

Draft: Integration Coordination

A manual for Departmental Software Platforms and Services (DSPS)
Infrastructure Software Development and Architecture (ISDA), IS&T, MIT

Document Summary

This is a manual. It describes policies and recommendations regarding the management of application platforms by the Departmental Software Platforms and Services team (DSPS). This will include the management of all assets required to reproduce and run server applications, such as source code, configuration, and procedures for release to production support.

This document is used to instruct work and for audit guidelines.

Revision History

08/25/08	Steve Landry	Source-Control policies
08/29/08	Steve Landry	Additions about managing access-control lists
09/07/08	Steve Landry	Changed installation procedures. Added Pipeline process information.
10/06/08	Steve Landry	Added information on support requests, email and ACL, RT queues.

Table of Contents

- Draft: Integration Coordination..... 1
 - Document Summary..... 1
 - Revision History..... 1
- Communication, Email, and Ticket Queues..... 3
 - Requesting Support..... 3
 - Email and Access-Control Lists..... 3
 - Request-Tracker Queues..... 3
- Process and Control..... 4
 - Pipeline Planning..... 4
 - Release Strategy..... 4
 - System Access Rules..... 4
- Configuration Management..... 5
 - Philosophy..... 5
 - Standard Topology Roadmap..... 5

Definitive Software Library	5
Application-Server Installation Standards.....	5
Installation Constraints.....	6
Source Control Management, Subversion.....	8
Access Control.....	8
Directory Structure.....	8
Tagging for Release.....	9
Other Directories.....	10

Communication, Email, and Ticket Queues

Requesting Support

Customers External to IS&T

Always instruct customers external to IS&T to use the list **isda-support@mit.edu** to ask for help. To triage any issue submitted by an outside customer, the ticket should originate in the Request Tracker queue **ISDA::Customer-Support**.

Internal, Departmental Requests for Support

IS&T staff who know they specifically require support or project resources from AMIT should submit requests to **map-support@mit.edu**. Depending on the request, AMIT will track the request in the **ISDA::Admin**, **ISDA::ProductionDefect**, or **ISDA::ProductionDeploy** queues in Request Tracker.

Email and Access-Control Lists

isda-support@mit.edu: Primary list for customers outside of IS&T to request help and support.

map-support: List to submit requests for support directly to AMIT by internal IS&T teams. We manually create tickets for these requests into the appropriate ISDA support queue.

amit-auto-deploy@mit.edu: Email sent to AMIT via the **ISDA::ProductionDeploy** queue or tickets forwarded from **ISDA::Customer-Support**.

amit-auto-defect@mit.edu: Email sent to AMIT via the **ISDA::ProductionDefect** queue or tickets forwarded from **ISDA::Customer-Support**.

amit@mit.edu: Master member list for the Application Management and Integration Team, used for internal team email and access-controls.

Request-Tracker Queues

ISDA::Customer-Support: ISDA's general customer-support queue.

ISDA::Admin: Work orders for any non-production application or system-integration work by AMIT.

ISDA::ProductionDefect: Specifically and only for tracking triage, troubleshooting, and resolution of problems with production enterprise services.

ISDA::ProductionDeploy: Specifically and only for tracking deployments to production enterprise services.

Process and Control

Pipeline Planning

ISDA processes change requests through the ISDA Pipeline Planning process. ISDA management chooses the members of this group. All tasks related to new systems implementations, enhancement, or problem resolution come through this group.

Release Strategy

DSPS' systems release alternately to NIST, Database Servers, or to Server Operations hosting for our own production support. System configurations are highly variable and therefore included in the run book for each given system.

Release Schedule

Productive Systems (staging or production): Release is variable. ISDA project coordinators negotiate a release window by coordinating with Pipeline, the project team, and the operational team responsible for the given system. This holds for systems defined as either “staging” or “production.”

Non-Productive Systems (prototype, development, test): Systems are installed in such a way that implementation teams should have full access to system configurations they need to modify. This is not accomplished by full root or full admin access. If permissions are not sufficient, coordinate change requests through Pipeline.

System Access Rules

- Only System Integrators in DSPS and System Operators in OIS root access to any server. Only System Integrators in DSPS and System Operators in OIS have interactive shell access to “staging” or “production” systems.
- System integrators perform all configurations and deployments to any staging environment in order to derive a reproducible build process for rollout to production systems managed by Operations.
- Application administrators responsible for provisioning can have access to web-administration consoles in production environments for their products.
- All developers can have shell access to prototype, development, and test machines as users other than root.
- All developers can have access to any machine they support as the “logs” user, which will allow them read-only access to certain parts of the system.

Configuration Management

MIT does not have an automated configuration-management system so the term *configuration management* is used to denote process and controls, not technology.

[Improvements to this process are under review via a project to build a configuration-management system in collaboration with OIS. ed.]

Philosophy

All systems are prepared as if they will be rolled out to data-center operations for production support. The reality is that DSPS supports some of its own production systems but formal release to Operations of all ISDA enterprise services is our ideal end state. Discipline in release management is guided by this principle.

Standard Topology Roadmap

We are moving toward the following standard configuration.

1. All systems virtualized on the VMware service offered by Operations.
2. All services should have four tiers: DEV, TEST, STAGE, PROD.

Most release-management processes in this document are designed around the assumption that ISDA will continue to move in this direction.

Definitive Software Library

System integrators store definitive versions of system installations and standard components on file-system trogdor in ISDA's co-located environment. Access is tightly controlled and by command-line. File management is manual. Run-book documentation for any give system is stored in the ISDA wiki.

[Improvements to this process are under review via a project to build a configuration-management system in collaboration with OIS. ed.]

Application-Server Installation Standards

DSPS supports a menu of standard application-server components. AMIT can redeploy these packages in any combination to provide any new project with a standardized platform upon which to build.

General	Java Application Servers	PHP Application Servers
RHEL4.x RHEL5.x Apache 2.x	JDK 1.6 SASHServer (SourceLabs Tomcat)	PHP 5.2.x PEAR

MySQL 5.x OpenSSL	Tomcat 5.5.x	
Development Environments Only: sFTP		

Installation Constraints

- Where we use standard packages (RPM, DEB), installed software stays in its default location.
- RPM or DEB based installations must use the package that comes with the operating system in question.
- For custom builds, /usr/local is our preferred location for software installation.

Hand Off, System Installation

Development teams do not have shell access to accounts on staging or production environments. Only system integrators in DSPS and system operators in OIS have that access.

This constraint is by design, to force a hand-off of installation and configuration, in order to evaluate installation procedures as operationally sound prior to release to data-center operators. We call this “vetting the run book.”

Web Console Access

Some applications have web-consoles for use in provisioning and system configuration. For products that provide this feature, DSPS will grant administration privileges to staff on the development team based on the requirements of the responsible team’s manager, coordinated through Pipeline.

System-Level User Accounts

We create standard system-level users on every machine: db, logs, repos, and www.

- Server applications run as these users, and not as root, for security reasons.
- Developers can use these accounts to access machines, since they cannot have root access. This access is permanent for development and test machines. It is temporary, for purposes of problem resolution, on staging and production machines.

The purpose of each account is as follows:

- "db" for MySQL or other database installations
- "logs" is to give developers read access to logs on staging or production servers.
- "repos" for Alfresco and other CMS installation
- "www" for Apache, Tomcat, and all other web-application server components

Only a fixed list of system integrators and system operators ever have root access to a web

server. This list is fixed for systems where ISDA provides production support, variable for systems released to OIS.

Source Control Management, Subversion

We assume the use of Subversion as provided by IS&T via the enterprise installation at svn.mit.edu. This rules will update as the capabilities of that system change.

Access Control

Our Subversion installation is integrated with Moira. Access control lists are stored there. Subversion's ability to send email notifications is also stored in Moira.

1. Always use a separate Moira lists for repository access control versus email notification.
 - a) Do not define the ACL as usable as an email list.
 - b) Do not define the email list as usable for access control.
 - c) The ACL can be the same as one used to control access to AFS lockers, remember that this list's membership should not spread past the project team or people to whom they report.
 - d) This email list is separate and discrete from human-to-human mail lists for projects. Not everyone on a project needs to be alerted about source-control updates. This is optional.

Directory Structure

Conceptual

[product]/[service]/[versioning]/[component]

Pseudo

[svn repos]/[top-level]/"trunk," "tags," or "branches"/[IDE Project or deployable unit]

Defined

Product: The arch, meta-project, brand, or other executive concept. This should not be a team or organizational unit.

Service: That which is useful, runs, or operates as a whole, a conceptual piece of an overall product that you are likely to manage as a coherent project.

Versioning: The standard SVN directories of "trunk," "tags," or "branches" denoting the primary version-control function of SVN.

Component: A collection of code that compiles as a single deployable unit (a war file, tarball), that which functions well as an IDE project, or both.

Example

We'll use a fictionalized multicomponent web service as an example.

- 📁 MitWebServices
 - 📁 GroupManager
 - 📁 trunk
 - 📁 GroupManagerWS
 - 📁 GroupManagerService
 - 📁 GroupManagerJNI
 - 📁 GroupManagerJMX
 - 📁 tags
 - 📁 Release- 0.1.1-20101104
 - 📁 Release- 0.2.0-20101104
 - 📁 branches
 - 📁 JohnDoesPrototype
 - 📁 PROD-0.1.1-20101104
 - 📁 CommonWsLibs
 - 📁 AnotherService
- 📁 SomeoneElsesRepos

By this strategy, you end up with a GroupManager source tree where the web service, underlying libraries, and instrumentation code are all managed and versioned together even though they have different deployment paths.

You also end up with a top-level repository that is used to manage the whole “product,” which can be vague but a meaningful categorization when projects change hands.

The content of the GroupManager sub-tree contains only that which is local to operating that particular service, but shared or dependent within that service. CommonWSLibs are tools shared and common to the overall MITWebServices product, like AnotherService.

Tagging for Release

Under “tags,” Version Capture

“Release”-[version number]-yyyymmdd

We do not specify any particular scheme for version numbers. Products, projects, and technology are too varied.

When you are done unit testing and prepared to move a service to a test system, tag that service directory as a “Release,” including your local version number and date.

Under “branches,” Production Support

Once the *tagged* release is approved for production deployment, *branch* it to a production release in the form:

“PROD”-[version number]-yyyymmdd

Again, discrete version numbers are up to the project team. Leave them out of the name if you are not using them or tracking by date.

Other Directories

This process is not designed to stifle the flexibility of source control. It is designed to give project managers and system operators a self-documenting way to find things across many projects. Beyond this basic organizational principle, other directories within a service, especially branching to prototype changes or tagging to capture mid-steps in the development process, is wide open.