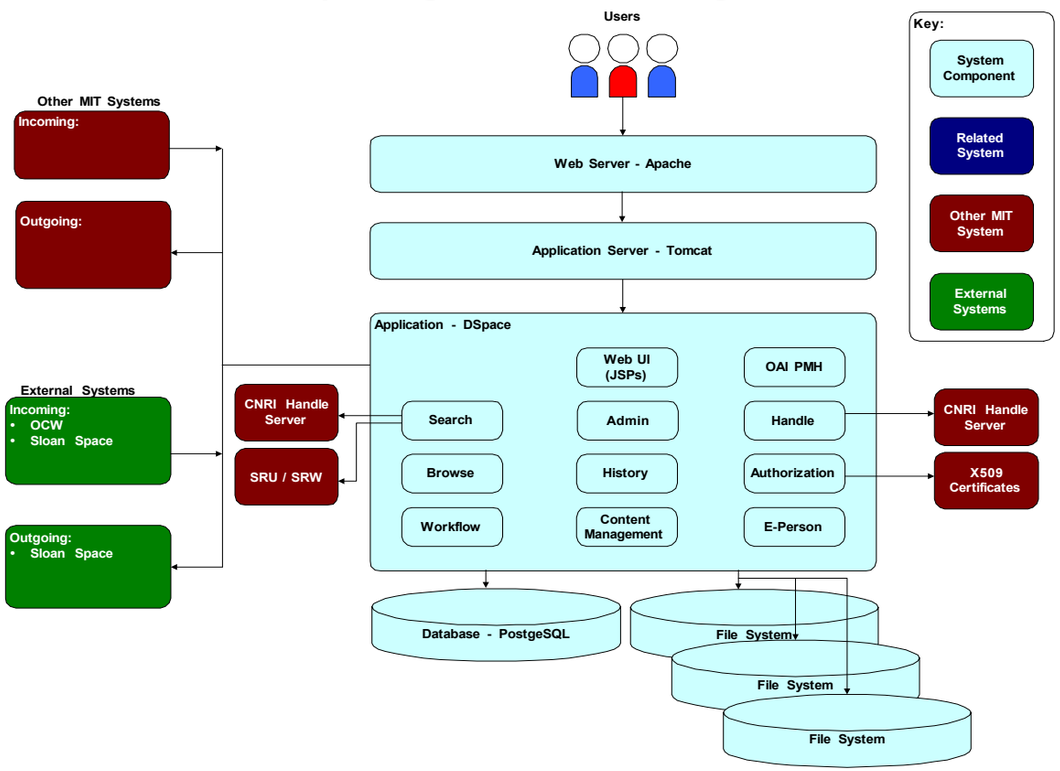

DSpace Physical Architecture Diagram

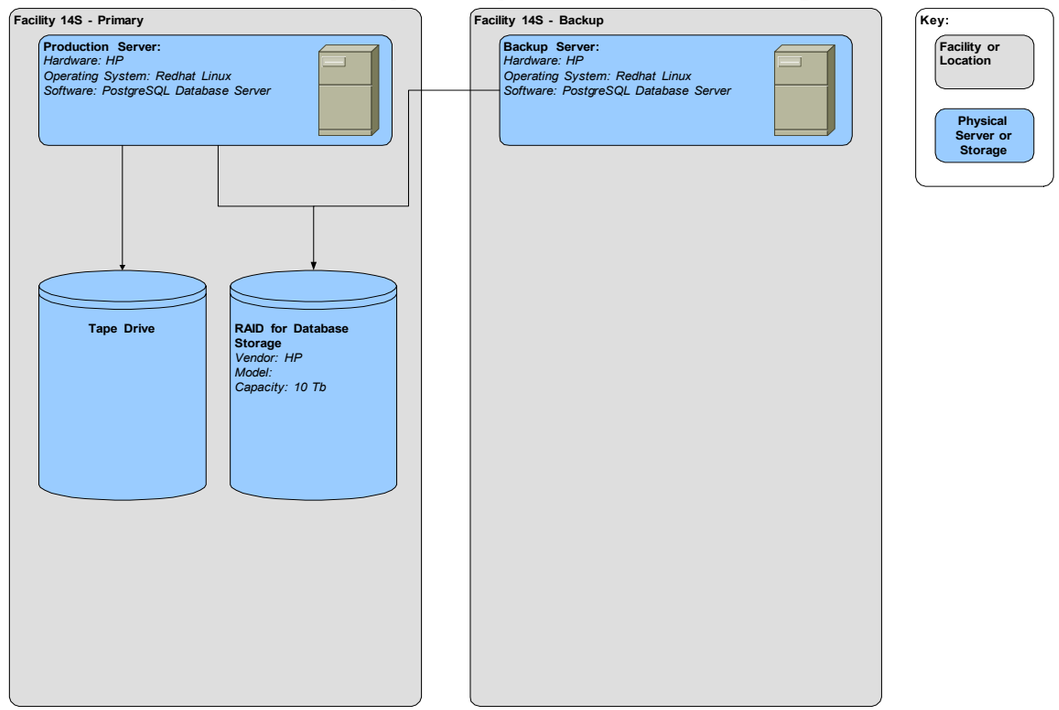

EHSWeb

EHSWeb is MIT's Environmental Health and Safety Management System. It provides a framework for DLCs to manage training and requirements for training for its employees. The EHS system was developed to help MIT with its policy towards safety in research and teaching. In addition EHSWeb helps provide MIT with accountability with environmental health and safety stewardship and still maintain an independence in research and teaching.

The EHSWeb system has various user interfaces to allow users to perform the actions outlined above. Each user group of the application can use the appropriate interface to conduct their daily business. These interfaces are shown below, indicative of their purpose.

The Diagrams below display the Logical and Physical Architecture of EHSWeb. The Logical Architecture diagram displays systems in addition to the EHSWeb application. These systems include SAP, PI/Space Registration, Sara and Netcaster. An explanation of these systems follows. SAP contains information about the training events MIT community members can attend and those they have already attended. Therefore this information must be fed into EHSWeb to gain a holistic understand of the training an individual has taken. The PI/Space Registration system allows Principle Investigators of various areas at MIT to keep track of the hazards in their area of responsibility. This information allows the EHSWeb application to gain an understanding of the training needs people may have due to the spaces they are currently working in. Finally, Netcaster is a web based application to allow MIT community members to take online training courses. It manages course information and completion and must feed this to EHSWeb to indicate people are completing their required training courses.

Currently there are plans to move the EHSWeb functionality into SAP. By February 2005 the inspection and audit phases are expected to go live, followed by the training module in June 2005.

EHSWeb Logical Architecture Diagram

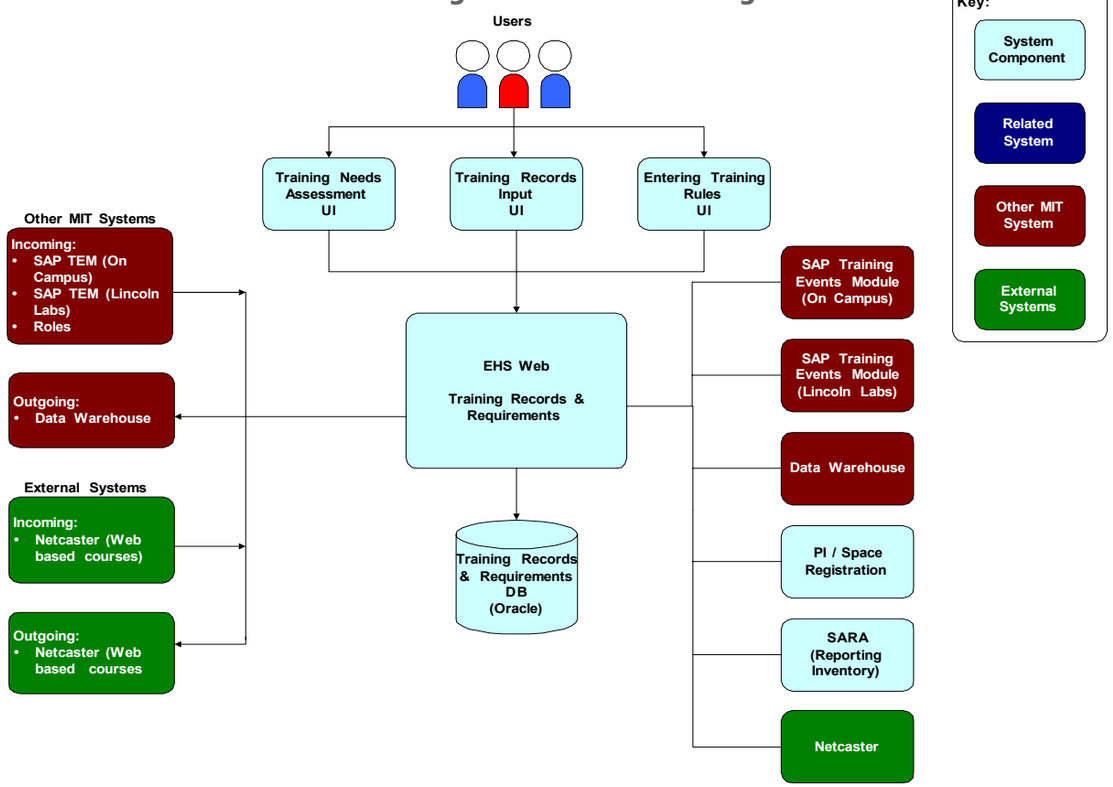

EHSWeb Physical Architecture Diagram

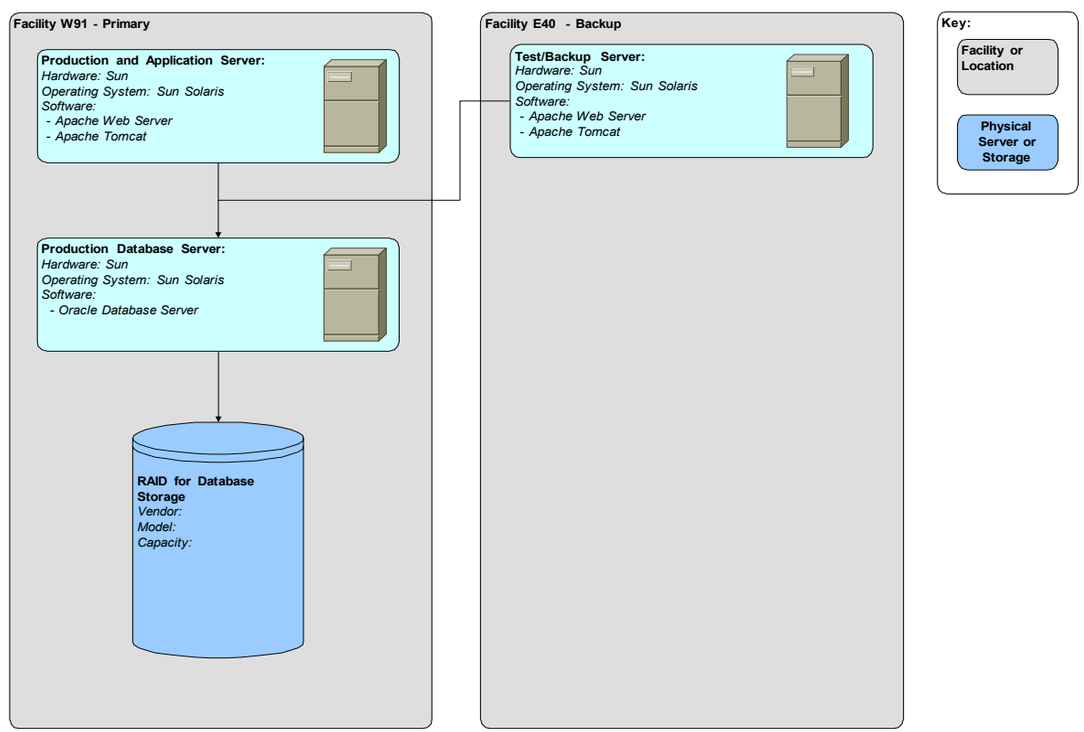

Medical Center Environment

The Medical Center at MIT is a complete Medical Center, combining both in-patient and out-patient facilities, based on the MIT campus and operated by MIT. The Center is similarly both a provider of medical services, and an insurer or payor, providing medical insurance for students and employees of MIT. As a result their IT environment resembles that of any large medical practice. The fact that the Medical Center is a part of MIT is reflected only in the fact that it's patients are affiliated in some way to MIT and that some patient and financial data flows between systems within the medical center and other systems at MIT.

The first key system at the Medical Center is the Practice Management System. This system is responsible for many of the operational aspects of the Medical Center. It contains patient data, manages work done under the MIT HMO organization, manages registration and enrollment information and the student health program. The Practice Management System is an implementation of a system provided by IDX, with several custom modules written to extend the functionality to match MIT requirements.

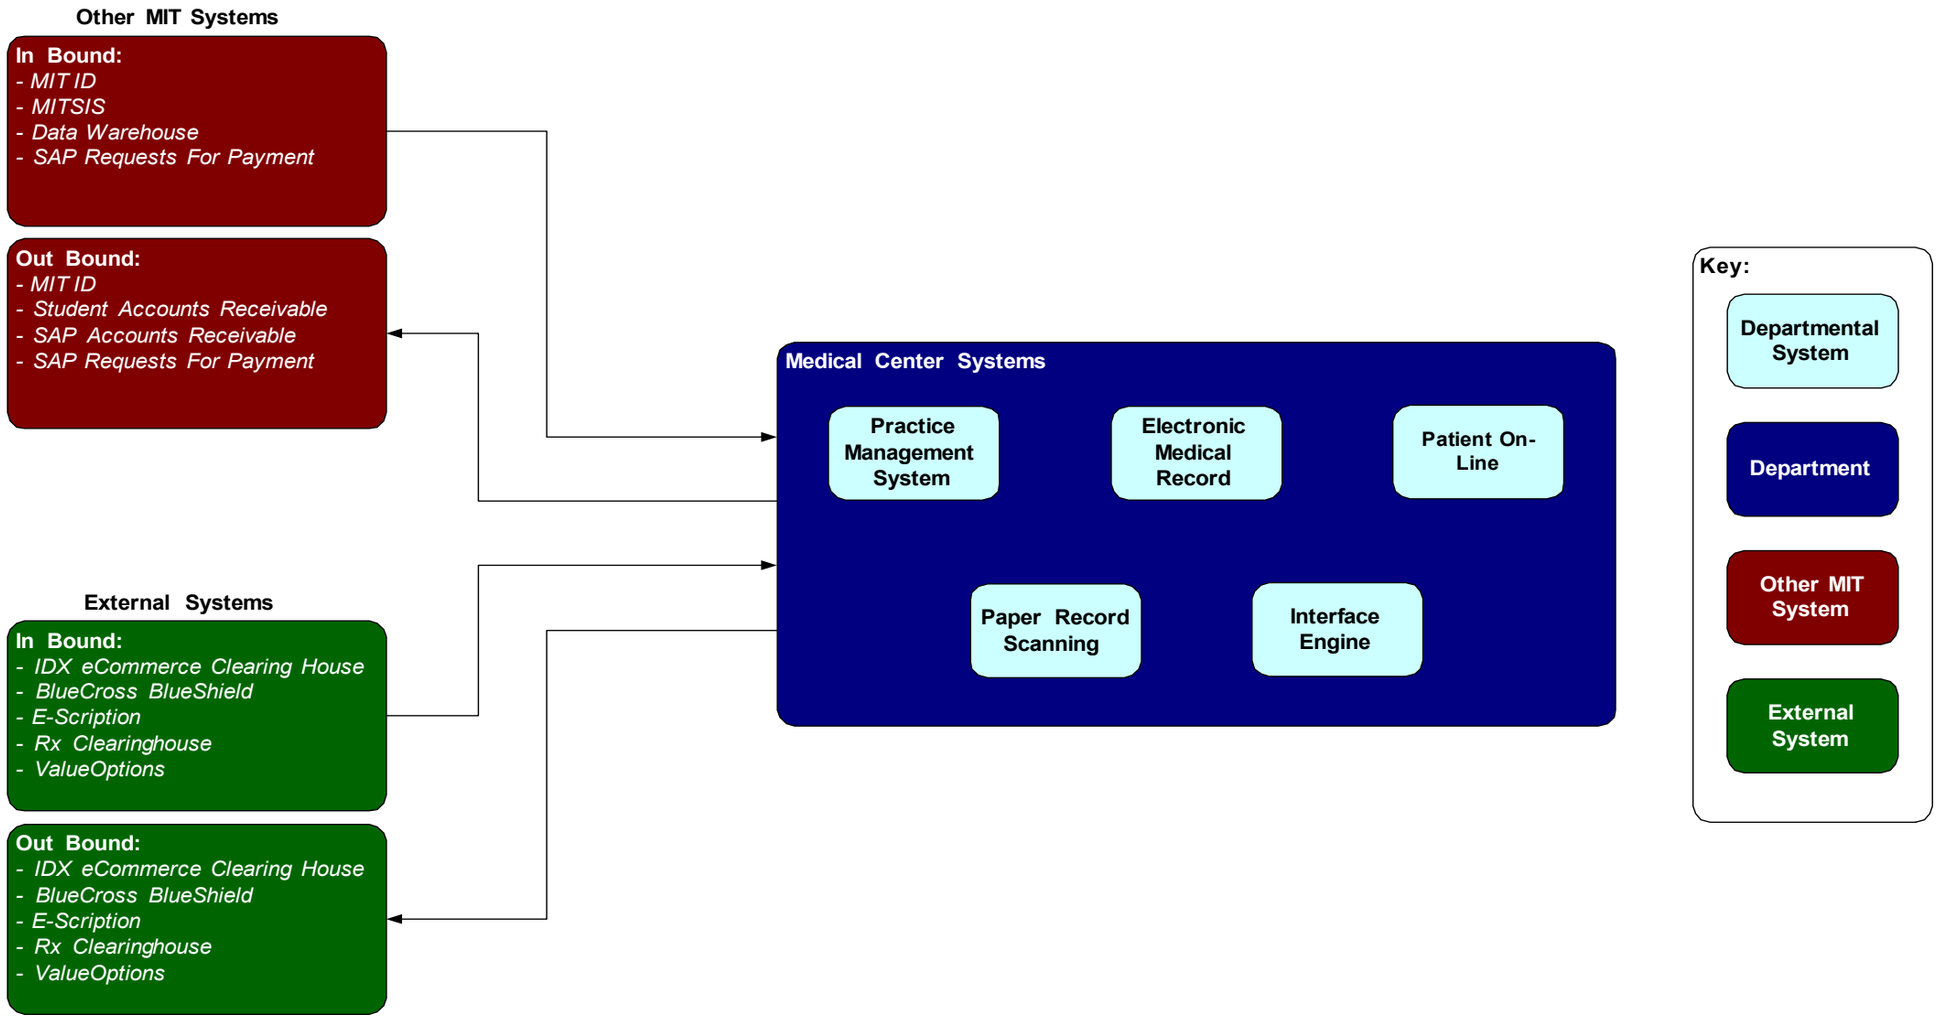

The Electronic Medical Record system is an implementation of Alllscripts' Touchworks product and is the central repository for all information pertaining to a patients medical record (visits, treatments, prescriptions etc.).

The Paper Record Scanning system is a system designed to digitally scan paper based medical records accumulated by the Medical Center and enter them in to the Electronic Medical Record System.

Patient On-Line is an online patient portal that is, at time of writing, limited in functionality. The vision for Patient On-Line is that it would become a significant channel for patients to view their medical information and interact with the Medical Center in a self-service manner that is also HIPPA compliant.

The Interface Engine is responsible for managing all interfaces to the TouchWorks electronic medical record. These include demographic and appointment data, prescription drug and diagnostic test data from Medical's custom systems and scanned documents from the Scan system. Of particular note among the external systems is E-Scription. This is a system maintained by an Application Service Provider for the purpose of transcribing medical information captured verbally by doctors into text which can be stored within an electronic patient records.

OpenCourseWare (OCW)

MIT OpenCourseWare (MIT OCW) makes the course materials that are used in the teaching of almost all MIT's undergraduate and graduate subjects available on the Web, free of charge, to any user anywhere in the world. As of May 2004 over 900 courses have been published on MIT OCW. A state goal is to publish virtually all 2000 of MIT's undergraduate and graduate courses on MIT OCW by 2008.

More broadly MIT OCW's goals are:

- Provide free, searchable, access to MIT's course materials for educators, students, and self-learners around the world
- Create an efficient, standards-based model that other institutions may emulate to openly share and publish their own course materials

The MIT OpenCourseWare (MIT OCW) technology solution supports a complex publishing process. This is a large-scale digital publishing infrastructure that consists of planning tools, a content management system (CMS), and the MIT OCW content distribution infrastructure. The current technical solution has been in use since April 2003 with a four-person technical support team managing all aspects of this infrastructure.

The planning tools used by the MIT OCW team to assist faculty in publishing their course materials include a custom application of FileMaker Pro, and several checklists and documents. For creating and managing content, we use several desktop tools (file conversion tools) as well as the CMS, an extensively customized version of Microsoft Content Management Server 2002 that fully supports our publishing process. Our content delivery infrastructure includes a sophisticated publishing engine, content staging server, and a content delivery network utilizing Akamai's EdgeSuite platform.

OCW Logical Architecture Diagram

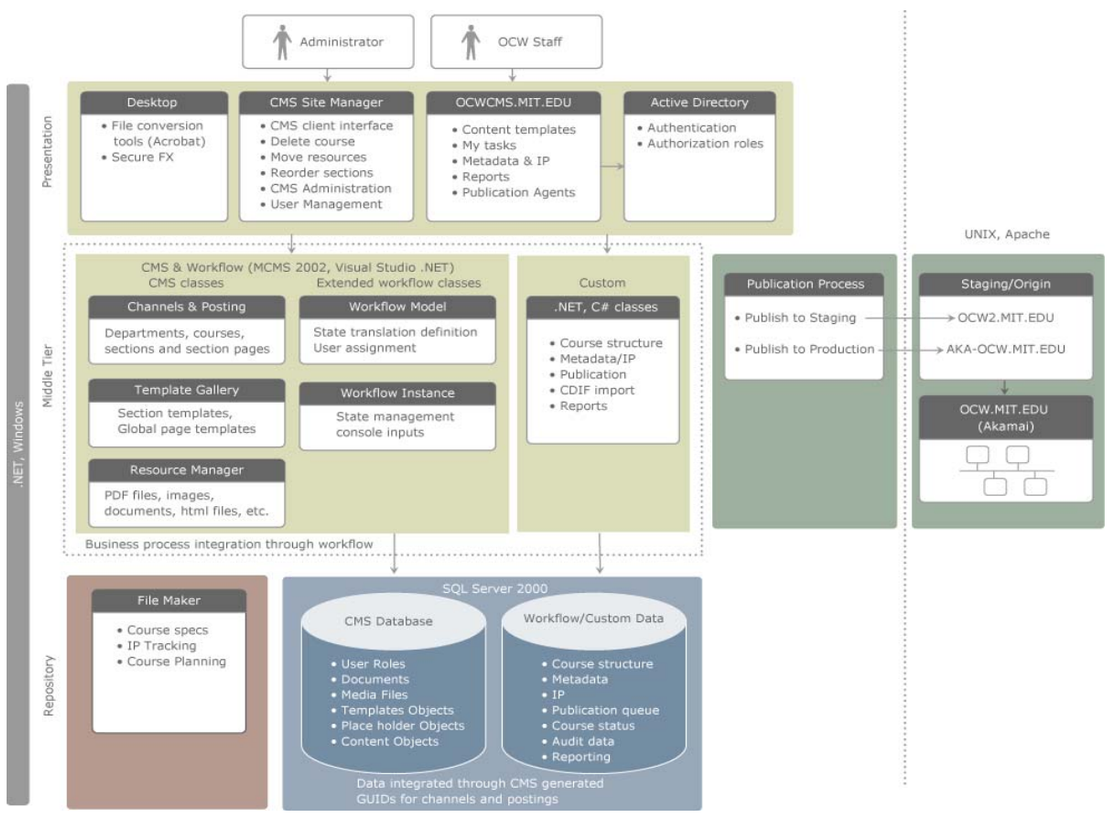

OCW Physical Architecture Diagram

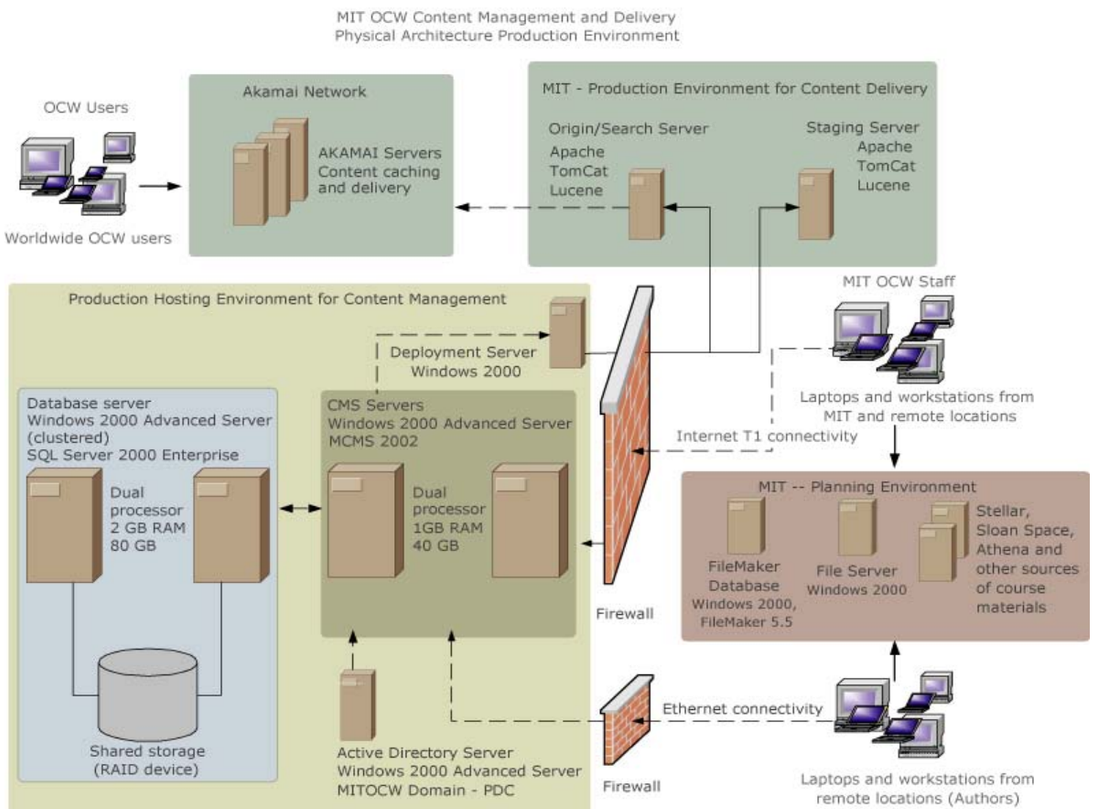

SAP

In the mid 1990's MIT chose SAP to be its ERP Solution. Initially the Financial Piece of SAP was implemented at MIT followed by an Administrative implementation. MIT has two separate installations of SAP. One installation is for its main campus, the other for Lincoln Laboratory. (Lincoln Laboratory is a Federally Funded Research and Development Center of MIT located in Lexington, MA). The main campus installation of SAP has the following modules currently implemented: FI: Financial Accounting, CO: Controlling, PCA: Profit Center Accounting, MM: Materials Management, SD: Sales and Distribution, CA-CL: Classification System, HR: Human Resources, PY: Payroll, LDS: Labor Distribution, CA: Cross Application, Plant Maintenance, and Training Events Management. Lincoln Lab has its own support staff for their SAP installation and maintains it separately from the Main Campus installation. Customized I-Doc interfaces from SAP at Lincoln Lab are used to communicate to the main SAP system. These interfaces help integrated the two SAP systems to fulfill the financial reporting relationship between the two. In the near future these interfaces will also help integrate the Hazardous Materials/Health and Safety reporting relationship between the two.

There are two different ways to access the SAP applications at MIT. Users can install the SAP GUI on their desktop and access it from there. In addition MIT has a web interface called SAPWeb which users can use for SAP requisitioning. Later this year there are plans to integrate a Web Application Server for SAPWeb.

Various updates and upgrades are planned for SAP in the next couple of years. A highlight of these include the following. SAP-BUD will signify the migration of the Nimbus (budgeting) functionality to SAP. This is scheduled to go live in December 2004 with a web front end interface. EHS is scheduled to have three phases rolled out. Inspection and audit are scheduled to go live in February 2005. The training and needs assessment phase will go live in Jun 2005. Employee Payroll will be part of SAP by Jan 2006. (Student payroll will still be managed by SSIT.) Finally a hardware renewal of the SAP platform is scheduled for 2005. There may also be an upgrade of the SAP version around the same time.

Shown below are the logical and physical architecture diagrams of SAP.

SAP Logical Architecture Diagram

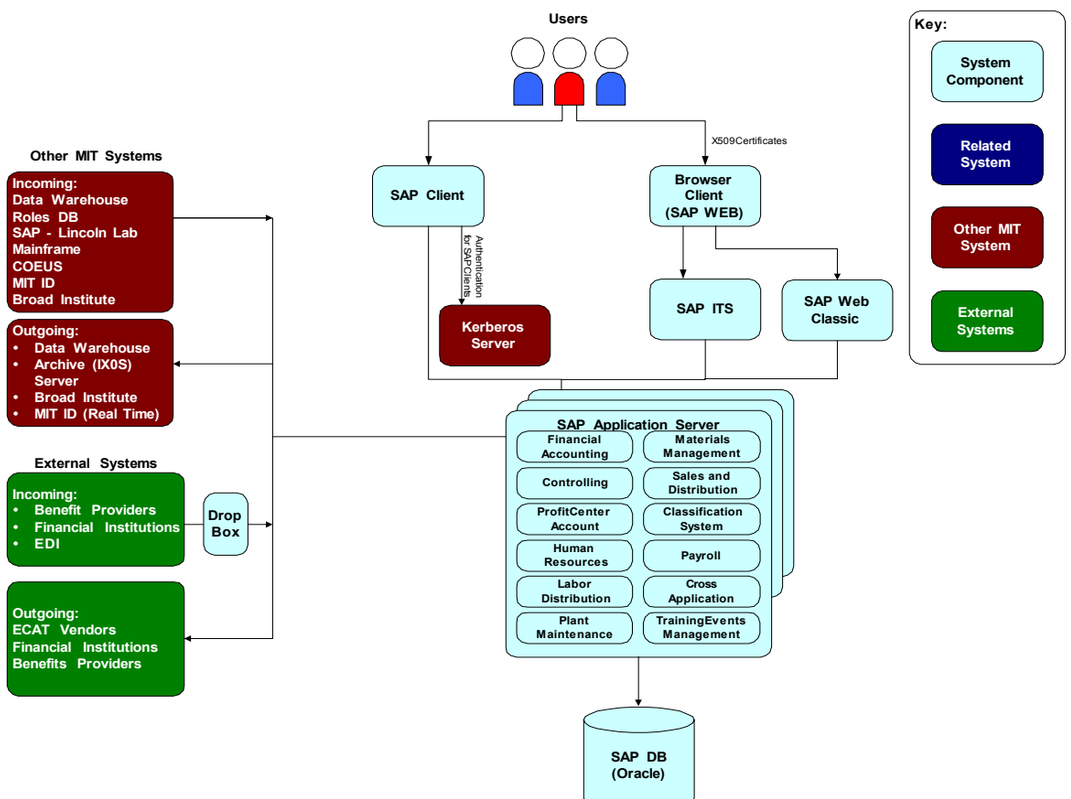

SAP Physical Architecture Diagram

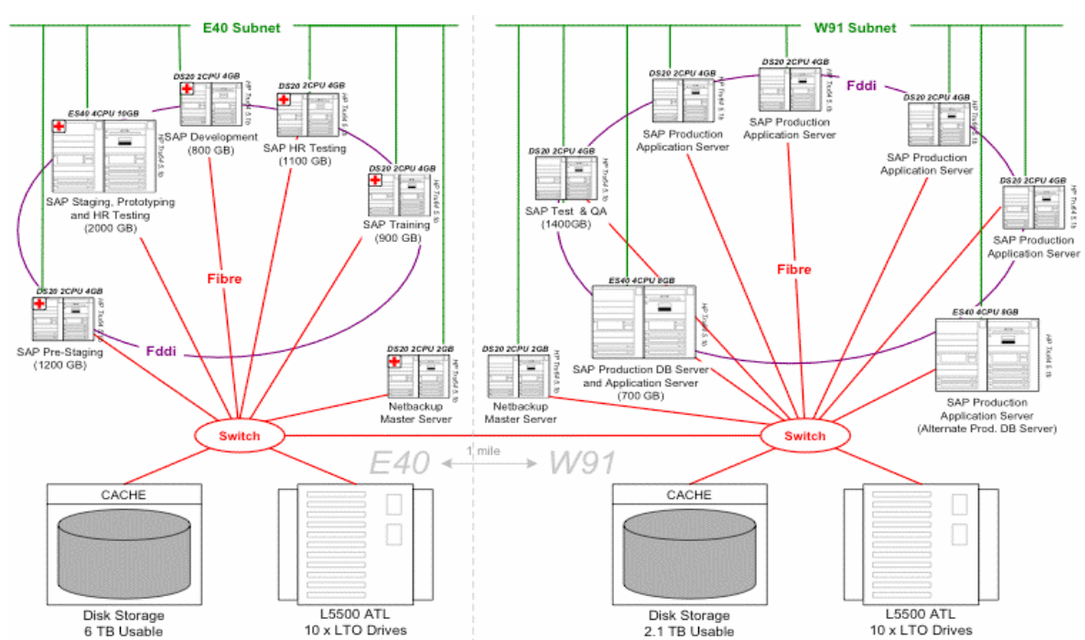

SloanSpace

SloanSpace is a portal for the entire extended MIT Sloan community. Its purpose is to allows its users course management facilities, provide research material/collaboration, and create online communities. Its users vary across a wide spectrum and include groups such as industry sponsors, researchers, alumni, students, and staff.

Version 2 of SloanSpace is based on the dotLRN project. dotLRN is an MIT open source platform which is currently being used by 25 institutions and organizations worldwide. All new rollouts made to dotLRN are first tested in production by one member of the community and then deployed. This ensures greater quality and allows its users to have confidence in the product.

SloanSpace has been in use at MIT since Fall 2000. Since then it has become an integral part of the Sloan School of Management. Currently there are over 5000 users of the application and about 100 courses at Sloan utilize SloanSpace on a daily basis. This represents over 90% of the total courses Sloan offerings. Individual faculty in other departments such as physics, aeronautics and astronautics, and mechanical engineering have also used the application to manage their courses. In addition to these courses, SloanSpace also has more than a hundred online communities. These communities range from student groups, research centers, and industry partners interested in collaborating ideas and work.

For all its users SloanSpace proves to be an indispensable too. Professors can use it to easily distribute teaching material to their students. Students can use it to easily access all course material they are currently enrolled in and submit homework assignments electronically. Researchers find great ability in being able to exchange ideas and store work for greater collaboration and quicker response times.

The diagrams below show the Logical and Physical Architecture of SloanSpace.

SloanSpace Logical Architecture Diagram

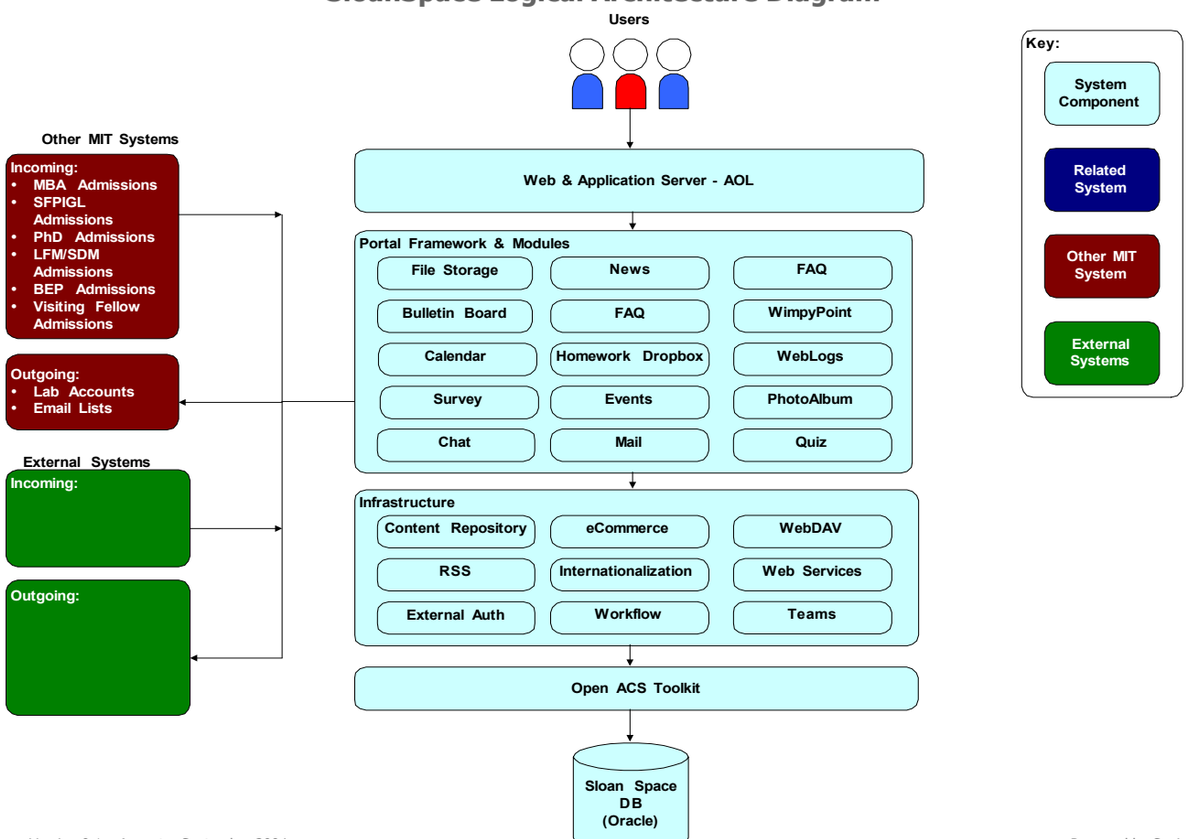

SloanSpace Physical Architecture Diagram

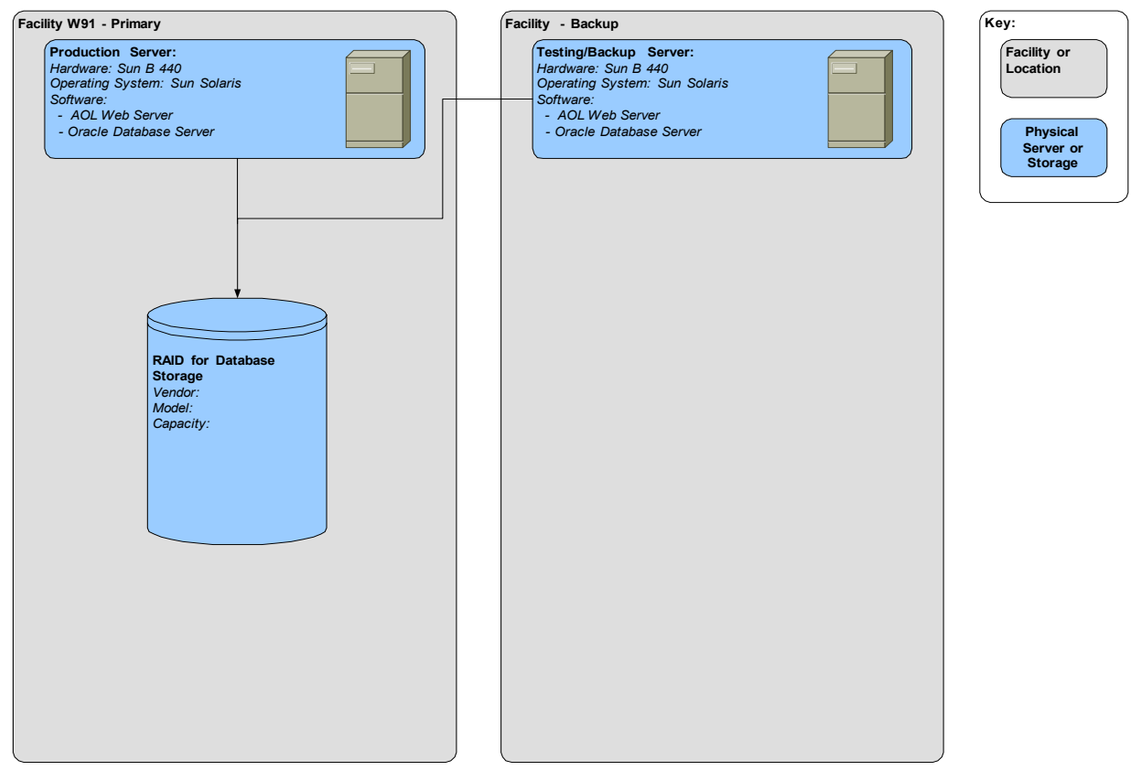

Stellar

Stellar is a Learning Management System developed at MIT for managing the authoring and delivery of courses. It provides faculty with the ability to:

- Add and manage a wide variety of multimedia class materials without technical knowledge
- Manage class website membership automatically with official data from the Registrar, and become acquainted with students through photos from the Card Office
- Simplify administrative tasks by managing homework submissions and letting students choose sections on the class website
- Communicate easily with students through announcements, email, and threaded discussions

Stellar is available for use by all departments and faculty at MIT. For the Fall 2004 semester there are 293 class websites, which represents approximately 40% of all classes taught at MIT.

Stellar is a java-based application, primarily a web application delivered through Tomcat and the Apache web server. There is a abstraction for accessing the database and also for accessing the file system where course materials are stored. On top of these abstractions is built a layer of domain objects. The domain objects contain reusable business logic and are accessed both by tools that are part of the online web application, and also by stand alone java programs which process feeds and other batch tasks.

Sakai is open source collaboration between several leading educational institutions, including MIT, to develop the next generation of educational software. Their goals include developing a common, open system for learning management. It is expected that over the next 12-24 months, parts of the Sakai toolset such as the grade book will be introduced at MIT. In the long term it is likely that functionality currently provided by Stellar will be superceded by similar functionality provided by Sakai, and that the Stellar application will gradually be replaced by an implementation of Sakai at MIT. More information on the Sakai project can be found at: http://www.sakaiproject.org

Stellar Logical Architecture Diagram

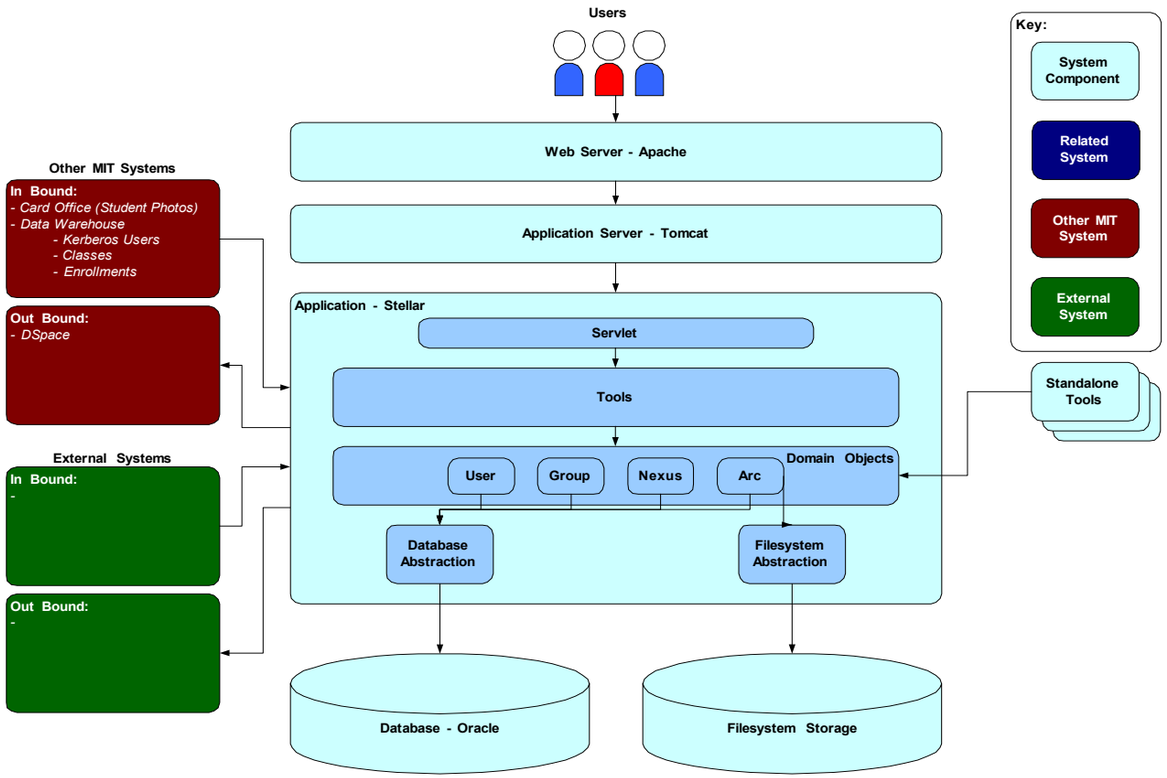

Stellar Physical Architecture Diagram

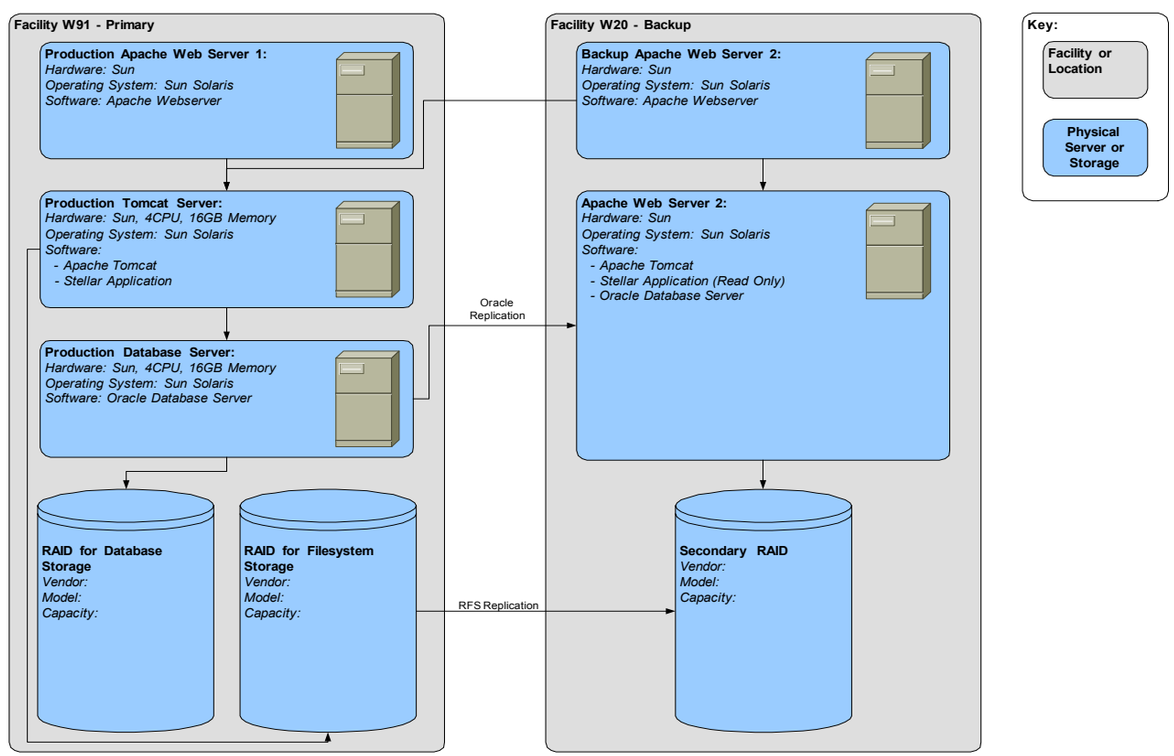

Student Systems Environment

The Student Systems at MIT, those maintained by Student Services Information Technology (SSIT), are responsible for providing a wide variety of services. The broad areas covered include: undergraduate and graduate admissions, registration, all student financial and accounting, grade processing, enrollment and the provision of basic services to students such as dining and housing.

The systems are built on four major platforms, supported by three shared databases. The legacy Mainframe Admissions Database contains admissions data for all students at MIT. The newer Admissions Database, using Oracle, is used by the Undergraduate Admissions system, and synchronizes admissions data with the mainframe admissions database. Lastly, the SIS database is the repository for the vast majority of student information.

The MITSIS platform is the established legacy platform for developing applications within SSIT. It is based on a highly customized version of a purchased package called Banner, which is implemented in SQL*Forms 3.0 and an extensive set of batch jobs written primarily in Brio SQR and Pro*C. The MITSIS database is supported by Oracle 7.3.4 – the last version compatible with Oracle SQL*Forms 3.0 an SQL*Net 2 upon which the Banner package depends. The MITSIS platform supports a wide variety of applications for use within student services. In addition there are several vital PowerBuilder applications that, while not part of the MITSIS platform, are contemporary with it.

The WebSIS platform is an outgrowth of the MITSIS platform created for developing web-enabled applications. It is a mix of technologies mostly using CGI technology to present a web interface to applications. The WebSIS platform primarily consists of Brio SQR used as both a transactional and reporting CGI technology. Brio SQR contains embedded SQL and also calls upon an extensive library of business logic implemented as Oracle stored procedures.

More recently SSIT have been using J2EE technologies to implement applications. Commonly referred to as the J2EE Non-Portal Platform, is an aggregation of applications that have been built, for varying purposes and at varying times, on the J2EE platform, but remain separate applications.

The final platform, and most recent, is the J2EE portal platform. This platform leverages BEA WebLogic Server and WebLogic Portal Server to create unified architecture and presentation strategy for new applications. The first application to be developed under the J2EE portal platform is the Undergraduate Admissions Portal. It is expected that the technological and physical infrastructure put in to place for the Undergraduate Admissions Portal can be used to support future applications developed within SSIT. The J2EE portal platform was explicitly built with the twin goals of supporting the Admissions Portal and constructing a Services Oriented Architecture (SOA) migration target for all existing SSIT platforms.

Student Systems Logical Architecture Diagram

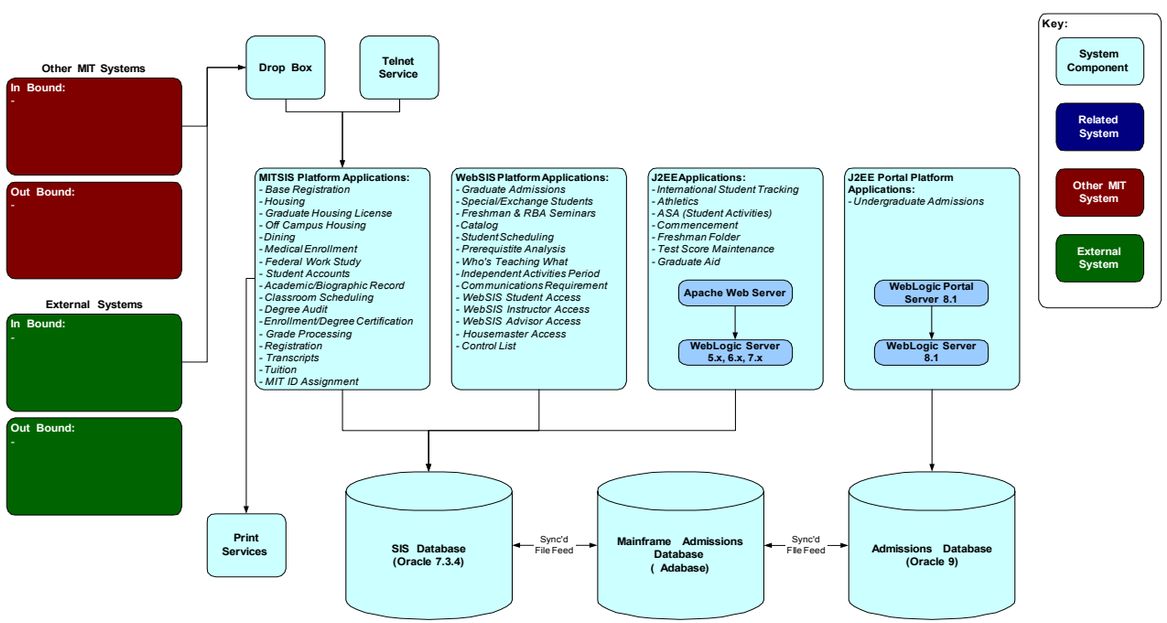

Student Systems Physical Architecture Diagram

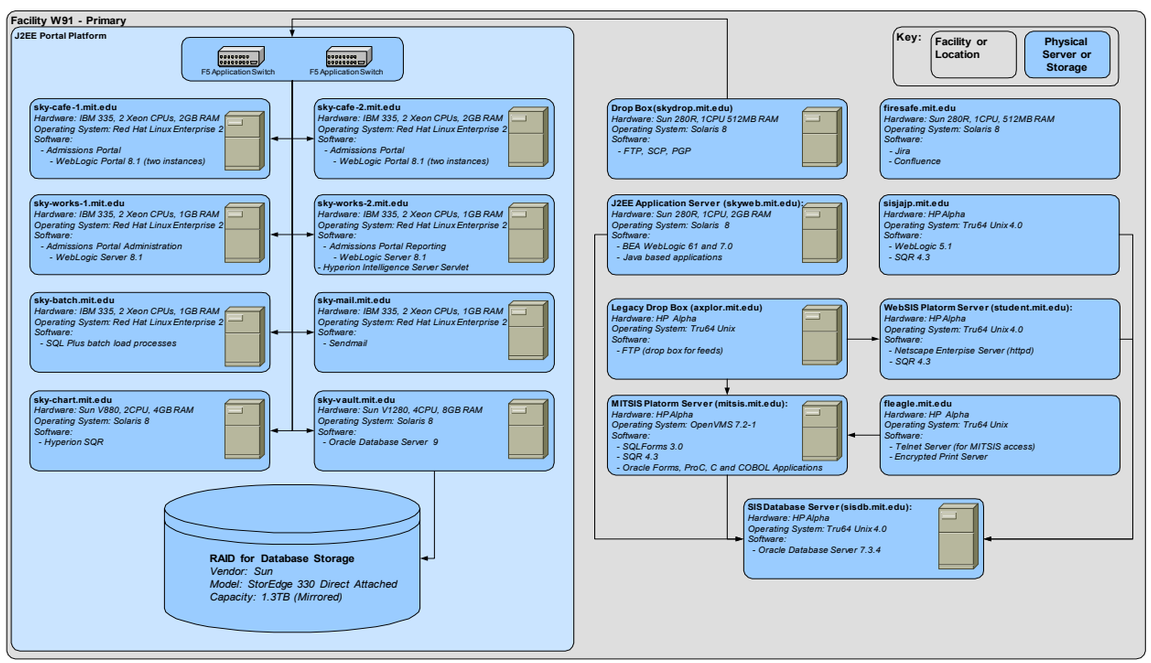

Instructions for SSIT Applications Inventory

Description	The SSIT Applications Inventory is the consolidated list of applications which SSIT support
Version	0.1
Worksheet Definitions	
SSIT Applications Inventory	The SSIT Applications Inventory; see above
Column Definitions	
Application	The name of the application which SSIT support.
Purpose	The purpose of the application listed in the application column.
Sponsor	The department which sponsored the application listed in the application column.
Users	The user group/ user department who primarily uses the application listed in the application column.
Technology	A listing of the technology involved in the application listed in the application column.

SSIT Applications Inventory

Application	*Purpose*	*Sponsor*	*Users*	*Technology*
Undergraduate Admissions	Web-based portal for prospects - personalization capabilities, registration for portal, events, online application, application tracking internal facing allows staff to find prospects, update data, search for subsets of applicants, target them in a variety of media, track contacts, Recruitment;tracking;campaign management; prospect information management;educational counselor tracking; Adabase/MITSIS interface	Admissions	Admissions office	J2EE; Cobol; Adabase, Natural
Graduate Admissions	for all graduate applicants: contact management, application requests and downloads, applications, applicant biographic record and addresses, educational history, Departments admissions decisions and applicants' replies to MIT admission offer, decision letters, admission reply survey displaying, reports and statistics, mailing house pass	Admissions	Graduate Admissions office, Provosts Office, Graduate Administrators (33 depts.), ISO, Office of Minorities	PowerBuilder, Web CGI, PL/SQL, UNIX
On-line graduate admission applications and recommendations	electronic form for application and recommendations. Special setup for course 6 on-line applications	Admissions	all graduate level applicants and their professors for recommendation letters, Graduate Admissions office, Graduate departments admins and faculty, Course 6 IT	CollegeNet partnership, SQR, FTP, pdf, mass printing
MIT ID Assignment	assigning MIT IDs to all graduate applicants	Admissions	Graduate Admissions Office, Registrar's Office	PowerBuilder,PL/SQL
Test scores (GRE and TOEFL) maintenance	for graduate applicants	Admissions	Graduate Admissions office, Graduate departments' administrators	IBM mainframe REXX programs, PowerBuilder, PL/SQL
Graduate Admissions data download to all graduate departments' databases	database template and download maintenance for all FileMaker Pro departmental databases. Specialized data feeds for Course 6 and Sloan	Admissions	Graduate Administrators (33 depts.)	FileMakerPro, Oracle db connection, PL/SQL
Graduate admissions mass emailing	for application requests, also for surveys for admitted-but-not-coming applicants	Admissions	Graduate Admission office, Admissions contacts, MIT Admitted graduate applicants	UNIX scripts and emailing, SQR
Base Registration and Financial Aid record creation	for all admitted graduate applicants	Admissions, Registration, SFS	Registrar's Office, SFS	PL/SQL, UNIX
Special, Exchange and Visiting Students admissions	for undergraduate and graduate level	Admissions	Graduate Admissions office, Provosts Office, Graduate Administrators (33 depts.), ISO	PowerBuilder, Web CGI, PL/SQL

SSIT Applications Inventory

Application	*Purpose*	*Sponsor*	*Users*	*Technology*
EC Tracking	Secure site navigated from the Alumni login site, allows Educational Counselors to view personal data as well as their assigned applicants and the application status information. Ability to submit interview online and upload to VMC.	Admissions	Admissions office	J2EE
ISO	Issue F1 & J1 documents and updates for students and dependents; Homeland Security reporting; travel letters for students and dependents; ssn letters; orientations and holds processing ;student data views; student employment authorization	International Students Office	International Student's Office	Java/Swing
Athletics	Physical Education class registration and requirement tracking;NCAA and MIT athlete eligiblity;PE class lottery entry	Athletics	Athletics, students, MIT faculty and staff	Java - Swing; Clipper; FileMaker;CGI
Housing (grad and undergrad)	Fee assessment; building & room definitions; room assignments; housing taxes, charges and refunds; reports, housing roll	Housing	Housing office, DSL financial administration	Oracle Forms; C; SQR
Graduate Housing license printing	Management of contracts for graduate housing tenants	Housing	Housing Office	PowerBuilder;PL/SQL
Off-campus housing	web-based interface to list off-campus rental and roommate offerings	Housing	rental owners, roommate seekers, MIT community	CGI; SQR
Dining	Dining program definition; fee assessment, Dining systems interface to MITSIS	Dining	Dining office, DSL financial administration	Oracle Forms; SQR
Medical	Medical program rules definition; insurance coverage maintenance for students and dependents, fee assessment; Blue Cross/Blue Shield Interface, On-line insurance waiver interface	Medical	Medical office	Oracle Forms; SQR
ASA (Association of Student Activities)	definition and management of student activity groups at MIT; request submission and processing for lockers, offices, mailboxes, bulleting boards, midway participation and early returning students; student activity groups' financial and facilities reservation signatories viewing and updates; president and treasurer lists and maintenance, display and management of amenities assigned to a student group	Student Activities Office	Association of Student Activities executive board, Student activity groups, Student Activities office, Campus Activities Complex, Scheduling Office, Audio/Visual Services, Athletics Facilities Scheduling Office, world-wide internet users, MIT community	J2EE, web-based
PowerFaids	Financial aid need analysis; financial aid disbursement	SFS	Student Financial Services	SQL Server
Federal Work Study		SFS	Student Financial Services	PowerBuilder
Student Accounts	Cashiering session;financial holds;cash advances;refund processing;statement processing; application of payments	SFS	Student Financial Services; students	Oracle Forms; C; J2EE

SSIT Applications Inventory

Application	*Purpose*	*Sponsor*	*Users*	*Technology*
Freshman and RBA Seminars	Subject proposal/approval work flow; student requests	Academic Services	Incoming freshman	CGI; SQR
Academic/Biographic Record	Maintenance of student academic information; enrollment status; biographic information; reports	Registrar	Registrar's office, academic departments, advisors	Oracle Forms; SQR
Catalog	Subject proposal/approval work flow; Online subject listing	Registrar	Academic departments, Registrar's office, Committee on Curricula, Students, World	CGI; SQR
Student Scheduling	Student subject scheduling	Registrar	Registrar's office, students	CGI; SQR;C
Classroom Scheduling	Classroom scheduling	Registrar	Registrar's office	C
Degree Audit	Student degree audit maintenance and calculation	Registrar	Registrar's office, students	Oracle Forms; C; SQR
Enrollment/Degree Certification	Certification for student attendance and degrees	Registrar	Registrar's office, students	Oracle Forms; C; SQR
Commencement	degree application, degree tracking and approvals, commencement book, diplomas;	Registrar	Registrar's office, academic departments, students	Web - J2EE
Final Exam Scheduling	Scheduling exam scheduling	Registrar	Registrar's office, students	Package - Strathmann Associates Final Exam Scheduler
Grade Processing	Grade rule maintenance, grade entry, grade tracking, grade reporting	Registrar	Registrar's office	Oracle Forms; C; SQR
Registration	Pre-registration; HASS-D lottery entry;class lists, picture class lists, registration form, registration holds, registration data entry, cross registration; subject add/drops; status of registration	Registrar	Registrar's office;students; academic departments; advisors; faculty	Oracle Forms; SQR; C; J2EE
Transcripts	Student transcript processing	Registrar	Registrar's office, Student Service Center, Students	Oracle Forms; SQR
Tuition	Tuition rules maintenance; fee assessment	Registrar	Registrar's office	Oracle Forms; C; SQR
Prerequisite analysis	Prerequisite rules maintenance; prerequisite deficiency reporting	Registrar	Registrar's Office, Academic Departments; Faculty	Web - CGI; SQR
Freshman Folder	Advisor and freshman online access to test scores, results, credit, and advisory messages	Academic Services	Academic Service Office, students, advisors	Web - J2EE
Who's Teaching What	Online instructor maintenance by term for enrollable subjects.	Academic Services	Academic Service Office, students, advisors	Web - CGI, SQR
IAP	IAP subject proposal entry and approval work flow; Online IAP guide and calendar	Academic Services	Academic Service Office, IAP coordinators, students, faculty, staff	Web - CGI, SQR
Testscore Maintenance	Freshman testscore maintenance and credit rules	Academic Services	Academic Services Office, Registrar's Office	Web - J2EE
Communications Requirement - Advisory Messaging	Rules definition; population selection based on requirement rules; email advisories; tracking	DUE	Communications Requirement	Web - CGI; SQR, J2EE
Graduate Aid	Graduate award maintenance	Provost	Academic Departments, Provost's Office	Web - J2EE

SSIT Applications Inventory

Application	*Purpose*	*Sponsor*	*Users*	*Technology*
WebSIS - Student Access	Pre-registration; HASS-D lottery entry; student schedule; grade report; status of registration; degree audit; financial aid statement; financial aid requirements tracking; student account activity; student account statement; address & biographic information maintenance; degree application; PE lottery entry; student email help	Registrar/ SFS	Students	Web - CGI, SQR, J2EE
WebSIS - Instructor Access/Department Administrator	Web view to: pre-registration picture class lists, pre-registration class lists & downloads, registration picture class lists, registration class lists & downloads, pre-requisite compliance reporting	Registrar	Faculty, Department Administrators	Web - CGI, SQR
WebSIS - Advisor Access/Department Administrators	Web view to student: pre-registration; HASS-D lottery entry; student schedule; grade report; status of registration; degree audit; picture; address; data downloads; department picture lists	Registrar, Academic Services	Advisors, Department Administrators	Web - CGI, SQR
Housemaster Access	Web view to student: pre-registration; HASS-D lottery entry; student schedule; grade report; status of registration; degree audit; picture; address; data downloads	Registrar	Housemasters	Web - CGI, SQR
Control List	Provides department administrators access to student information as reports or downloads to support monitoring of registration activities, and to feed department data bases with accurate student information.	Registrar	Registrar's office, Academic Departments	J2EE