



Overtone Software

Overview of **Overtone AbilitySuite[®]** **Email Content Management & Archiving**

Architecture and Operation:

Author: Ross Mitchell BE (hons),
NZCE – Chief Technology Officer,
Ability Software
January, 2007

Overtone Software[™]

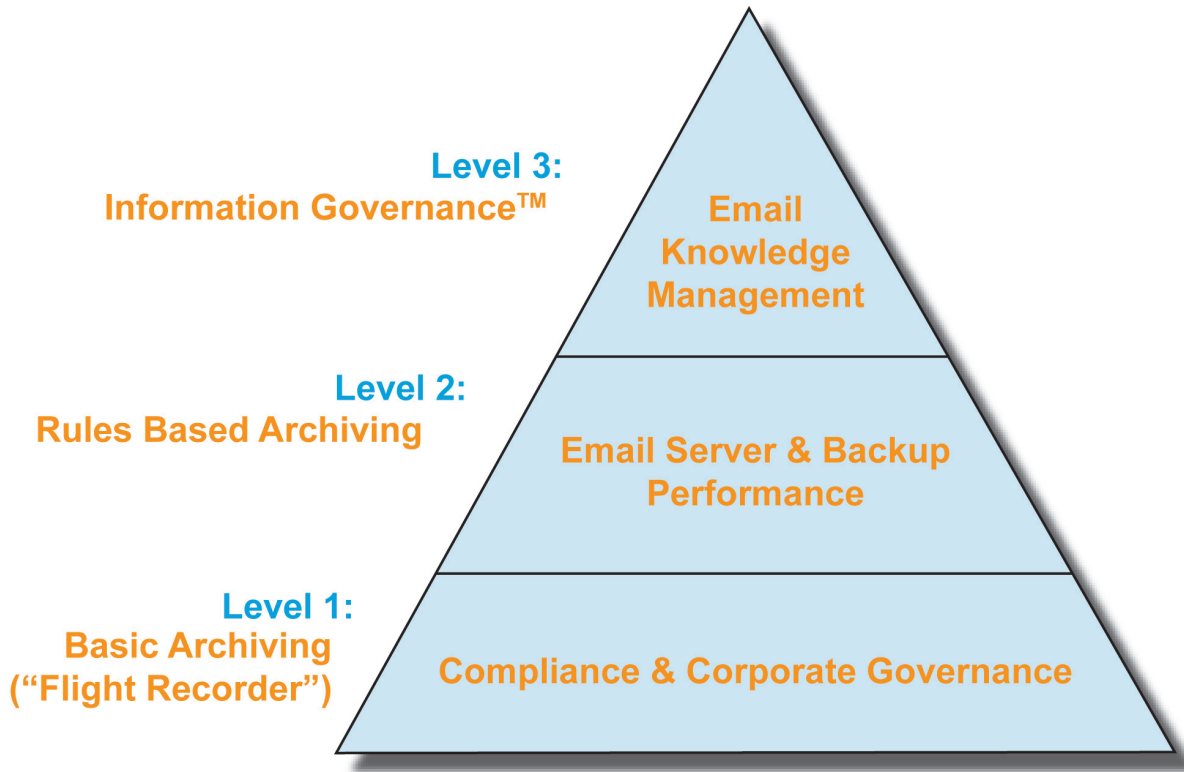
3 Bethesda Metro Center
Bethesda, Maryland 20814
Phone: 301-280-5902
E-mail: info@overtonesoftware.com

contents

Email Content Management & Archiving	3
What are the key drivers for Email Content Management & Archiving (ECM&A)	3
Reason One: Corporate Governance & Compliance	3
Reason Two: Mailbox Management & Backup Performance	3
Reason Three: The High Value of Information Contained in Emails	3
Overview of How <i>AbilitySuite</i> ® addresses these key requirements	4
Requirement One: Basic Email Archiving (“Flight Recorder”)	4
Requirement Two: Rules Based Archiving	4
Requirement Three: Email Knowledge Management	4
<i>AbilitySuite</i> ® Architecture	6
General	6
<i>AbilitySuite</i> ® MailStore™	6
Central Search Portal (CSP)	7
Category Instantiation Database	7
Configuration Manager	7
Rules Based Archive Manager	8
Other Key components of the <i>AbilitySuite</i> ® Solution	9
When Mail is Archived	10
How the Basic Architecture is exercised to Implement the Core Functionality	10
Level One: Flight Recorder	10
Level Two: Rules Based Archiving	11
Level Three: Email Content Management	12
<i>AbilitySuite</i> ® MailStore™ Architecture	13
Rolling CSP Architecture	19
Future Development	20
DB2 CSP Option	20
Removal of attachment	20
DXL XML data format	20
Configurable Trivoli Storage Manager (TSM) migration functionality	20
One-to-Many Relationship between CSP’s and Open Archives	20
One-to-Many Relationship between Configuration Manager and Main CSP	20

EMAIL CONTENT MANAGEMENT AND ARCHIVING

What are the key drivers for Email Content Management & Archiving (ECM&A)



Overtone *AbilitySuite*® for Email Content Management & Archiving

There are three key reasons organization wish to implement email content management and archiving systems these can be broadly defined as follows.

Reason One: Corporate Governance & Compliance

To ensure your company can meet strict legal and industry regulations it is often essential to have a record of all email messages. That is, automatically capture all inbound, outbound and internal email messages and store them safely outside the email files of users.

Reason Two: Mailbox Management & Backup Performance

Individual mail files are growing exponentially and mail servers are struggling under the load. IT departments are enforcing 'mail file size quotas', which are encouraging users to delete potentially important email messages. Backups are taking longer and longer. Transaction logging for incremental backups on Lotus Notes mail files create heavy server loading so often organization backup the whole mail file thus backup windows are large and tape usage, inefficient.

Reason Three: The High Value of Information Contained in Emails

Email use today has extended beyond simple messages and notes. Internally organizations use email as a default workflow tool to transport high value intellectual property between users. Email is used in HR negotiations and for agreements and resolutions between colleagues. Perhaps more importantly, email is also used to replace traditional external communications channels. Often

resolutions, commitments and agreements are made with external parties on behalf of organizations using email. Organizations should and can utilize this vast 'Knowledge Base' by empowering users to categorize, set security and centrally store emails so that the organization has access to this important information asset.

Overview of How AbilitySuite® addresses these key requirements

Requirement One: Basic Email Archiving (“Flight Recorder”)

Compliance can often be aided by securing copies of all email messages in a central and secure database.

At its most basic level, the *AbilitySuite*® for Email Content Management & Archiving (ECM&A) supports a 'flight recorder' like capability to automatically have copies of all inbound, outbound and internal email messages (including attachments) secured into the central *AbilitySuite*® MailStore™.

Requirement Two: Rules Based Archiving

AbilitySuite® for Email Content Management & Archiving supports a flexible scheme of rules for fully automatic archiving of messages from personal mail files into the central *AbilitySuite*® MailStore™.

Rules can be set - for example - to archive messages older than XX days, replacing them with a 'stub' in the user's mail file. Such 'stubbed' messages are substantially smaller than the original message and possible attachments.

'Stubbed' messages appear just like 'normal' messages in the inbox or folders within a mail file. When a user opens such a message, the original is instantly retrieved from the archive providing the user with a consistent experience. It is as if the message had never left the mail file.

Clearly, large numbers of 'stubbed' messages would still make mail files grow in size continuously - albeit much slower. Further rules can be configured to eventually remove 'stubs' from mail files - for example - after a further number of days, essentially creating an 'infinite mailbox'.

It's not hard to see how this process significantly reduces and effectively limits the size of hundreds and thousands of mail files without user intervention.

For existing large mail files *AbilitySuite*® for Email Content Management & Archiving comes with an administrators tool that can slowly migrate historical messages out of email files and replace them with stubs or delete them from the users email file. In this way you can automatically and incrementally migrate existing data into the *AbilitySuite*® archive.

AbilitySuite® can also detect duplicate emails in the repository. Duplicates are detected when an email has the same message ID. By using a single instance of the email in the central archive between users further net space savings can be achieved.

Finally, dependent on users unique hierarchical storage management requirements, messages can be migrated out of the *AbilitySuite*® MailStore™ using storage management solutions like IBM's Tivoli Storage Manager (TSM) onto offline media.

Requirement Three: Email Knowledge Management

Categorization, Meta Data and Corporate Taxonomy

While one can imagine searching for an email based on the standard meta data, which is inherent in any email message such as date sent, sent by, to whom and subject, using *AbilitySuite*® this

can be extended to include meta data, which represents lines of enquiry, which are unique to an organization. These custom meta data are typically values, such as project number, claim number, customer reference, committee name, product reference etc., dependent on the nature of the organization's business.

To easily and instantly locate email messages based on these custom meta data it is essential that messages are categorized using consistent terms which are unambiguous and standardized across the organization or department. Essentially, the categorization of email messages needs to be consistent with an organization's unique corporate taxonomy. For such a system to work effectively values for the categorization need to be supplied to users via a meta data management system capable of connecting to and brokering trusted values to each meta data field. Further, this needs to be implemented in a way that is simple and logical to users.

AbilitySuite® delivers and integrates such a meta data management system to allow for the tagging of emails as they are sent, composed, closed or any combination thereof.

Further, *AbilitySuite*® allows for a number of time saving capabilities to make this process simple to use such as:

1. User, group or pervasive email types with mandatory and option meta data fields.
2. Personal and pervasive pre-populated email types so that once a profile is created it can be re-used easily.
3. Auto-categorization of replies so that once an email "thread" is created all subsequent emails in that 'conversation' default to the categorization of the original.
4. Association of meta data with folders so that once an email is dragged into the folder it will be categorized according to the profile.

Security and Access Control is Paramount

If an organization is to open up its email archive to all users, security and access control must be a fundamental aspect of the archiving system. Because of the strong integration of *AbilitySuite*® with the Lotus Notes environment it is able to leverage the industrial strength security and authentication mechanisms inherent within the environment.

Security and access must be granular and the set at the message level. There is no degree of granularity above this that will work. *AbilitySuite*® uses Lotus Notes reader fields on every message to determine readership. Reader fields are set by defaults set-up in the definition of message types. These are used to extend readership to people or groups beyond the sender and valid recipients of the email. In this way, once a user is authenticated into Lotus Notes their access to messages in the archive is also set and if they search the archive via *AbilitySuite*® Central Search Portal (CSP), they will only see what they are permitted to see. In this way the archive can be made available to all users as a 'Knowledge Base'.

A fully integrated Central Search Portal

The *AbilitySuite*® Central Search Portal (CSP) allows the exposure of parameterized searches based on the configured custom meta data as well as standard meta data. No programming is required to close this loop. Once a meta data category is set-up and available through mail types to users it also appears on the search screen of the CSP for the construction of parameterized searches.



AbilitySuite® Architecture

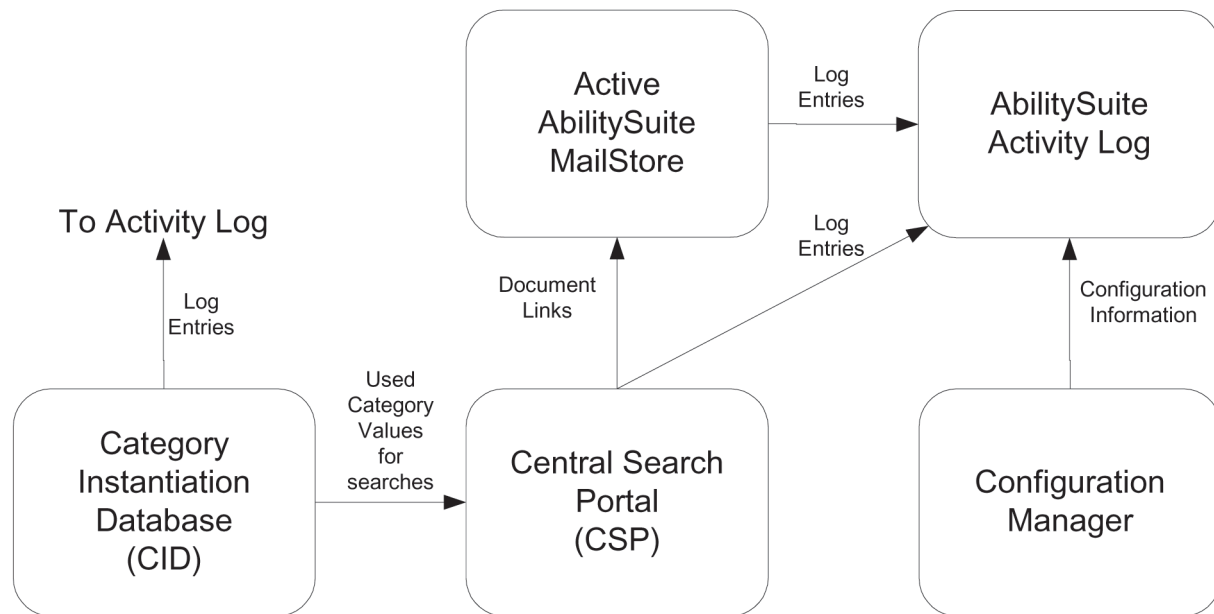


Fig 2 - Basic AbilitySuite® Architecture

General

At the most basic level the AbilitySuite® ECM&A solution consists of the following five databases:

AbilitySuite® MailStore™

The AbilitySuite® MailStore™ is in fact made up of three key components: The AbilitySuite® MailStore™ Repository(s), the AbilitySuite® MailStore™ Processor and the AbilitySuite® MailStore™ Interface. The AbilitySuite® MailStore™ Repositories themselves are simply the databases, which contain the email content. In themselves they contain very little in the way of sophistication. All code for the execution of rules for single instantiation, determining which repository an email should physically reside in and the resolution and correct handling of security (including BCC) is carried out by the AbilitySuite® MailStore™ Processor module. Subsequent access to emails is controlled by the AbilitySuite® MailStore™ Interface module. This separation of code and data (which is not always an architectural attribute of Lotus Notes development) makes maintenance of the system for upgrades and fixes a simple proposition. At any one time there may be several active “open” MailStore™ Repositories for each instantiation of the AbilitySuite® solution. Each open AbilitySuite® MailStore™ Repository may have rules governing the treatment of emails within them, that is, with regard to retention and storage management. On receipt of email for archiving the AbilitySuite® MailStore™ Processor evaluates rules based on sender and recipients as well as classification that determine, which open AbilitySuite® MailStore™ Repository the email should be physically stored in. Each AbilitySuite® MailStore™ repository has different rules for retention, storage management, roll-over and integration. This architecture provides a great degree of flexibility in the management of emails of different types or involving different people and is seamless to the end user.

Roll-over from ‘open’ to ‘closed’ occurs when the rule for the size or duration of an open archive are met. When roll-over occurs the new open archive assumes all the attributes and rules of the previous archive.

Finally, all systems, which archive email messages (Server Journal – Level One, Rules based

Archiving – Level Two, Mail Client – Level Three) to *AbilitySuite*® send emails to be archived to the MailStore(s)TM Processor by way of normal email routing. Thus, disconnected systems do not need to have a connection to the system in order for archiving to operate correctly.

Central Search Portal

The Central Search Portal (CSP) is a Lotus Notes database, which is designed to contain light-weight index cards for each email contained within active or closed *AbilitySuite*® MailStoreTM Repositories. The CSP is directly analogous to an index card system in a library. Index cards contain the standard and custom meta data for each email as well as reader fields and an auto launch document link to the content within the respective active or a closed *AbilitySuite*® MailStoreTM Repository. The CSP thus provides one holistic view across all *AbilitySuite*® MailStoreTM Repositories regardless of whether they are open or closed or what rules apply to them for retention and storage management.

Category Instantiation Database

The category Instantiation database (CID) contains category values for each category that has been used to categorize emails. The CID is the source of the lists used in the drop down boxes in the CSP, for each category. The reason the CID exists is that while category values may have been retired by the systems to which the Category Manager connects for classifications of emails, historic emails may exist in the archive using those retired category values. Thus, the CID contains a single instance of each category value that has ever been used to classify an email.

Configuration Manager

The Configuration Manager is a Lotus Notes database containing key configuration information for the operation of the *AbilitySuite*® ECM&A solution as follows:

1. General Configuration and Profiling
Documents for the general configuration of the system reside in the Configuration Manager. Additionally, settings for agent running options, roll-over of archives, migration of emails, rules management etc. are contained in the Configuration Manager.
2. Meta Data Manager / Broker
The Meta Data Manager / Broker is used for Level Three functionality. This system provides the functionality for the set-up of categories to be used for the classification of emails. Categories may be configured so that they source their values from Lotus Notes views, CSV file imports, direct entry into a managed list or through ODBC query. In the case of ODBC query client machines must have the ODBC drivers installed.
3. Mail Type Manger
The Mail Type Manager is where an administrator may set-up “mail types” for Level Three functionality. Mail types determine mandatory and optional fields as well as default security. Mail types may then be made available to any entity in the Lotus Notes Name and Address Book (NAB). In this way different mail types may be made available and limited to different groups or individuals.
4. Preference Profile Manager
The Preference Profile Manager is used to control the delivery of *AbilitySuite*® capabilities to users. Once again, preference profiles may be delivered to any entity in the organization through NAB entities. Preference profiles may be hierarchical so that



both horizontal and vertical groupings based on the organizations hierarchy may be supported. That is, if a person is the member of the HR group (vertical), the management group (horizontal) and the executive group (horizontal) and preference profiles are set-up for each of these then a ranking of each preference profile is used to determine which should apply to the user.

5. User Interface Behaviour Manager

AbilitySuite[®] provides the capability to perform custom, programmed actions on *AbilitySuite*[®] events in the lifecycle of an email. These are called code extensions and categorization actions. In addition, a set of configurable user data may be used similarly to “global variables” to determine and control the action of the categorization actions. Finally, a facility to create XSL style sheets exists where categorization actions are used; along with user data (and general categorization information) to generate content in the email. An example of this kind of functionality is the generation of unique signatures and disclaimers based on the categorization applied to an email.

6. Rules Based Archive Manager.

The Rules Based Archive Manager (RBAM) contains the agents, which test mail contained in individual user mail files against the rules, which apply to them, for Level Two functionality. The RBAM is configured by the Configuration Manager to run ½ hourly, hourly or daily to execute the graceful migration of mail from user mail files into the archive. The RBAM maintains a record of where it got to the last time it ran and takes this up at the next invocation. The agent, which runs in the RBAM, may stub or delete emails from user mail files according to the rules, which apply to them.

Individual users’ rules are maintained in the Configuration Manager in the form of user preference profiles. These are pushed into the preference profiles contained in each users mail file each time they open their mail file, if connected to the server. The RBAM interrogates the rules contained in the users mail file when it runs. The reason for this is that some preference profiles allow users to change the default rules, which apply to them, and thus each users rule may be unique (at the discretion of the administrator when the preference profile is set-up). Finally, an administrator can reset all user changes to rules and preference profiles if desired.

7. *AbilitySuite*[®] MailStore[™] Manager

The *AbilitySuite*[®] MailStore[™] Manager is where each new *AbilitySuite*[®] MailStore[™] is registered and the rules for their management (both retention, storage management) are set-up and maintained. Additionally, the *AbilitySuite*[®] MailStore[™] Manager allows the administrator to set-up scripts for the integration to and or export of emails to other systems based on rules.

Finally, if desired hierarchical storage management is configured in the *AbilitySuite*[®] MailStore[™] Manager for each registered repository.

Activity Log Database

The Activity Log Database is, as the name suggests, a database, which contains logs of the actions carried out by the various agents and systems, which operate to deliver *AbilitySuite*[®] functionality.

Other Key components of the AbilitySuite® Solution

Level Three Email Attachment

When Level Three functionality is enabled the categories, thread ID, readership and other key information, which is additional to the email, is formatted into a structured text file which is attached to the email. When any *AbilitySuite*® component of the *AbilitySuite*® solution receives an email with this attachment it:

- Knows that Level Three categorization and security extension has been applied to the email.
- May interrogate the attachment to determine the attributes associated with the email.
- May add replies to this email into the thread of the original and use the categorization of the original for default values in the reply.
- May create an index card, or concatenate attributes to an existing index card, based on message ID, in the CSP based on the information contained in the attachment.

AbilitySuite® attachments are recognised on a Lotus Notes mail domain basis. That is, attachments from outside the Lotus Notes mail domain will not be acknowledged by the local *AbilitySuite*® instantiation.

User Preference Profiles

User preference profiles are the vehicle by which an administrator determines how the *AbilitySuite*® ECM&A solution will operate for each user. Preference profiles may be delivered to any entity in the organization through NAB entries. Preference profiles may be hierarchical so that both horizontal and vertical groupings based on the organization's hierarchy may be supported. That is, if a person is the member of the HR group (vertical), the management group (horizontal) and the executive group (horizontal) and preference profiles are set-up for each of these then a ranking of each preference profile is used to determine, which should apply to the user.

AbilitySuite® ECM&A Mail Client

AbilitySuite® can be deployed for Level One journaling functionality and Level Two storage management functionality **without replacing the standard Lotus Notes client**. This deployment is termed "Zero Touch".

In order to access the advanced Level Three knowledge management features the *AbilitySuite*® mail client is used.

The *AbilitySuite*® ECM&A mail client is a modified Lotus Notes email template and is used to deliver advanced *AbilitySuite*® functionality to users. Thus, roll out of advanced *AbilitySuite*® functionality can be carried out centrally using standard Lotus Notes administration actions.

The *AbilitySuite*® mail template may be incorporated with existing modifications to the mail template although some care should be taken to ensure that customizations on similar areas of functionality are not introduced. It is possible to add customizations to the *AbilitySuite*® ECM&A template or to add the *AbilitySuite*® design elements to an existing modified template.

Currently, the advanced *AbilitySuite*® functionality is only available in the Lotus Notes client, however, deployment to the web client is in beta development phase.

As stated Level One and Level Two functionality can be supported using the standard browser interface to the Lotus Notes mail file.



When Mail is Archived

All systems, which archive email to the *AbilitySuite*® solution use normal mail routing to send mail to the currently active *AbilitySuite*® MailStore™ Processor, which is a Lotus Notes mail-in database.

When mail arrives at the *AbilitySuite*® MailStore™ Processor an “on mail receipt” agent runs, which carries out several actions. The key functionality of this agent, however, may be summarized by the following flow chart.

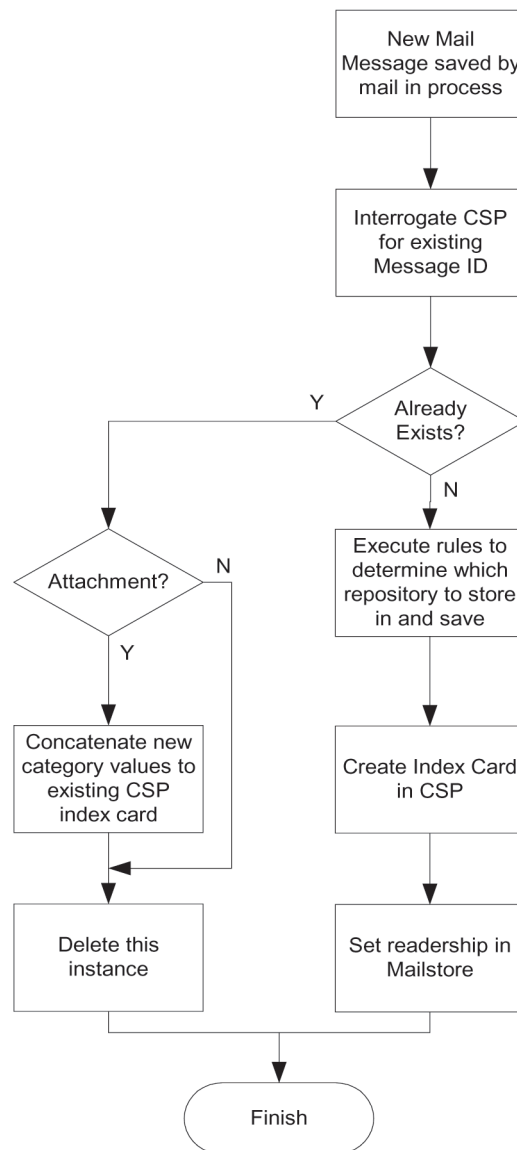


Fig 3 - On Mail Receipt Agent Flow Chart

How the Basic Architecture is exercised to Implement the Core Functionality

Level One: Flight Recorder

Level One “Flight Recorder” functionality is implemented by the *AbilitySuite*® ECM&A solution in the following way. Journalling is enabled on the R6 or above Lotus Domino email servers. In the Configuration Manager, Level One capability is enabled and appropriate *AbilitySuite*® MailStore™

templates, path and roll-over policy are defined. A new active archive is created automatically and the Interface and Processor need to be registered in the Configuration Manager database. It will then register the Processor in the NAB as a mail-in database. The journaling entry for servers is set to the mail-in database. When any active archive repository (*AbilitySuite*® MailStore™) is closed and rolled over a new active MailStore™ Repository is created (automatically by the Configuration Manager) and assumes all the rules and attributes of the closed repository.

When mail arrives at the *AbilitySuite*® MailStore™ Processor the “on mail receipt” agent runs and performs the actions outlined in the previous section. If no valid *AbilitySuite*® attachment is present the email has its security set (using reader fields) as the sender and any valid recipients for which valid internal Lotus Notes names can be resolved. In addition, special discovery roles managed by the administrator in the CSP and MailStore™ are added to the reader fields and thus will have access to search for and retrieve emails.

Clearly emails, which are recorded in this way, are only searchable using the standard meta data (To From CC, BCC and dates) or full text searching of the repositories themselves. If, however, a user later categorizes the email using Level Three functionality this categorization will be applied to the single instance of the email already recorded in the repository. Similarly, if the rules based archive agent later activates a rule to stub or delete the email from a users’ email file, the email will already exist in the archive due to the journalising action.

Level Two: Rules Based Archiving

AbilitySuite® is capable of both manual and automatic rules based archiving. In the case of manual archiving the administrator centrally defines whose emails to archive and what the date range of emails to archive is, then sets this running as a batch process. The best practice section of the *AbilitySuite*® Administrators Guide advocates this approach as the best in order to carry out a controlled migration of historic email, in the first instance, after deployment of *AbilitySuite*®. Following this, the administrator may choose to control the migration of subsequent historic mail or set automatic rules for archiving for different groups of people based on NAB groups.

Note: When the *AbilitySuite*® is deployed in “Zero Touch” client mode only central manual archiving by the administrator is possible.

The Rules Based Archive Manager (RBAM) is part of the Configuration Manager, which contains the agents, which test mail contained in individual user mail files against the rules, which apply to them, for Level Two functionality. The RBAM is configured by the Configuration Manager to run ½ hourly, hourly or daily to execute the graceful migration of mail from user mail files into the archive. The RBAM maintains a record of where it got to the last time it ran and takes this up at the next invocation. The agent, which runs in the RBAM, may stub or delete emails from user mail files according to the rules, which apply to them.

Individual users’ or groups’ rules are maintained in the Configuration Manager in the form of user preference profiles. These are pushed into the preference profiles contained in each users mail file each time they open their mail file if connected to the server. The RBAM interrogates the rules contained in the users mail file when it runs. The reason for this is that some preference profiles allow users to change the default rules, which apply to them, and thus each users rule may be unique (at the discretion of the administrator when the preference profile is set-up).

The rules based archiving agent works in a two part way. In the first instance it works through a view in each user mail file, which is keyed on creation date and a “processed” flag.

If the email being examined activates a rule then the agent interrogates the CSP to determine if the email



exists in the *AbilitySuite*® MailStore™. If the email does exist then the rule action is completed, i.e. to stub the email using an auto-launch document link to the copy in the MailStore™ or deletes the email from the users mail file. In the case of the former the 'processed' flag is set. If the email does not exist the agent sends a copy of the email to the *AbilitySuite*® MailStore™ and moves to the next email. On the next pass of the agent the email will be found in the *AbilitySuite*® MailStore™ and the action completed.

Finally, the RBA agent may be distributed between various servers in a distributed environment so that the load is also distributed. Documents are set-up in the Configuration Manager for each server, which define the mail files on which that server should operate. Thus, the burden of processing mail files may be shared out.

Level Three: Email Content Management

The *AbilitySuite*® ECM&A mail client can be configured, through preference profiles, to prompt users with the *AbilitySuite*® Level Three dialog:

1. As an email is composed
2. When it is sent
3. When the user exits a message
4. Combinations of the above
5. From action buttons in individual emails and views (e.g. inbox for received mail).

The *AbilitySuite*® Level Three dialog presents to the user a selection of mail types, which the user is permitted to see, including global pre-populated profiles and personal pre-populated profiles.

When one of these methods is used to categorize an email the *AbilitySuite*® ECM&A solution attaches to the email an instantiation of the Level Three email attachment. This attachment contains the categorization information, security and readership, which the user has determined for the email message as well as some automatically, calculated items such as a thread ID for the 'conversation'.

When emails are sent, with Level Three enabled, a BCC copy of the email is routed to the *AbilitySuite*® MailStore™ (This only happens if Level One Journalling is not enabled, else all email is captured anyway) Thus, once received the categorization applied is then recorded in the CSP as the email message is stored, as outlined the "on mail receipt flow chart" earlier in this document.

If replies to a categorized message are received from outside the organization with a valid attachment then the "On Mail Receipt" agent uses the attachment to default categorization of the new message to that contained in the attachment and sets the security and thread ID to that of the original.

Categorization

The category values used to categorize each message are accessed from the source determined by the Category Manager. Thus, trusted lists are used at all times. It is possible to configure each category to allow users to select "not available" so that a placeholder is inserted if a mandatory value, which the user requires, is not currently available. In this case the user can send the email as normal but will be prompted by the archive to fill in the category on a daily (or multiples of a day) basis until this has been done. A separate notification can be configured after a different period of time to a special role, e.g. your administrator.

In the main it is considered best practice to use Notes views with a sorted column so that type ahead is available in category selection.

Security

Mail types set default readership extension as configured by an administrator. Readership extension is functionality designed to allow other people to see the email if they search in the archive even if they were not originally sent the email. A mail type can be configured to allow “confidential” flags to be set by the user (in the categorization dialog), which removes the default readership extension. In addition users can extend readership to other users by using a specific readership extension property in the categorization dialog. Clearly, anyone sent the email will have access to it in the archive.

The effect of all of these security values is to set the contents of reader fields in both the *AbilitySuite*® MailStore™ and the CSP entry for each email.

BCC

The *AbilitySuite*® handles BCC correctly while still single instancing emails. That is,

- Users may only search for emails that they themselves have been BCC'd on.
- Special roles may search for all people BCC'd
- Users when viewing an email or search results will not see any other BCC recipients while special roles will.
- Users can see the full BCC list in any email they have sent.

In order for this to work correctly the server setting in the Domino server configuration settings document “ if each recipients address does not appear in any address header then add their address to the BCC list “ must be set to “Yes” so that BCC names are recorded on emails.

AbilitySuite® MailStore™ Architecture

One of the goals of the *AbilitySuite*® ECM&A solution is to provide a **completely Lotus Notes based option** to knowledge management, server load and backup performance-independent of the wider technology context within an organization. At the same time, however, if desired *AbilitySuite*® allows organizations to integrate email into wider Information Lifecycle Management (ILM), content management or storage management paradigms if desired. Further, it is a goal of *AbilitySuite*® to provide a superior email solution without this being effected by or tied to other decisions in relation to these solutions.

Many solutions for email archiving, which exist in the market are either part of a large vendors wider document management, content management or technology strategy and thus require a buy-in, by increments, to that strategy to implement an email archiving solution.

AbilitySuite® provides an open highly specialised and capable alternative. The only technology, which is prescribed by *AbilitySuite*® is that it is a system designed for a Lotus Notes email infrastructure. We believe that this is the correct association for an email solution to have, not an operating system, not a DMS, ILM or backup solution but the core system, whose capabilities *AbilitySuite*® seeks to extend, i.e. email.

AbilitySuite®, however, also provides techniques and out-of-the-box integration to other systems in order that it may be integrated with initiatives in other connected areas within the organization.

AbilitySuite®, in the first instance and in it's most simply constellation, focuses on resolving the email challenge within the Lotus Notes technology environment. In order to achieve this *AbilitySuite*® can deploy a totally Lotus Notes based solution using a rolling MailStore™ architecture.



Keeping it in Notes: AbilitySuite® Rolling Archive Technology

In the first instance, for those organizations who wish to deploy a completely Lotus Notes based solution, this is possible using what is called the *AbilitySuite®* Rolling Archive Technology.

The theoretical maximum size for a Lotus Notes .NSF file is 64GB. Common practice however, sets a usable limit at something less than this. In the main the maximum practical size of NSF's is determined by usability, view refresh times (user experience) and FT index limitations. It is also an accepted fact that the hardware and operating system platform on which any Domino system runs is a significant factor in these considerations.

To a large extent the key limitation, (user experience), is mitigated in the *AbilitySuite®* MailStore™ since access to documents are based almost exclusively by document links, which take a user directly to the document in questions, without the need to traverse views or interact with the database from a User Interface (UI) perspective. This allows *AbilitySuite®* MailStore™ Repositories to be run at sizes above what would normally be considered usable. However, we do still recommend values below the maximum supported size.

The *AbilitySuite®* MailStore™ Repositories may be configured to close themselves as an archive when a pre-set size is nearing and to open a new archive for mail to be archived to. Thus, an *AbilitySuite®* implementation will look like that, which is represented in the diagram below. Note: *AbilitySuite®* MailStore™ Repositories may also be configured to roll-over by date, which is very useful in smaller installations.

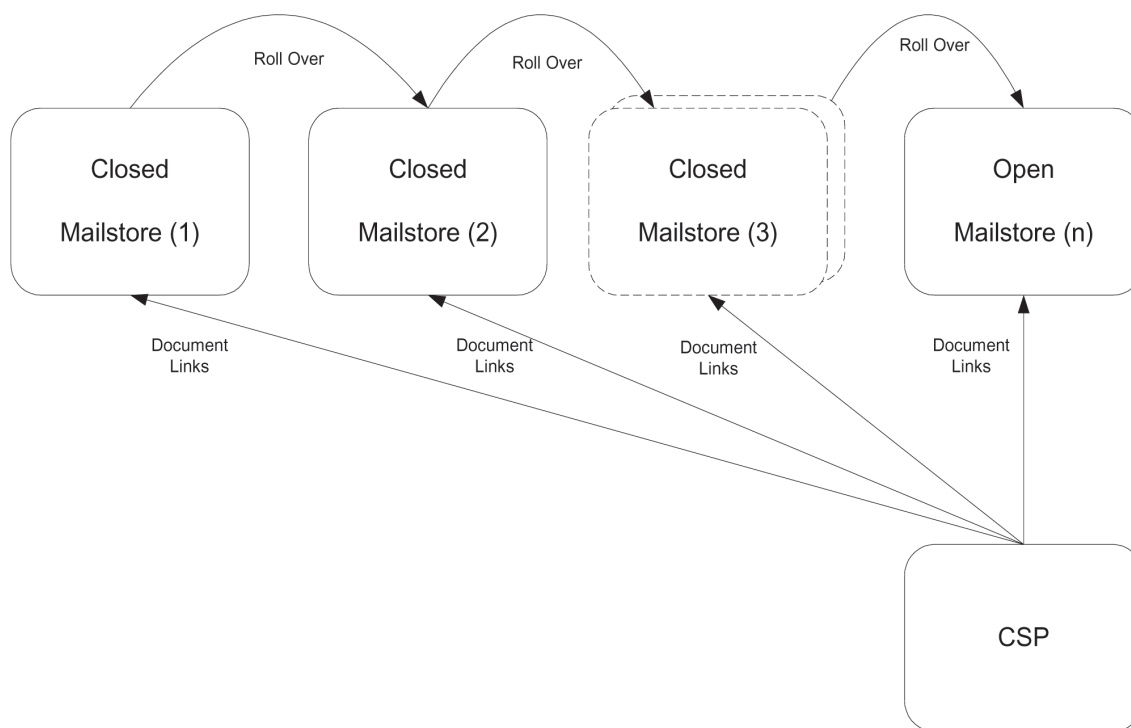


Fig 4 - MailStore™ Roll-over

The *AbilitySuite*® MailStore™ Repository architecture outlined has three important benefits to the Lotus Notes environment.

1. 100% Capture of Emails – but not in a separate proprietary system.
To provide a scalable solution completely within the Lotus Notes technology context. Essentially using *AbilitySuite*® the same amount of data is kept but it is organised better so that it's use and the management of it is greatly improved.
2. Optimise Backup Windows
One of the issues with Lotus Notes mail files is that they are unique individual files on the server. While from a number of perspectives this is a good architecture for a mail system, from a backup perspective it can prove problematic. Without transaction logging and a backup system closely aligned to Lotus Notes like Trivoli Storage Manager (TSM) then the entire mail file often needs to be backed up daily if so much as one change is made to the mail file. This clearly is likely to occur if the user is sending or receiving email daily. Also, transaction logging is known to place significant burden on servers. This means that many organizations may have a closely coupled backup system but have simply turned transaction logging off. Using the *AbilitySuite*® rolling archive strategy closed MailStore™ repositories may be backed up once and then excluded from further backups. This means that only the open archive (MailStore™) needs to be included in daily backups. If Level Two functionality is used on user mail files then significant reductions in user email files are achieved and the major proportion of email content is in closed archives. The combination of these two factors means that backup windows are significantly reduced, as is the requirement for high numbers of tapes. As one can see this significant effect on the environment is achieved very simply using *AbilitySuite*® ECM&A.
3. Improve Mail Server Performance
Every time a document is added to a Lotus Notes database a number of activities must take place, which place a burden on the Domino server. Examples of this are view refreshes and full text search index refreshes. The burden that these tasks place on servers goes up exponentially with the size of the Lotus Notes Database to which the document belongs. In addition, sheer size of mail files has an effect on server performance for the simple reason that when its opened the complete file is loaded into DB cache. Thus, the reduction of size of mail files introduces a significant performance saving for Domino messaging servers. With the *AbilitySuite*® rolling MailStore™ repository system only one archive is open at anyonetime, which requires server processing, while the, remainder and bulk of emails are in closed MailStore™ repositories, which are read only and thus do not require view refreshes and FTS index updates etc.

Moving it out of Notes: AbilitySuite® Integration Technology

Many organizations have Information Lifecycle Management architecture or a document / content management strategy and wish to have emails included in this. Alternatively, an organization may not have any of these but the sheer volume of email that they need or want to keep requires some form of storage management beyond keeping all the email on the mail server or in primary storage.

AbilitySuite® provides the capability to do this in an effective and importantly in a simple and open way.



Hierarchical Storage Management

Given the large amount of data, which is within emails systems and the predicted growth in this, many organization wish to further migrate email data onto low cost storage after a period of ageing.

The effect of this is that data may be further compressed and loaded onto slower but ultimately cheaper media. The supporting argument for this approach is that the value of the emails, the likely number of required accesses justifies the slower access to these emails.

AbilitySuite® does not prescribe a hierarchical storage management solution but rather provides the capability for utilising existing infrastructure if this is in place or comes with out-of-the-box integration to leading solutions, such as IBM Tivoli Storage Manager.

The *AbilitySuite*® MailStore™ repository allows for rules to be set-up for the ageing of the repository and can trigger events when criteria are met. For instance, a rule which activates after 100 days from closing the repository can be set-up to migrate the content (including attachments) of emails to tape library via, for example, Tivoli Storage Manager, where it is placed on one of a range of supported storage options, and compressed.

The way that this is implemented is that each rule triggers a sequence of events, in the case of the out-of-the-box integration with TSM the sequence is:

1. Render to content of the emails to XML (DXL i.e. Domino XML)
2. Write the XML as a compressed object to TSM either through the script or API interface (We support both)
3. Verify the object and receive a pointer to it in TSM
4. Delete the content from the body of the email message and replace it with the pointer.

When users subsequently wish to retrieve the content they are presented with a prompt that indicates the content is stored on tape and asks if they wish to retrieve it. If “proceed” is selected this will render the content back to the user in Lotus Notes.

The architecture used to implement this, however, is completely open. That is, you can write your own scripts, which are run in sequence when the rule activates if you wish to use a system other than Tivoli Storage Manager. The standard steps, for example the rendering of the email to XML, can be used by you to assist in this. Importantly, this gives you the choice to the supported out-of-the-box system or you can interface to whatever HSM or backup (tape library) system you already have.

Most hierarchical storage management and backup systems will have some form of scripting interface, which is very easy to use so the odds are that you can use the one you already have. Note also that once you have done this, those email archives need never be included in a backup again with the additional bonus that the information will persist longer than in a typical father/grandfather backup regime.

Finally, for the small, medium or distributed organization very simple solutions exist using *AbilitySuite*® that do not have to include tape libraries and HSM. The simplest of which is, of course, to leave it in Notes. However, if greater compression than that which you get with Notes is required or if you wanted to migrate the content to other storage options like XML rendering and file zipping can be utilized

Certified Non-Tamperable Storage

Often organizations have a requirement to migrate some emails to certified non-tamperable storage. *AbilitySuite*® has the capability to redirect emails to different archives based on rules such as who has been involved with the email and also what “type” of email it is.

This capability makes it possible to store important emails on certified storage. *AbilitySuite*® comes with out-of-the-box integration to IBM SSAM, such as IBM’s DR550 in the same way as it comes with out-of-the-box integration with TSM (see above). Indeed the mechanism is exactly the same for both as is the open architecture. Thus, the ability to migrate and retrieve emails from other certified storage devices such as EMC’s Centera and NetApp’s Snaplock are possible.

Export to Discovery Systems

In order to understand the philosophy behind the *AbilitySuite*® approach to discovery the reader should access the document “I know we need an Email Archive, but which one?”

In brief this white paper makes the following points:

1. The email archive is not the best place to carry out legal discovery, because:
 - a. Normally discovery requires an aggregation of information from many systems of which email is only one.
 - b. Email archives do not and should not support many of the capabilities specific to discovery such as sharing information with outside litigators etc.
 - c. When a document is registered as of interest in a discovery process and is loaded into a discovery system a legally defensible process for the management of the document complete with chain of custody must be instantiated. The email archive is not the place to do this.
2. What litigators or forensic discovery teams need is to be able to easily gain access to emails of interest and have these exported in an acceptable format to their discovery systems for processing.

AbilitySuite® thus provides the capability to export documents from a search result set into a separate Lotus Notes file for import into discovery systems. Most discovery systems will have the ability to support files in this format.

Emails can be supplied in Lotus Notes .NSF format or XML.

Integration With Document Management, CRM or Other Systems

As can be seen from the previous sections export of content from the *AbilitySuite*® system is a relatively easy task. This task can be automated so that export or registration of documents can be based on rules. However, *AbilitySuite*® best practice promotes a services oriented approach to the provision of email knowledge management and archiving capability to other systems like DMS, CRM and other core line of business solutions.

That is, we advocate that emails are most useful in an email system. Thus, we believe that when you access an email it should be via your mail client so that you are able to work with that email (forward, reply, print etc). If access to the emails are required by other systems then this access should be provided ‘as a service’ from the email system.



Using the technology for integration described already one can see that it is a simple task to integrate the *AbilitySuite*[®] to other systems in a general sense.

In the case of registration within a system like a DMS, which is not Lotus Notes based *AbilitySuite*[®] advocates to simply write the meta data (email header information) to the DMS record and insert a Domino URL, which would enable the launch of the document through the Notes client.

In this way the user can immediately use that content (reply, forward etc.) as if they had found it within their own mail file.

Alternatively, if desired, *AbilitySuite*[®] can simply write an XML version of the document to the DMS system.

If the system to which *AbilitySuite*[®] is to be integrated is a Lotus Notes based system (for example a Lotus Notes CRM system) then other very simple but powerful capabilities of the *AbilitySuite*[®] system can be exposed through a Lotus Notes based, configurable integration point.

For example, your CRM system can be used as a category source for emails (Level Three) and can also run contextual searches from a customer record directly into the email archive. For example “Show me all emails that have been sent to OR received from ACME and ourselves in the past three months that I am allowed to see”.

In order to implement this, simply copy the *AbilitySuite*[®] Software Development Kit (SDK) components (views, script, libraries, agents) to your CRM database. Then set-up the relevant field mapping between fields in your CRM customer record and the *AbilitySuite*[®] CSP using the Configuration Manager and call the search agents from within your CRM database,

As previously stated, given the number of other systems, which may wish to use the information in *AbilitySuite*[®], we advocate the document registration approach and services approach – we believe that this is the most user practical method even if the DMS or other system is itself viewed as an aggregator of content and itself can expose services for a holistic view of documents. We say this because ultimately if a user finds a email record it is likely that they will want to USE IT not just view it and the right place for that to happen is in the email system.

Regardless, the *AbilitySuite*[®] can do both and for both stands out as a specialist email management system for several reasons:

1. The first and very important feature, which will assist in this area, is the advanced capabilities, which *AbilitySuite*[®] has for knowledge management in the first place. *AbilitySuite*[®] has extremely advanced knowledge management and security features for the sharing of emails. Additionally, use of the Level Three capabilities enables an organization to separate the ‘wheat from the chaff’ when it comes to filtering the high-value emails for registrations or storage in a DMS or other system.
2. The second is that *AbilitySuite*[®] is implemented in Lotus Notes, so the format of documents can be either native Lotus Notes rendered through the Notes client or XML exported to the DMS or other system. While one can argue for formats that are not proprietary you will only hear that from vendors who’s format is not one of the de-facto document standards like Lotus Notes. That is, name one credible document management solution, which does not in some way support integration and / or rendering of Lotus Notes documents. Additionally, Lotus Notes has a native XML renderer into DXL (Domino XML), which is a highly understood and supported dialect of XML. *AbilitySuite*[®] is capable

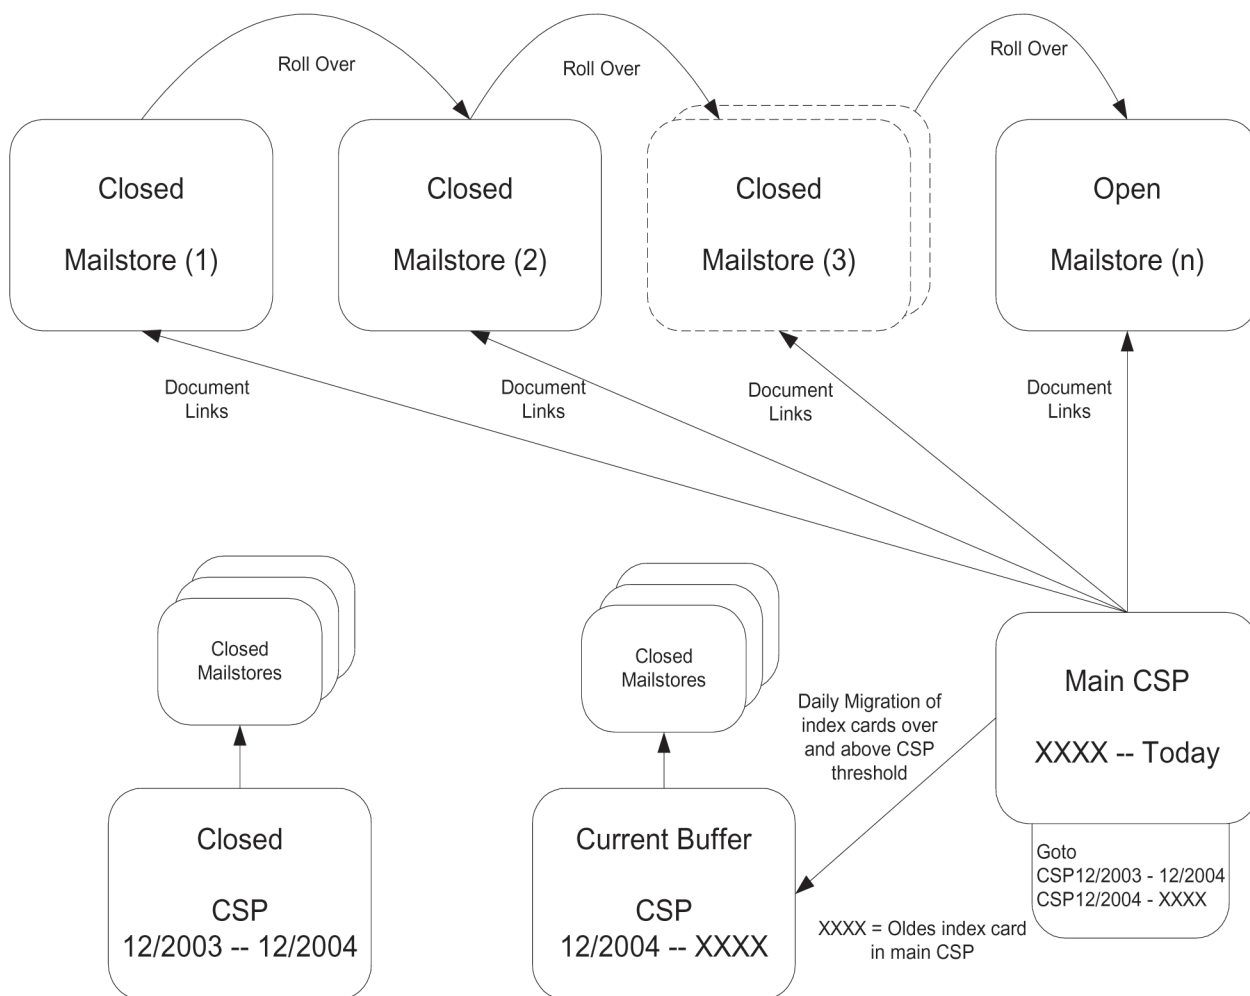
of supporting either standard and so access to documents within *AbilitySuite*® by other systems and even search engines is a given.

- The third reason is that the scripting capability which *AbilitySuite*® has enables the export or registration of documents within the DMS or other system with almost trivial ease. For example, it would be infinitely easy to write a Domino URL and meta data to a CRM, DMS or other type of system if this system has a command line interface or API for doing do.

Finally, if it desirable to treat different emails differently (and it almost certainly will be) rules can be written, which route each type of email to a different MailStore™ repository. While this is seamless to the user accessing emails using the search capability it has profound implications for the ease of management or integration of emails to other systems. Using the rules engine for the routing of emails into the separate MailStore™ Repositories allows *AbilitySuite*® to treat the emails within these in different ways from an integration, retention and information lifecycle management perspective.

Rolling CSP Architecture

Clearly, as a Lotus Notes database (Note: DB2 option in version 6) the CSP has an upper limit for the number of index cards, which the current open system may hold, and in a similar but not exactly the same way as the MailStore™, a rolling storage paradigm is maintained.



In the Configuration Manager an upper limit for the number of index cards, which the Main CSP will hold, may be set.

Each index card has an average size of between 1.5Kb and 2Kb. Thus, at the theoretical maximum any one CSP can contain is $64 \text{ GB} / 2\text{Kb} = 32$ Million documents. This theoretical maximum is, however, somewhat reduced by the number and complexity of views in a database. Additionally, the platform and memory available on a Domino server has a significant impact on maximum size. Thus a method of migrating older CSP index cards to closed CSPs is also supported in the *AbilitySuite*®. It should be noted (and is expanded on later in this document) that for large centralised systems a DB2 CSP option will be available in *AbilitySuite*® ECMA R6.

As stated a maximum number of documents may be set in the Configuration Manager for the CSP. What this figure should be depends on the infrastructure on which the *AbilitySuite*® system runs, however, once set the CSP will automatically migrate it's oldest documents into a buffer CSP. Once a buffer CSP is automatically created it is registered as an archive in the "Go To" options of the main CSP and may be searched similarly. The name the buffer CSP is given by default is the bounding dates of the index cards it contains. These dates are based on "date archived".

Once a buffer CSP reaches the upper threshold configured in the Category Manager it is closed and a new buffer is opened. The administrator may rename closed CSPs and they and their corresponding MailStore™ may be migrated off main storage, if desired.

Future Development

DB2 CSP Option

As stated previously, *AbilitySuite*® ECM&A offers a complete email archiving solution within the Lotus domino technology context.

Since Domino R7, Domino ships with DB2 and there is an option to "back-end" Lotus Notes databases with DB2 as the storage medium. This option removes the restrictions on Lotus Notes databases in terms of view indexes etc. since these all translate into SQL queries.

AbilitySuite® will offer opportunity to take advantage of this new feature to improve the options, which are available to *AbilitySuite*® installations. Key amongst these opportunities is the option of a DB2 back-ended CSP. The ability for the CSP to grow indefinitely offers an opportunity to service large centralised monolithic Lotus Notes topologies better since the main CSP will hold all index cards rather than a sliding window of these. This option will be available in version 6, which is scheduled for release in Q3 2007.

Removal of Attachment

The *AbilitySuite*® Level Three email attachment is used to carry meta data, security and thread information. Tests have shown a significant client performance increase where users have networked home drives and are using Level Three functionality when the attachment is removed and replaced with another paradigm. This is because creating and reading the attachment requires disk access and networks may slow this operation down. Additionally, the increased usage of the "in reply to" field in the SMTP standard by mainstream mail system has meant that reliance on the attachment to match message threads is becoming less important. In version 6 of *AbilitySuite*® ECM&A the attachment will be replaced with fields written directly to the Lotus Notes document for meta data and security values while the "in reply to" field will be used for message thread matching for emails received from outside the organization.

One-to-Many Relationship between Configuration Manager and Main CSP


In keeping with the philosophy of the *AbilitySuite*[®] it is intended that both distributed and centralised architectures for Lotus Notes systems are supported.

While the DB2 option for CSP back-end due in version 6 is targeted at centralised large systems it is still desirable to support distributed architectures, which use a single name and address book. *AbilitySuite*[®] ECM&A has a unique advantage over many email archiving systems because of its support of distributed architectures (which is a key feature of Lotus Notes). That advantage is simply because *AbilitySuite*[®] is implemented completely within the Lotus Domino environment. A large and new technology footprint is not required to implement this. If an organization is made up of a number of regional offices, departments or shops which are geographically separate then organizations often have two choices with traditional archive systems: (1) to implement a new technology footprint in each regional centre with all of the implications of this; or (2) to ship all email to a centralised archive creating large and growing network demands simply for the purpose of archiving email. That is, should all London email have to be moved to New York simply to meet archive requirements or should they be archived locally in London.

AbilitySuite[®] offers the option of a local archive WITHOUT a new technology being introduced. In this case a complete *AbilitySuite*[®] installation is implemented in each location and archiving is kept and managed locally.

The only limitation on this currently is when regional CSPs are deployed AND Level Three capabilities are used AND the organization uses a single Lotus Notes domain spanning all the regions. The limitation is that the Level Three attachment is designated with the Lotus Notes domain as part of its name. This is to prevent attachments from other organization using the *AbilitySuite*[®] being recognised by an organization's own *AbilitySuite*[®] installation. Thus, attachments from each regional centre in a single domain will be recognised as valid if received from another region. This in fact is acceptable if the category and mail type configuration is the same in each region but this is something, which must be maintained by administrators currently. In version 7 it is intended that multiple main CSPs will be able to be supported from a single replicated Configuration Manager.





Overtone Software™
3 Bethesda Metro Center
Bethesda, Maryland 20814
Phone: 301-280-5902
E-mail: info@overtonesoftware.com