daemon
Directory of Business

The Organisational Directory for Government

Technical Overview

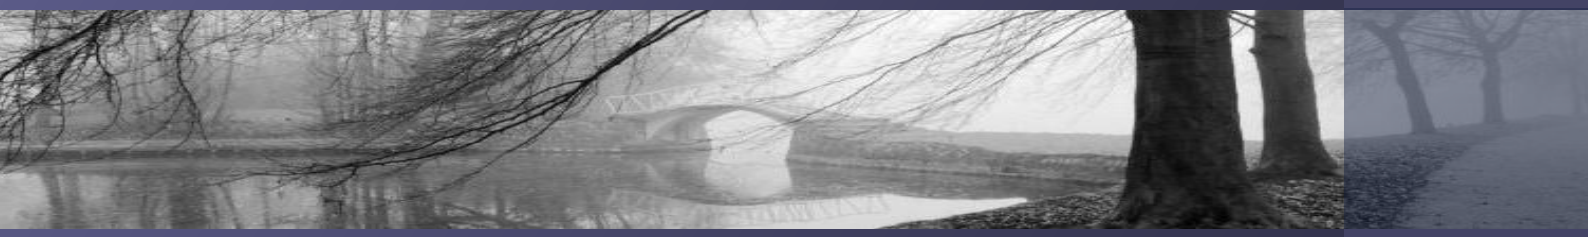

daemon
DIRECTORY SERVICES

- POWERFUL WHITE AND YELLOW PAGES DIRECTORY
- INTUITIVE GOOGLE LIKE SEARCH CAPABILITY
- FAST AND EFFICIENT SEARCHING
- ENTERPRISE LEVEL INTEGRATION
- LDAP, ACTIVE DIRECTORY & SECURE TOKEN SERVICE (STS) INTERFACES
- FLEXIBLE UPDATE MECHANISMS
- WEB SERVICE DELIVERY AND SUPPORT
- SECURITY ACCREDITED
- FAST IMPLEMENTATION

dæmon
DIRECTORY SERVICES

Gaston House, East Bergholt
Colchester, Essex
CO7 6SD

Tel 01206 299288
Email admin@dds-labs.com
Web www.dds-labs.com

For further information

Daemon Directory of Business
Organisational Directory for Government

Technical Overview

Document Version
DOB - Tech Ovr v6.0 Sep 2011.doc

Readership

This paper describes the technical architecture the DDS Directory of Business (DOB); a database developed to meet the Government's need for a secure repository of personal and organisational information that is easy to search and easy to update.

The document is expected to be read by Technical Architects needing to understand how the system operates, typically with an interest in potential use of the system in their own environments

Further Information

To contact the DDS Government Directory Team to ask for a demo, discuss the business implications or talk through the technical architecture, Email info@dds-labs.com or call + 44 (0)1206 299288 to

Visit http://ww.dds-labs.com for more information on the DDS Directory Suite of products

Contents

.

1. Overview

The DDS Directory of Business provides comprehensive directory information for an organisation and its partners.

The scope of the information is organisational; i.e.: the system holds data beyond the standard person-only directory such as Outlook or Active Directory 'address books', holding organisational data including people, jobs, roles, skills, organisational functions, teams, office location, etc., thus providing a comprehensive body of knowledge about the organisation that can be accessed in a consolidated way with simple Google-like searches.

The application can be installed to run from local customer servers, or is available to Government customers on a Software-as-a-Service (SaaS) basis through the GWS Crown Framework contract.

Data in the DDS Directory is maintained through a combination of user self-service, designated team administrator updates and synchronised data imports. These update mechanisms combine to provide a secure white and yellow pages directory service, that's a rich source of corporate information

Core person-data in the system will typically be synchronised with an organisation's Active Directory and Email address-books. Other organisational information sources with relevant information can be included in the synchronisation process; thus business rules can include imports from an HR staff database, user SharePoint "My Site" web sites, or even separate spread-sheets holding relevant staff or project details.

The Directory allows basic personal data imported from external directory sources to be enriched with additional information provided by staff and team authors. Links between externally sourced and internally enriched content are automatically maintaining. Business rules can be defined so that certain data sources are treated as 'authoritative' over segments of the data; e.g.: imports from the organisation's Active Directory can be regarded as the list of active staff in the organisation, with new additions and omissions defining the on-boarding and off-boarding process for people in the Directory.

Most organisations work with partners, with whom the staff typically need to engage. The DDS Directory caters for this by allowing imports of external partner directories into the main directory database, with partner data partitioning off and subject to different data management rules. This enables a single search request to span an organisation plus partners whilst the data itself is subject to different sets of import, field definition and security redaction rules.

Keeping the system easy to use without losing functionality has been a goal of the system design. The system is simple to use - experience has shown that no training is required; searches are easy to understand, being based on Google type logic with results presented as a ranked list. Where users need more granularity with their searching, there's an advanced search function allows users to use combinatorial logic, with field-level filters, and Boolean (AND..OR..NOT) functions to define more focused searches. Output options exist to deliver results as spread-sheets, distribution lists, or org charts. For organisation searches, where free text search strings are used the Directory employs an optional thesaurus function to find acronyms and alias terms; thus the thesaurus helps ensure that searches for "Information Assurance" also pick up "IA" and names searches for "Elizabeth" pick up "Beth" and "Liz".

The DDS Directory support user single-sign-on, so that users identities can be automatically recognised, and them made responsible for their own personal profile information. A "data aging" mechanism detects aged data and warns the data owners to update or at least confirm the data.

All aspects of the organisation and its partners are represented in the organisational tree, and team updating is done by designated team authors, who are granted rights to parts of their teams in the tree by locally appointed administrators, who in turn are granted their rights by managers 'above' them in the tree. In this way management of the system can be expanded to cover very large organisational groups, with many millions of staff and 000's of teams.

Directories hold large numbers of personal records and so security is a major concern. The Directory includes a role-based redaction system that restricts the view of the data according to who is doing the viewing and from where. The system also complies with the strict privacy rules as set by the Data Protection Act apply and has been fully penetration tested and security accredited to meet those requirements, and is currently accredited by the Government to IL3 (RESTRICTED) level.

2. Application Architecture

All services are provided through a componentised service orientated design that produces a modern system architecture that is modular, agile, and scalable.

The logical architecture is illustrated in the Figure below:

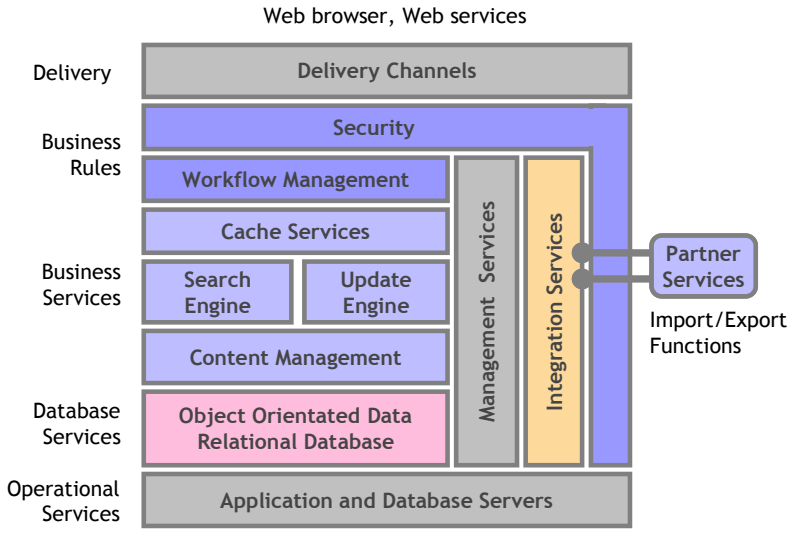

Figure 1– Service Orientated Logical Design

Delivery

Browser Support - The system is accessible from any user's standard desktop web browser, and self-adapts to different browser capabilities, thus exploiting the latest browser technology whilst still working with early browsers back to IE6. The system is fast – with sub second responses to searches, due to the clever design of the system and extensive use of caching technology.

System Interfaces - The Directory supports other systems that need directory information and is commonly used by other Intranet systems to provide staff details for user personalisation, and for information on organisational teams and offices. All user level searches that are possible have equivalent web service interfaces that other applications can call. There are three primary interfaces supported for this:

- Web services – there are Web service interfaces on TCP 80 that enables other services and applications to share the Directory service data.
- LDAP v3 – the Directory supports LDAP services on TCP 389 for person, team and build searches
- Secure Token Server (STS SAML 2.0) – the Directory supports STS services so allowing it to participate as an Identity Provider in a SAML based integrated Federation solution

Accessibility Access Compliance - The Directory has been accredited for Accessibility Access (AA) in compliance with the Disability Discrimination Act and supports systems that serve special needs access, such as the 'JAWS' text to speech converter.

Business Rules

Workflow Management - The Directory has Workflow, Object Caching and Security services presented to the system as sets of Business Rules, configurable by a hierarchy of users to allow for the flexibility of control necessary to manage diverse directory content according to strict security rules.

Security (Accreditation) - The Directory is used in government where personal information is covered by stringent security rules. The Daemon Directory of Business has the ability to present limited views of the data according to who is viewing it and from where. The application software has also been fully penetration tested and security accredited to IL3 (RESTRICTED) level environments.

Security (Redaction) - The visibility users get of Directory information can be limited by sets of business rules. The business rules are highly configurable, allowing the view to be limited according to the type of

user and the network through which they have opened the session. Business rules can be applied at both 'row and column' level to give a high degree of flexibility in defining access rights.

Business Services

The DDS DOB combines 'white and yellow pages' information and allows users to search for staff and organisational information. This creates a richer data set which users can access to find information about individuals, groups of staff, teams, buildings and roles.

The Directory database uses a directory tree' to represent the organisation's structure teams, people and offices. The directory tree can be browsed and 'drilled down' and the entire content can be searched using 'free text' keyword searches to identify content. These directory trees help to provide a visual representation of the organisation that staff can use to place the content they have searched for in context, to provide a better understanding of the organisation as a whole and how different areas interact.

Search Engine - There is an internal search engine that provides Google like searching based on simple, Google-like logic that most users will already be familiar with. Experience shows users do not need training to use the Directory and quickly find it indispensable. Advanced searching facilities are available for users who need greater granularity of searching. These include multi-field and Boolean (AND..OR..NOT) searches, and search output can be viewed as lists, team-views, org charts, or downloaded to spread-sheets, email distribution lists, label documents or the clipboard, for access via the users' standard desktop. There's also an optional 'thesaurus' to ensure expansions of acronyms or equivalent terms aren't missed.

The Search system can be configured to use either an internal or external search engine, a set of Web service calls being available to interface to search tools such as Autonomy, Lucene, or FAST.

Update Engine - The data in the Directory can be updated in a number of ways, e.g.:

- Self-service; staff keep their own records up-to-date (with limited rights over what can be updated)
- Team Authors; who maintain the content of team pages and the structure of the organisational tree
- Local Administrators; who manage sub-trees and have rights to appoint other administrators and authors beneath them
- Automatic imports of content from other systems, e.g.: the HR and Email systems, with
- Data aging; there's an inbuilt workflow system that alerts users when the information they're responsible for hasn't been updated for a while, reminding them to update, or at least confirm their details. This is combined with innovative options for displaying aged data, visibly indicating that there is a lower confidence in these results.

When updates are made in the Directory, the changes can be synchronised with other systems in the organisation (e.g.: the Email and HR systems) so they can keep in step, either programmatically or manually.

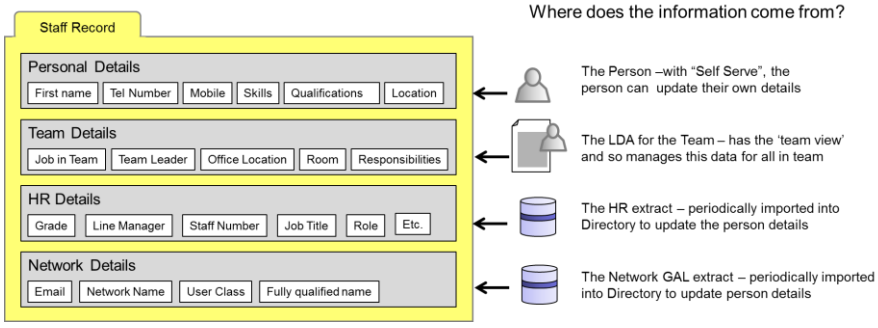

Figure 3 – Who updates different parts of the staff record in the Directory

Content Management – The Content Management functions are central to the Business Services layer of the system. These maintain the referential integrity of the various parts of the information store whilst ensuring that all updates are properly sequenced, content backups and audit logs written, and multi-user access managed. There is a powerful object caching mechanism that ensures delivery of content is rendered and delivered in performant way, re-using pre-rendered objects content wherever possible.

Database Services

Database design - Directory content is organised logically as an Object Orientated data model and held physically in a relational database. This makes it accessible to standard reporting and data management

functions whilst maintaining the flexibility to represent a rich set of directory objects, (people, teams, offices, narrative, etc.) .

All data access is provided through a set of Database services implemented as SQL Server Stored Procedures. This provides extra performance and security allowing the application and database functions to be clearly separated and potentially operated on different systems in separate security zones.

XML Metadata - The Directory provides for local variations in data between datasets through the use of metadata. This allows different agencies in the system to have different data structures. Holding information in this way allows the Directory to maintain information on the natural 'matrix management' situations that occur in the modern organisation.

Operational Services

The Daemon Directory of Business has a long pedigree – it was initially developed in the mid '90's to meet the needs of large central government departments and agencies, clients including the Home Office, Dept. for Work and Pensions, Dept. of Health, Dept. for Transport, and Department for Communities and Local Government.

The Directory has now been re-developed to operate on the Microsoft .NET 4 platform with IIS7 and SQL Server 2005/8 and uses standard Microsoft clustering technologies for scalability and resilience of web, application and database servers.

The Directory has now been re-developed to operate on the Microsoft .NET 4 platform with IIS7 and SQL Server 2005/8 and uses standard Microsoft clustering technologies for scalability and resilience of web, application and database servers.

The system anticipates high volumes of data stored ad retrieved and has been tested with 2-3m staff records. The largest single implementation at the moment has information on 125,000 people, operating in thousands of teams, over a dozen partner agencies. Usage is continual during the working day, varying from 200-300 hits/minute average to 600 at peak time. Responses under these sorts of load are still virtually instantaneous, the software having been designed with efficiency under load as a prime objective.

Import / Export Functions

Partner Services - Modern organisations work with partners and staff need to contact colleagues across these partnerships. The Directory 'Integration Engine' links information from partner organisations into the corporate Directory database providing an immediately searchable and cross-referenced view of the enterprise.

Extensive batch data conversion options exist to migrate and include data from other directory sources in thus giving users fast, consistent access to directory information from external agencies.

Integration System - The Directory has a flexible content integration and reconciliation engine that allows directory information from multiple sources to be managed and/or integrated via a variety of input feeds. This provides a configurable set of tools, driven by business rules, through which directory information from external systems in partners' organisations can be imported and reconciled. Customisable export routines also allow the Directory to be user as a Single Point of Update for external systems.

4. Information Architecture

The Directory information store is a relational database. There are four main entities: Agency (organisational grouping), Team (organisational element), Staff (person), and Location (Building).

The application makes extensive use of database stored procedures as the data service entry point and will operate on either Oracle 10g or SQL Server 2005/8.

Because the directory is capable of holding personal information it must be protected. This is done at the data service level with a field level Roles Based Access Control (RBAC).

An example of the logical data model (LDM) is provided in Figure 4 below.

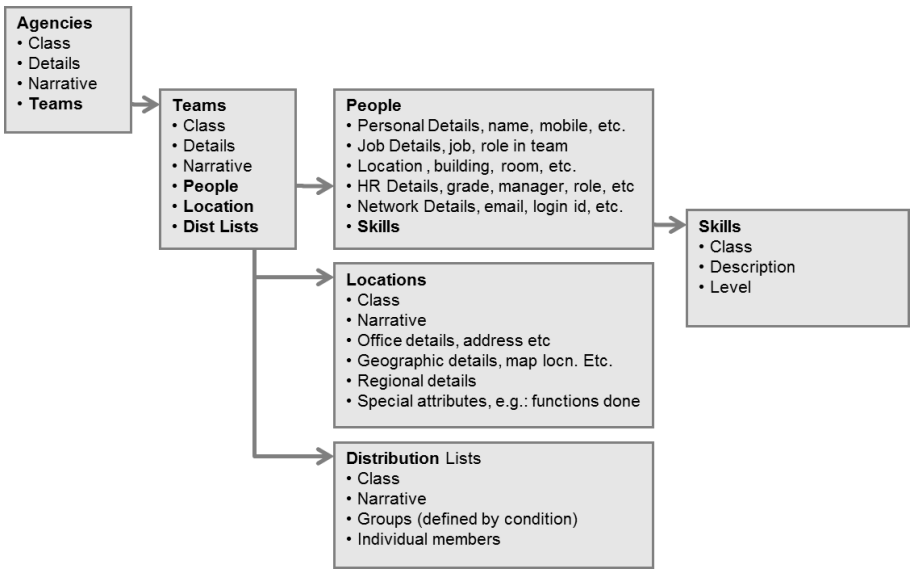

Figure 4– Directory Logical Data Model

The information is stored in a relational database with a fixed record structure. For efficiency the database is partially de-normalised and uses table structures with fixed field definitions for the core data entities.

That said, the business requirements for managing data demand more flexibility than can be accommodated by fixed field structures and so all of the key entities allow for extension fields to be included that are defined by XML metadata. Thus the system can easily handle different structures for staff, team, building and skills in different partitions of the database, each representing the requirements and availability of information from the agencies contributing their data to the directory.

The internal Search engine uses a word list which is maintained as a set of data tables. The data held on words includes pointers to their parent objects in the database, their proximity to each other (so proximity word searches can be made) and their type (so that field level searches can include/exclude them.

Search results are also filtered through the redaction engine so that searches are similarly constrained by redaction rules as the content they point to.

The search system includes a thesaurus database that links an alias list and a list of familiar names to the Word index. The thesaurus is managed separately, either by importing external data or updating the system manually.

5. Interface Architecture

The Directory delivers content to users for viewing and updating through a standard web browser. Both Microsoft and Open Source browsers are supported, using JavaScript and AJAX for an enriched user client experience.

Interfaces Supported

The application has a rich set of user service interfaces to support the integration of external data from both full and 'partial' partners' organisations. This is illustrated in figure 5 below:

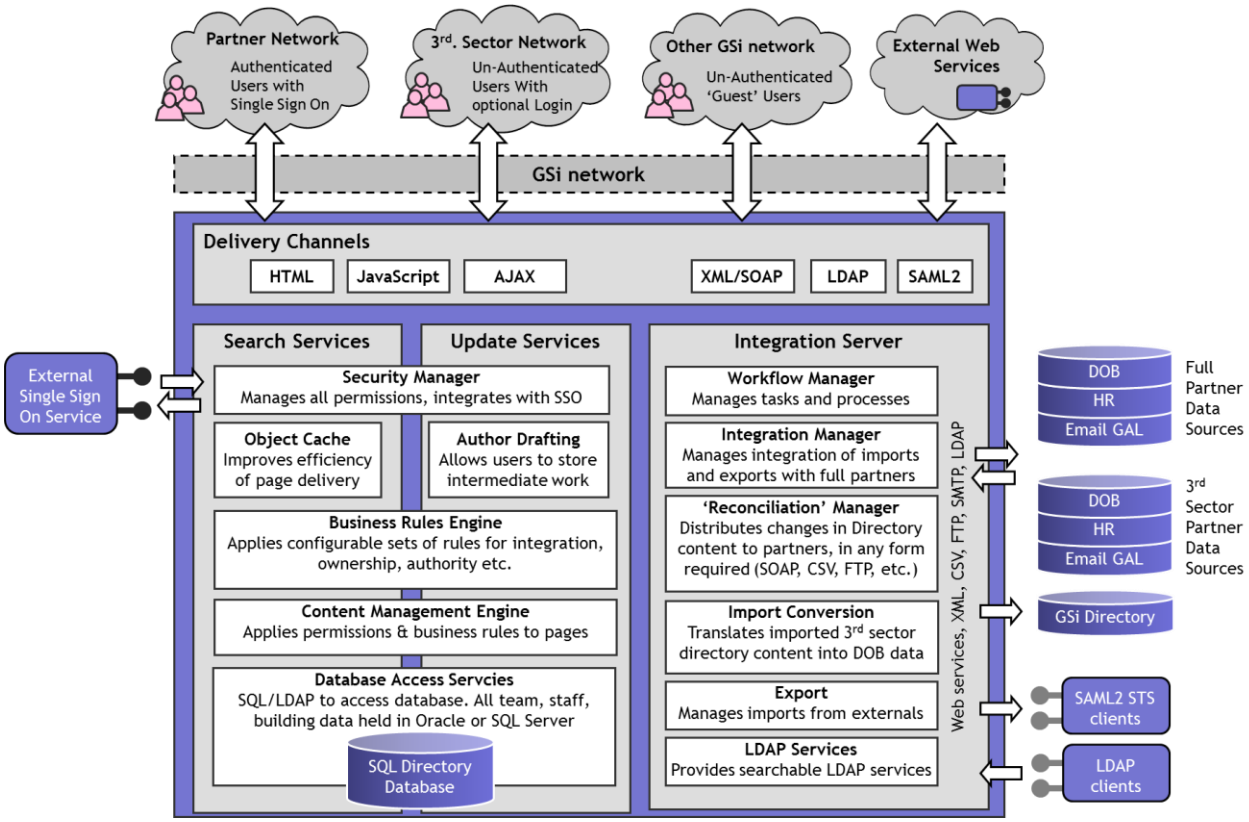

Figure 5- Interface and Component Architecture

The interfaces include:

Partner Directory Imports - the Directory can import and synchronise directory information from virtually any source, including Active Directory and eDirectory, using XML configuration tags to map fields, all managed by the business rules in the DOB Integration Engine.

LDAP v3 Server - the Directory supports the capabilities of LDAP v3 for those systems that need this. There are LDAP schemas for people, organisational structure, teams, and offices.

Secure Token Services (STS) - The Directory has an integral STS server that supports identity federation systems using wither SAML 2.0 or WS*-Trust. The DOB STS server can be configured to work as an Identity Provider (IdP) with identity federation brokers such as Microsoft Forefront ADFS 2.0, IBM Tivoli, or Oracle Identity Federation Server

Government Interfaces - The DoB can both import and export directory data as text CSV files, in a format compatible with the C&W GSi Local Government Directory

Directory Update Exports - The Directory can be configured so that when updates are made there's a neat synchronisation capability that can feed the updates to other systems in the organisation (e.g.: the Email and HR systems) so they can keep in step, either programmatically or manually.

Integration with Partners

Partner Integration - The Directory will work in a 'federated' context, where the organisation view is composed of staff & organisation directory information from multiple sources. This includes the situation where some data is 'mastered' on system operated in separately secured networks.

The Integration Server - Integration of diverse data sources is performed by the Integration Server. Once integrated, the Directory provides a single, consistent and searchable view of the data across the entire enterprise.
The Directory is a fully Web services enabled application that delivers information on staff, organisation and offices for users to view through their web browser or through XML service-to-service exchanges.

Mastered vs. Copied data - The Directory provides read and update facilities for locally 'mastered' directory data as well as imports of read-only information from partner directories. This is illustrated in Figure 6 below.

Directory Federation

Being able to link directory content from different agencies enables a single 'federated' view of the directory to be presented to the users, whether agency data is being mastered on the directory or not.

This is illustrated in the following figure, which shows the Directory working with content from multiple agencies, some mastered, and some not. The users searches and retrievals would work across the entire directory space, or if required could be 'filtered' down to work on single or groups of agencies.

Figure 6 below illustrates the way full and partial partner data is managed in the Directory:

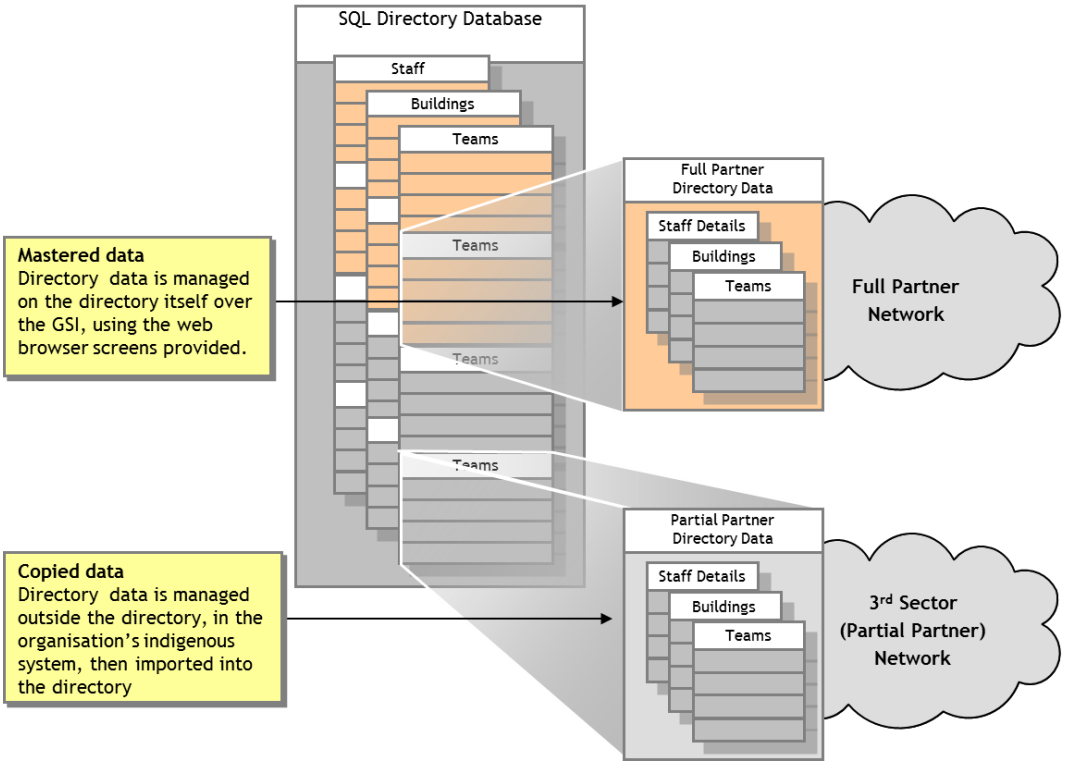

Figure 6 – The 'federated' context, with Mastered and Copied data

6. Security & Redaction

Accreditation

The Directory has been accredited to hold CESG IL3 (government protected marking scheme – (RESTRICTED level) data and has been Penetration Tested to that level.

Redaction Engine

The Directory includes an integral Redaction Engine mechanism to limit the visibility of directory data. This enables a different view of the data to be presented to users from different agencies.

The most common use of this is to provide a restricted view of an agency's data to users from outside that agency (i.e.: 'guest' view of users) whilst internal agency users see a full view of the data.

Redacted views can be defined on any set of fields within the directory structure. The fields to be redacted are set by the administrator of a team or agency.

Commonly redacted field are personal information within Staff records, where a team or agency may want to limit the ability of a guest user to see staff details such as name, email address, role, grade etc.

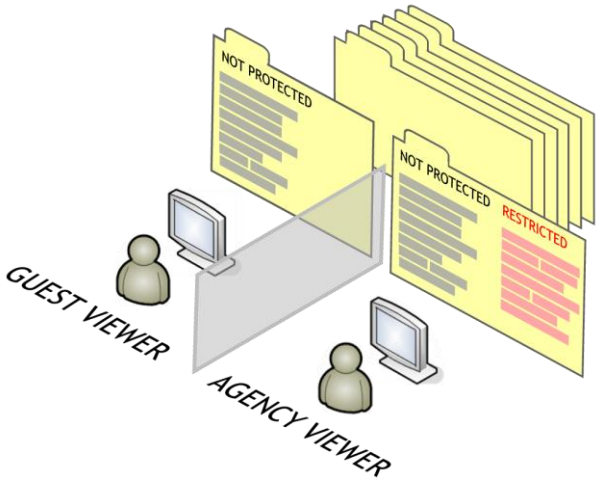

Figure 7 – Different Viewing Permissions for personal data

The most common use of this is to present a more limited view of personal or team information to 'guest' viewers accessing the directory over the GSi. Figure 7 above illustrates the concept.

Redaction Policy Management

The views available are set by the author of the information defining different viewing restrictions for different user contexts. The figure below illustrates the level of granularity available. The figure below illustrates how permissions can be set and the granularity available :

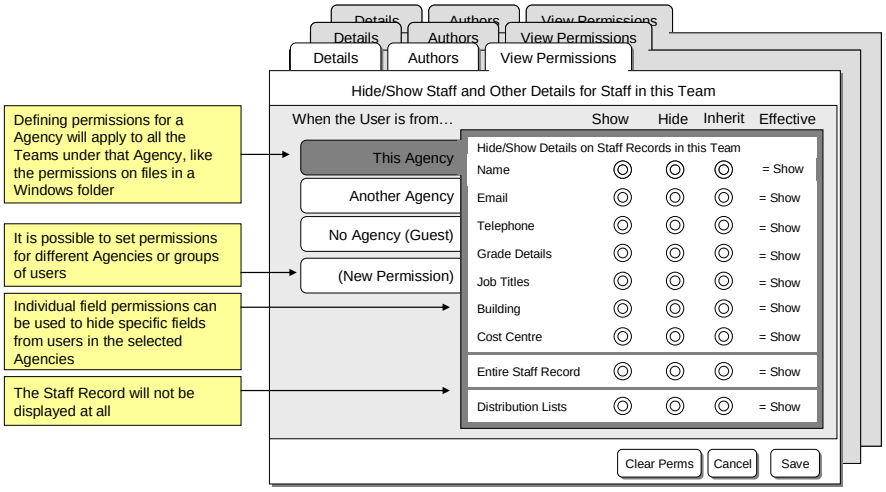

Figure 8 - Setting Viewing Permissions for fields in the Directory

User Authentication and Single Sign On

This architecture allows the Directory to act as the 'user-facing' component in an integrated user identity and authentication management system.

The Directory can be used as a SAML2.0 Secure Token Server for integration with a Federated Identity Management service (such as Microsoft's ADFS2 or Sun's Open SSO Identity Management service) or with applications requiring a Claims Based Authentication identity management service. As a part of the Directory package this service has been security accredited to IL2 and IL3 level user management.

Managing Different Data Sources

Support for the federated enterprise is focused on the ability to manage directory information according to the needs of the different business units comprising the enterprise. Some units will wish to 'master' their directories on the Directory, whilst others may master their data on their own systems and simply provide copies of content for the Directory.

Of these, some may prefer to handle updates outside of the Directory whilst others may wish to use the Directory updating mechanism to route updates to their host systems.

For the systems mastering their data on the Directory, there will be a need to integrate parts of that data with other directory masters, e.g.: staff email addresses should be mastered on the email global address list, staff grade information will be mastered on the HR system, staff network account information should be mastered through the network active directory service.

The Directory Integration Service is based on the Microsoft BizTalk integration server. This provides for the development of sets of business rules to support the arbitrarily complex and diverse relationships needed to combine separate directory data entities into a single Directory view of the enterprise directory.

7. Physical Architecture

Guidelines for operating the Directory are as follows:

Recommended Configuration, per user population	< 20,000	20,000 - 40,000	40,000 - 150,000
Hardware			
Database Server - 'Wintel' muliple CPUs (e.g.: HP Proliant DL370 equivalent)	1	2	2
Recommended # CPUs	2	2	4
Recommended Minimum RAM	4GB	8GB	16GB
Recommended Minimum Disk	100GB	100GB	200GB
Application Server - 'Wintel' twin CPU server, 4 GB RAM, (e.g.: Proliant DL370 equivalent)	1	2	4
Recommended Minimum RAM	4GB	4GB	8GB
Note for up to 20,000 users all the above functions can be combined on a single server.			
Operating System Software			
Database Servers Windows Advanced Server 2003 (Clustered, 64-bit preferred) SQL Server 2005/2008/2008R2			
Application Servers Windows Advanced Server 2003/2008/2008R2 (Clustered, 64 bit preferred) .NET 4.0 Framework			

Notes:

The performance guidelines above have been based on more than 3 years of running a high-usage site where 100,000 staff have been successfully served by two twin-CPU servers – one a database server the other the application server.

Web logs show that at this site, server configuration comfortably processes bursts of up to 10 requests per second without noticeable delivery degradation.

Scalability and resilience is achieved through standard Microsoft load balancing and clustering.

The application makes extensive use of 'object caching' to ensure that the server CPU power is used efficiently. Almost every object delivered to the user is cached in XML format, thus making extensive savings on server CPUs, especially the database servers.

The Directory can be operated on a secure network, such as that proposed for Home Office's Group Web Space (GWS) and offers interfaces to support read write access to other network resources. These have been organised so that the Directory will be capable of being accredited compliant to GSi standards. The interfaces on the Directory support protocol and business rule based workflow orchestration for:

- FTP and related file transfers for import export
- SMTP for system alerts
- LDAP and Web services client/server
- SAML2.0 Secure Token Service (STS) server
- File transformation, Importation and conversion of legacy Directory data, via FTP and/or Web services
- 'Active' (real-time updateable) links to common master repositories via LDAP and/or Web services. These include the HR ERP system, the email global address list, and network the active directory service,
- Semi-active (message based) links to partner system databases, via SMTP and/or Web services
- User and team author level updates to the data, via application web forms

8. Capability Summary – The User's Perspective

Capability –User's Perspective
General Usability
Users use the Directory most often to search for people, with standard searches based on name or any other personal details if known and free form text searches if not
Searching is easy to simple to learn with all searches based on a Google like syntax
Searching also supports advanced needs - more advanced searching is available at 'field' or 'free text', and 'synonyms' levels as users' experience advances
Searching is possible by field, free text or combinations, with input via quick, query-by-example or Boolean user input
The scope of any search can be limited by 'agency' (i.e.: limited to parts of the organisational tree)
As with Google, the user's most recent search activities are memorised between sessions
Searching - Field based
All entities and their can be searched, (i.e.: people, teams, and offices, etc.) by any field relevant to that entity
Where relevant, searches can use pick lists for selecting from sets of objects, (teams, grades, offices etc.)
Jumps are provided in the results list (i.e.: find a person then jump to their team and/or office)
Searching - Free text
Free text search on all content in the directory
Words and phrases can be combined using Google-like search syntax
Option to distinguish context of word searches, (e.g.: free text search for word 'Church' but exclude staff names)
Option to define and include a thesaurus in the search, to ensure free text searches don't miss acronyms, synonyms and other defined aliases
Free text search results are listed by order of 'relevance', based on Google-style search algorithms
Searching - Results
Results from person searches are retuned as for Google searches in pages and can be displayed in different order
Most columns in the results lists can be clicked on to jump to the dial of that item, e.g.: person detail, office detail, etc.
Results can be returned as in different formats, (e.g.: Excel, Email, cut'n'paste lists, Word labels etc.) for including in other office applications.

Organisation Tree Structure
The organisation is represented as a 'tree' , with the main organisation and it partners represented as top level 'folders' in the tree The tree is presented in the users' browsers using sophisticated caching and dynamic display logic to allow quick access around even the biggest tree structure The display of any organisational page (i.e.: a 'team') is synchronised to the display of the tree – a user can always see 'where they are in the tree' Users can set the scope of the tree, selecting between predefined 'agencies' to include or exclude as the root of the tree they see All users' tree and search details are held in browser 'cookies' to give inter-session consistency
Updating
Staff can easily change ('self-serve') their own data, working within predefined business rules set by the configuration, (e.g.: room, tel., job title can be changed but not name, email or grade) Sets of 'team authors' can be defined with rights to edit team content Team authors can be limited to authoring parts of the organisation tree. The entire authoring process is designed to be distributable, thus authors can be defined with different levels of administrator rights, allowing them to administer sets of authors within their organisational scope Team editing is web form based, but adds extensive user-editing features to make the process as easy to use. This includes find, replace, undo/redo, saving as draft, validating content, etc. Updates are applied to database in real-time All updates are fully audited Business rules can be defined on the updating process so that, where required, the Directory can act as a central update integration point, routing update requests to other systems using Web services XML messages
Special Needs Support
The user interface has been designed to support sight impaired users supported, e.g.: via JAWS for text to speech translation and has been Accessibility Assessed to AA standard. All user screens allow users to vary screen readability, e.g.: font size

9. Capability Summary – The Customer's Perspective

Capability - System Manager's View
Agency Support
Ability to view multiple agency's directories
Agencies able to be loaded through diverse ETLs
Able to manage agencies as master or secondary sources of data
Data Structures
Database holds Person, Team and Office details. Data attributes can be extended and are variable over different agencies.
Data managed internally as relational database with referential integrity
Virtually unlimited number of data items possible. The system has been successfully tested for performance with 10m staff, ½m teams, hundreds of agencies, thousands of offices
Unlimited number of organisational teams possible
Update Management
Updates come from three sources; user self-service, team authors, and data imports
Users use self-service to update their personal profiles; the fields they can change is pre-set on a per-agency basis.
Authors update teams (and people within those teams). They can be managed centrally or hierarchically
Authors can be assigned to many teams and classified with different levels of trust
Most changes can be undone.
All changes are fully audited and logged
Aged data items are flagged as such to their viewers as being less reliable; owners are alerted to this via email to encourage them to maintain them
Update parameters, (e.g.: what fields to lock, aging procedures, etc.) are managed on a per-agency basis in the Business Rule section of the directory.
Security Management
An inbuilt redaction engine allows information to be redacted depending on the viewer's agency, making the directory ideal for sharing with partners (who may not need a detailed view of the organisation's data). Such rules are managed in the Business Rules engine.
All directory software components have been developed for HMG IL3 (RESTRICTED) level operation and have been fully penetration tested and accredited to that level.
Additional security, limiting the user's ability to export large record sets, reduces the aggregated and associative risks of holding large numbers of personnel records.
The Directory can be fully integrated with any corporate Identify and Access Management system (i.e.: to provide Single Sign On)
All access and searches of the system are fully audited and logged

10. Capability Summary – The System Integrator's Perspective

Capability - From the System Provider's View
Platform
Totally web browser based service able to be run and 100% managed suppliers network domain
Web browser client designed to work with web browsers secured to government standards, and does not require ActiveX or Java applet plug-ins
Web based application server able to run in discrete network as managed service
Application based on standard Microsoft Server 2000/3 and SQL Server 2000/5 platform
Reliability
Server resilience utilises standard Microsoft Server clustering for automatic server failover and network load balancing
Clustering supported at 3 levels: IP switch, application and database
Minimal resources needed from client workstations (browser only needed)
Scalability, Performance & Volumetrics
Track record of use in the largest of government departments.
Server clustering to support increased usage
Extensive object caching technology maximises server performance and minimises network traffic
Highly efficient application with unlimited number of users supportable, evidence shows >50,000 users possible on single Intel server, twin CPU, 2GB server (e.g.: DWP)
Integration
Supports XML interfaces for integration with portal and other service consumers
Supports integration with other government directories, (Email Address Lists, HR systems, etc.), both triggered and scheduled
Support multiple interfaces for Web services, LDAP, SAML2.0 STS, FTP, and SMTP to other systems
Has standard ETL routines for integrating partner and external directories
Links to integration to business rules and workflow engine for linking to other systems
Support
All support services centralised
Replacement of legacy system, with difficult to support obsolete components
Can be deployed to clients as a managed service, simplifying SLA relationships

Daemon Directory Systems Ltd.
Gaston House
East Bergholt
Colchester
Essex
CO7 6SD
☎ 01206 299288
✉ admin@daemonlabs.co.uk
💻 www.daemonlabs.co.uk