

UPortal Developer's Meeting 9/24-9/25 at Columbia University

Attendees:

*Bramhall, Susan	Yale University
Durfee, Bernie	Interactive Business Systems
Ellentuck, Dan	Columbia University
Farmer, Jim (9/24)	IM+M
Fox, Dan	University of Maryland, College Park
Fracapane, Don	Columbia University
Gertz, Patty (9/24 only)	Princeton University
Gilbert, Howard	Yale University
Johnston, Andrew	Columbia University
Karchenko, Peter	Interactive Business Solutions
Lindholm, George	University of British Columbia
Markow, Jonathan	Columbia University
Mosca, Maria	Columbia University
Oltz, Mike	Cornell University
Pickens, Dave (9/24)	Sun Microsystems, Inc.
Rundle, Debra	Princeton University
Tilton, Justin	IM+M
Vigdor, Alex	Columbia University
Wallace, Dave	University of Delaware
* notetaker	

Timeframe for 2.0 is JA-SIG in December. All functions being discussed will be in 2.0. There is no plan to release without these features completed.

Monday 9/24

- 1) uPortal "State of the Union" [Jim Farmer]
What is going on in the world of JA-SIG and uPortal?
 - a) The release of version 2.0 getting close. (See above note.)
 - b) Performance testing began at I-Force and will continue.
 - c) Massachusetts Virtual Education System delayed uPortal implementation on advice that uPortal is not yet ready to support 1 million users. Delayed until April/May 2002.
 - d) Possible use by USA Dept of Ed
 - e) "Authentication is a barrier." UPortal is engaged with Shibboleth/I2 discussion.
 - f) Blackboard has tried to initiate standards for channels; they want to do their own portal. Were worried that XSLT would prevent meeting response time requirements.
- 2) Corporate Portal Conference Review [Justin Tilton]
What is the status of other portal software?
How competitive is uPortal in relation to these products?
 - 1) Most offerings are web front end to existing systems or application suites as opposed to general purpose framework. These focus on creating the ability to present application in small area in a browser. Almost all of the portals have a singular page, with very tiny channels. Whenever you pick one it becomes full screen.
 - a) Epicentric is the closest to the generality of uPortal. Their 4.0 release, not out yet, has XSLT channels. Ken spoke to them for about an hour. Very well designed and fleshed out.
 - b) PeopleSoft was talking about a university portal but would not show it.
 - c) Plumtree has 60% of the market and Epicentric 30%.
 - d) No one supports wireless devices – Epicentric works with 3rd party which takes html and does best it can to convert – also they don't feel it is important. Maria said there is a recent Gartner report to the effect that there is no market for wireless devices support in the portal.
 - e) IPlanet portal does support wireless - tour to be presented later
 - f) Portals are becoming a commodity.
 - g) Any features which offerings had but uPortal doesn't?
 - Management
 - Interchannel communication

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- Workflow via portal
- h) Blackboard may decide to use uPortal
- 3) Channel Efforts [Bernie Durfee]
 - What organizations are developing channels that can be shared?
 - What organizations are developing channels that can be purchased?
 - Who should keep track of these efforts? [Did not decide discuss this.]
 - Where should these efforts be recorded?
 - a) Princeton – EJB time collection system includes rules engine for union rules.
 - b) Announcements – Columbia - will share.
 - Add, subscribe, delete by topic
 - Needs to be Integrated with permission / groups
 - Servlet based function to share announcements with other portals(?)
 - c) Bulletin board - UBC
 - d) IMAP email channel - UBC (+IBS?)
 - e) News channel - Athabasca University
 - f) Course management - WebCT and Blackboard
 - g) Financial Aid [who?]
 - h) User Management - Stanford
 - i) IPlanet calendar – IBS
 - j) Briefcase – IBS
 - k) Notification – IBS
 - l) NNTP based Forums –IBS
 - m) WSUI – web services User interface invented by Epicentric – SOAP based + XML – no java needed use XML only
 - Take advantage of all the web services components
 - JA-SIG wants to be observer on WSUI standardization
 - n) Repository – UDEL will host site
 - Content as well as new classes
 - In progress channels
 - Desired channels
 - Post XSL / XSLT as well as java
 - Get updates periodically
 - o) IBS channels will be available in binary
 - p) Jim: A digital books channel? (Making of America project and other similar)
- 4) Performance Enhancements [Peter Kharchenko]
 - Caching/Multithreaded channels: How are ICachable, IMultiThreadedChannel, IMultithreadedCachable used?
 - How do they work?
 - What are the benefits?
 - What is SAX event caching versus character caching?
 - a) All changes made before the performance testing:
 - Cache both stylesheet root and stylesheet set - turned out to be not very effective.
 - One thread per channel – expensive so added a threadpool object.
 - ICachableChannel interface which channel can implement
 - (1) key, object – scoped by channel instance or system wide – other users get cached version – example=login channel (before auth)
 - (2) Key represents unique state
 - (3) Key must also be specific to stylesheet used
 - (4) Implementation uses soft hash map – garbage collection can remove objects if necessary
 - (5) What does it cache? SAX output stream
 - Multithreaded channels
 - (1) E.g. for every rss channel have one shared instance of the rss channel class
 - (2) Need a way for channel class to be able to handle multiple channel content
 - (a) Like IChannel but takes a unique id for channel on each call
 - (3) Exercise care to make sure all methods are thread safe.

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- And IMultithreadedCacheable – CGenericXSLT implements this.
 - Advice: do cacheable first – easier and biggest improvement – next multithread
 - Finally created character cache – treats output as a blob - all done in the framework
 - (1) Were using Apache serializer classes
 - (2) Replaced this with caching serializer – extends org Apache packages
 - (3) Issue about whether to keep classes in Apache tree or whether we rename – class Peter extended is declared final so had to reimplement it – ugly
 - (4) Each channel output is cached as a string if use proper key – channel can just return same output.
- b) Would like to cache big XSL transformation of layout (structure / theme) (done?)
- c) Guest specific acceleration
- LayoutBean is now UserInstance and has a child called guestUserInstance. Right now it looks in a fixed table of guest usernames, if it's there it uses the latter class; there's only ONE instantiation of it. Has only one instance of the Layout object for all guest users. (Layout object is a big DOM object.)
 - How to determine if Guest user? Open question – in IPerson?
 - Shared layout object for all users
 - Cache as much as possible
 - Channels on guest layout should be multithreaded and cacheable
- d) Peter got some ideas from Zed's UBC report but didn't do everything the same way. - version 2 vs. version 1. Some further optimization may add more of the UBC ideas.
- e) Holly Cross performance problem?
- eventually sped up their performance by 25 fold.
 - It's the channels, NOT the infrastructure, that makes it slow.
 - Their decision to pull the plug on uPortal was made by Holy Cross president before they completed tuning.
- f) Jim Farmer summarized:
- Concentrated on getting XSLT to be faster.
 - The character caching made a bigger difference than expected, others less.
- g) Someone asked what if we could write channels in JSP? Peter said, in current architecture that would be very difficult to do. This design was made to be cleaner and more closely integrated, and flexible output. Maybe it could be pluggable channel architecture in some later release.
- h) Michael Kay (author of XSLT Programmer's Reference) expects XSLT to get factor of 3 faster (for a given processor) over the next 2 years, but not much more than that.

5) Howard – IPerson Attribute Mapping [Howard Gilbert]

How do channels obtain person attributes?

How is this related to the EduPerson standard? Iperson /eduPerson

a) Presentation by Howard:

=====

At Yale they have students, alumni, employees, etc. For idiosyncratic reasons most groups at Yale are in the human resources database except for alumni. Yale uses kerberos netid's. But it may vary by institution what the format of the userid is.

When someone 'logs in' the first thing that happens is the authentication step. Then there is additional information in the category of "directory information". Yale has checked in a new service called PersonDirectory driven by XML configuration called personDirs.xml which indicates sources (plural!) of the user info. Each piece comes via LDAP/JNDI OR some other database via JDBC. Map each incoming attribute to an IPerson attribute - if it has the same semantics as (e.g. "eduPersonAffiliation") the use the eduPerson attribute name. You can also have another analogous variable which is site specific. "eduPersonPrincipalName" (EPPN) should become be a globally agreed on user identifier in the Shibboleth system. Yale doesn't consider the thing used to look up users in the database, to be the same thing as the person's logon ID. Howard recommends that we all use two things, the eduPerson part, AND associate it with some other identifier in case we want to change the latter thing later, without losing the global name.

=====

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- b) Questions about eduPerson as a standard
 - There is a question of whether eduPerson is important to EduCause.
 - Does MACE have some kind of LDAP directory thing in it?
 - (1) Jim Farmer said yes but there's no eduperson kind of structure to it.
 - (2) There was a discussion of need for a reference implementation for eduPerson – the consensus from people at the table was that it is not in scope but needed by other schools.
 - Patty asked -- What about the OKI project? They have come to the point where they need APIs. The first one is Directory, the second is .. Authentication. They were asked if they tested their APIs against existing implementations, e.g. Directory against LDAP? Nope. [these responses from Jim who spoke to the OKI people a few days ago]
- c) Summary
 - GlobalUID - use of this attribute in Principal object to represent local portal ID is incorrect. We may eventually use a globally unique identifier in uPortal but it may reside in IPerson EPPN instead of principal.
 - New PersonDir Service driven by XML file which maps attributes from existing sources using LDAP or jdbc. Any attribute can be mapped to any IPerson attribute. When the semantics of the attribute are the same as the eduPerson semantics, the eduPerson attribute name should be used. Process for mapping and specifying driver and connection is document in comments in the code.
 - Optimization to pool jdbc connections and jndi contexts is a good idea. Currently no caching of directory connections (either jndi or jdbc).
- 6) Sun I-Force Load Testing [Peter]
 - Overview.
 - What were the results?
 - What were the lessons learned?
 - What are the bottlenecks in performance and how can we fix them?
 - What further testing is needed?
 - a) User Actions
 - Page refresh
 - Guest init
 - User init
 - Typical guest session
 - Typical user session
 - b) Layouts
 - Empty
 - Single lightweight channel
 - Single Heavy channel
 - Typical
 - c) Caching
 - Stylesheet set
 - Stylesheet root
 - Channel cache
 - Transformatio cache
 - Character cache
 - d) Log verbosity
 - DEBUG
 - ERROR
 - e) Results – get report from Peter
 - Trans / sec (pages rendered)
 - (1) Got about 40/sec on first test
 - (2) Got up to 230-240 with character caching
 - Response time – ends less than .2 seconds
 - Helps to run multiple tomcats per apache web server
 - Pooled database connections using JA-SIG pooling – works well.

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- DB considerations
 - (1) Joins
 - (2) Prepared statements
 - f) Next round – multiple web servers, ssl
 - (1) How to test non-sun hardware
 - g) Debra at Princeton ran load tests last week. They were running Weblogic 5 on NT, and Weblogic 6 on Unix, the latter gave better results, both in the neighborhood of 35-40 pps. Lots of issues about hardware and software configuration
 - h) What is a typical user? Student vs commercial user – we don't know this yet
 - i) Where should our efforts go next?
 - Database.
 - next version of XALAN
- 7) Layout Management [Justin Tilton]
What happened to the subscribe channel?
What do we do about Netscape problems?
Is the current design suitable?
Should we build more layout management features directly into the layout itself? Example: [X] button
- a) Subscribe channel – old one works and is fast but
 - visual analogies may be hard to understand, especially first page
 - preview not implemented
 - save changes needed
 - independent of structure or theme
 - works for old Netscape browsers
 - b) Newer preferences channel is a whole portal with channel
 - Can move channels and columns around, add tabs, columns
 - Can also add a new channel (=subscribe)
 - (1) Browse through categories and place channel
 - (2) Problem with Netscape 4.7 browsers because of too many nested tables
 - Maybe the best way to do it is to do it in portal view – button switch to more for changing preferences or layout.
 - Add ability to supply channel parameters at subscribe time is intended to be implemented.
 - Will add save / cancel instead of on the fly commit
 - c) Other ideas
 - Use JavaScript? But the only place that uPortal uses JavaScript is in the Detach to separate browser window. Very hard to get it to run on all browsers.
- 8) User Administration
- a) Yale UserIdentityStore
 - Implementation takes Boolean. If no user_id in uPortal after authentication then create user data on the fly from a template user. Template user determined from attributes in IPerson object.
 - Feedback indicated a number of changes which will be implemented and documented.
 - (1) Don't create new layout – use the layout for the template user until the layout is saved with changes (Needs a change in RDBMUserLayoutStore to retrieve layout for another user.
 - (2) Skip use of UP_ROLE table and map straight to user (roles are being replaced by groups, see later topic.)
 - (3) Send code to list for comment / integration
 - b) DBLoader enhancements from Newfoundland should be incorporated
 - c) Stanford's user channel – not for reference implementation
 - d) Other administrative interfaces – needs coordinator to ensure they will all integrate
 - Users
 - Logging
 - Profiles
 - Properties

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- Groups
 - Stylesheets
 - permissions
- 9) IPlanet portal – Sun [David Pickens]
- a) Supports wireless devices. WAP comes into the server on a different port, which is how it knows to use different UI.
 - b) <http://wportal.demo.iplanet.com> for a demonstration portal.
 - c) Looks VERY similar to uPortal
-

Tuesday 9/25

- 10) Channel Administration [Justin Tilton]
What happened to the publish channel?
New features and workflow.
How should we implement the modification of channel parameters during publishing and subscribing?
- a) Publish channel is still there - New “publish” more like other preferences channels
 - new channel or modify channel
 - workflow schematic + wizard-like feel
 - driven from XML file known as CPD file
 - can review or cancel at any time or back up in workflow diagram
 - passes whole channel registry in XML to get all categories
 - categories presented hierarchically not as one scrolling list
 - Presents all channels from the channel registry
 - Needed transaction concept because of dynamic building of request
 - (1) Builds up an XML document element by element – cancel discards the document
 - b) Will consider preview but it's difficult because requires new path through render – may not be so necessary since can now go back and edit or delete.
 - c) Where should category management go? Same or different channel under discussion
 - d) How to determine display order of channels? Current sort by alpha in XSLT – could be priority
 - Decision – sort by priority from db to XML
 - Allow user to sort alpha
 - e) Publish date, approve date
 - George Lindholm asked what is the "publish date" in the database? It is supposed to be "do not publish the channel until a specified later date". What is the "approved" field? There will also be a "who published it" field.
 - f) Considering conversion of CPD file to XML scheme – deferred to “3.x”
- 11) Profile Administration [Peter Kharchenko, Justin Tilton]
What is this screen supposed to do?
Should we allow users to have separate layouts for each profile?
What should we do about current problem which lets users delete their current profile?
[Peter, I wasn't sure what to fill in here - maybe you can summarize questions and what you heard.]
- a) New user does not need any profile – use system profile only
 - b) Users do not need UA mapping – can use system / guest
 - c) Should user be able edit profiles or create new ones?
 - No – user can choose from system profiles only – add image for easier selection
 - Need edit ability for administrative users to create new system profiles
- 12) Authorization/Permissions and Group Management [Bernie Durfee]
What is the API for determining user roles and permissions?
Can channels define roles?
Can administrators define role mappings?
What authorization "hooks" are missing from the framework?
- a) Split role management from authorization

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- Permission manager
 - Group manager
 - Maybe a help class between them
- b) Permission System: Bernie - (get ppt file)
- Permission means: Can principal X perform action Y on target Z?
 - Semantics depend on the creator/user of the permission.
 - Principal: Who does this permission effect? Actually a 'Principal' is just an actor -- a string, not anything formal. Call it "actor".
 - Action: What activity?
 - Target: What is the action affecting? ChannelID, GroupID, whatever
 - Owner: Who does this permission belong to? They belong to a Channel.
 - Manager has setPermission and getPermission to get and set permissions for "actor" (which could be user, group or other id)
 - PermissionManager instance is created by the framework for the channel.
 - Channel A can't edit channel the actor's permissions.
 - Where do you check? Bernie is planning to do this when a person tries to render the channel. At least once per login. Because the person may have subscribed a year ago and the permissions have since changed. The infrastructure will call a method in the channel to actually do the permission check.
- c) Needs to be done ...
- Performance tuning: This gets called very often by nature, so it must be very efficient.
 - Administrative interface.
 - Integration with Publish/Subscribe.
 - Integration with GroupManager. -- need a helper.
 - Can channels define roles?
 - Can administrators define role mappings?
 - What authorization "hooks" are missing from the framework?
 - In the login channel there is an example of Authorization using permission manager.
- d) Group manager – Dan Ellentuck
- Given a user, what are the permissions?
 - Use role hierarchy
 - Multiple inheritance
 - Can use groups for other things than authorization
 - Every permission has an owner and name of actor
 - Groups can belong to groups and inherit permissions – must check for circular references
 - Roles will be converted to group (I need to talk to Dan about how to get default layout from the interface.)
 - Use user id from IPerson.getID not whole person object
- e) User manager needed to construct IPerson based on a uPortalUID (Howard)
- New implementation of code from PersonDirectory service
 - Don't present all information gathered at login – some subset
- 13) Database issues [George Lindholm]
- Which databases does uPortal support?
Does DbLoader need improvement?
Are our database interfaces appropriate?
What best practices should we be aware of to make uPortal work with as many databases as possible?
- a) Which do we support?
- Oracle
 - MySQL – some problems with lengths, no transactions – is there a way to do this? JDBC 3.0? Not our issue. Need to make it work in DbLoader
 - Sybase – problems with names and text which have 255
 - Informix – doesn't handle long strings – needs byte arrays
 - SQL Server – IBS resolving some problems

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- PostgreSQL works fine
 - b) DBLoader – what should be changed – needs owner
 - DBUnload
 - (1) Would help forward migration
 - (2) Might help creation on new layouts for users
 - Enhancements by Newfoundland
 - (1) Should be integrated
 - c) Specific database issues in table creation
 - Example oracle tablespace, index tablespaces etc - how would we support these?
 - Referential integrity – useful during development but becomes sticky later.
 - These are really not worth how hard the work to do it
 - Add Separate XML to create sequences for oracle
 - d) Interfaces
 - merge some of oracle specific impl up to generic
 - merge IUserLayoutManagerStore and IUserPreferencesStore and IcoreStylesheetDescriptionStore
 - also merge implementations
 - e) Howard will optimize directory access independent of rdbm services
 - f) Best practices
 - Stick to singular (as opposed to plural) names for tables
 - Should get datasources from JNDI – George may have time to do this from Bernie's start.
 - What about nonstandard join syntax? George wrote Oracle optimized a version of the RDBM. He wants to add JOINS and PreparedStatements to this, and then merge his changes back into the base RDBM classes.
- 14) UPortal Best practices [Peter Kharchenko]
What code has been written in UtilitiesBean and org.jasig.portal.utils package that we should all be utilizing?
What are the meanings of the org.jasig.portal subpackages: utils, services, tools, helpers?
- a) XSLT helper class made more like a bean –
 - Old – used static methods with many parameters (These should be retained for those who prefer.)
 - New – instantiate and use set to specify XML, stylesheet etc.
 - b) Packages - Bernie
 - Factories – put the one class from here into utils (Bernie)
 - Tools – external to runtime
 - Utils – generic for all portal classes
 - Helpers – move the one class from here into utils
 - Services – channel services – also framework services
- 15) Housekeeping [Peter Kharchenko]
Should we consolidate properties files?
Should we reorganize uPortal files to comply with the WAR-file standard?
Should we modify existing code (and IChannel!) to use Xalan 2/SAX 2?
Can we remove all direct references to Xalan/Xerces so that a properties file controls the XSLT engine and XML parser?
Should we deprecate GenericPortalBean and UtilitiesBean and find new places for their methods?
- a) Properties files are growing – consolidate? (not sure of decision - Was this a yes?)
 - b) Go to XALAN 2/SAX 2? Go from documentHandler to contentHandler – YES!!! Peter
 - c) Should framework be able to switch implementations of XML parser – independent of XERCES etc.
 - d) Some schools wanted the portal to be completely XSL/XML implementation independent, but that would be difficult before JAXP is out. Peter doesn't know if JAXP will solve Xerxes dependence.
 - Look at it and describe issues – say why we don't - Peter
 - e) Deprecate genericPortalBean and utilitiesBean – yes – use a util class – Ken or Peter
 - f) Comply with WAR-file standard?

UPortal Developer's Meeting 9/24-9/25 at Columbia University

- UBC and Princeton will determine what's needed to allow WAR file distribution – what needs to be moved – properties should go to WEB-INF
 - Have to copy classes and rename classes from serializer of apache – Peter will talk to apache folks
 - Implications for repository
- g) Bernie – work on cleaning up LogService and removing deprecated Logger.
- h) Clean up database table names where unclear – Peter and George
- 16) Documentation – [Debra Rundle]
- What current documentation do we have?
- What is missing or outdated?
- Who is working on it?
- a) Princeton –
- Lee Novak is working on uPortal doc part time. She has talked to Bill Brooks.(Cal-Poly).
 - User doc
 - Portal installation doc
 - Administration doc.
 - Enhancement / channel development
 - Princeton will also will create relational model