



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 to 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.

