




eMail Archiving and Records Compliance Foundation Planning Guide

Version 2.2.6

Windows



Note

Before using this information and the product it supports, read the information in [Appendix. Notices](#).

If you are not an IBM Business Partner, you are only authorized to use these materials for your internal use and for the sole purpose of evaluating products and services from IBM and the IBM Business Partner. You may not make the materials available to any third party.

Fourth Edition (March 2008)

Licensed Materials - Property of IBM

Copyright International Business Machines Corporation 2006, 2007, 2008. All rights reserved.

US Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP

Schedule Contract with IBM Corp.

Preface

How to Use Starting Point Assets

This set of solution enablement materials provides a starting point that helps you develop and deploy a solution for a particular type of customer problem. It provides a framework around which you can build similar solutions that meet your customer's unique needs.

Solution Overview

▶ *Helps you understand solution concepts, business value, and system architecture considerations*



Solution concepts
and business
value



Solution
architecture



Recommended
software,
hardware, tools

Solution Planning Guide

▶ *Helps you plan your customer engagement*



**This
Document**

✓ JSP programming
✓ Networking
✓ SQL

Skills required to
implement the
solution



Spreadsheet of time
estimates to plan services
engagement



Task descriptions
and best practices

Implementation Guide

▶ *Structured learning opportunity that shows you how to set up an instance of the solution and helps you learn important techniques that you can use to develop and deploy a custom solution*



Implementation
checklist



Instructions, tips,
techniques for
solution installation



Customization
information



Use cases that can be
used as part of demo
or education



Sample code, data,
scripts, configurations
to get started quickly

Demo Toolkit

▶ *Provides customizable demonstration assets to help you sell the solution*



Customizable presentation:
Solution benefits and overview



Video clips of the
solution



How-to document to
assist with using and
customizing the assets

Comments Welcome

Your comments are important to us. Please send your feedback about this and other IBM Solutions Builder starting points directly to us at sbuilder@us.ibm.com

Legend



Used to identify worksheets, sample code, or other reusable files that are provided with this specific solution.



Used to identify areas where there is more than one implementation option for the solution.



Used to identify tasks or places that require special attention.



Used to identify websites referenced in this solution.



Used to identify IBM Redbooks referenced in this solution.

What's New in Version 2.2.6 of this Solution Starting Point

In version 2.2.6 of *eMail Archiving and Records Compliance Foundation*, you will find:

- Version updates to product fixpacks
- Updated link to product fixpacks
- Additional instructions on installing fixpacks
- Updated instructions for creating user IDs and configuring products

Contributors

eMail Archiving and Records Compliance Foundation was produced by members of the IBM Solutions Builder Express Portfolio team with the assistance of the following specialists:

Business Partners

- **Jens Hildebrandt** holds a Degree of Public Administration at the University of Applied Sciences in Mayen, Germany. He has worked with IBM DB2 Content Manager (formerly ImagePlus) since 1989. He is a partner at Udit GmbH and an advanced Business Partner in Neuwied, Germany. He leads the Archiving and DMS department and is responsible for planning, installing, customizing and product integration of archiving systems. His experience includes IBM DB2 CommonStore for Lotus Domino and SAP, as well as deep programming experience with DB2 Content Manager OnDemand on iSeries and Multiplatforms.
- **Steven Muchow** holds an engineering degree from the Johns Hopkins University. He is an Information Lifecycle Management Specialist with CAS Severn, an IBM Premier Business Partner in Laurel, Maryland. Steve has over 14 years of systems development and integration experience. His areas of expertise include cross-platform systems development, enterprise content management, and integration of high volume workflow and document management solutions. Steve may be reached through email at smuchow@cassevern.com or by accessing www.cassevern.com.
- **Chris Tillman** holds a computer science degree from Towson University in Towson Maryland. Chris is an Information Lifecycle Management specialist with CAS Severn, an IBM Premier Business Partner in Laurel Maryland. He has led development teams in the Business Process Management (BPM) and the Information Lifecycle Management (ILM) areas for over six years, specializing in the IBM portfolio of products. Chris can be contacted at ctillman@cassevern.com or by accessing www.cassevern.com.

- **Chris Turner** is a Presales Technical Specialist for IBM VAD Meier Business Systems in Melbourne Australia. Chris has over 25 years experience in systems development and database support. His areas of expertise include database design and implementation, in-depth support and training for database administrators and managers in the multi-value database arena and service team management. Chris may be contacted at chris.turner@mbs.net.au. IBM Services

IBM Services

- **Peter Elford** is a Senior Content Management Specialist for IBM SWG in Sydney Australia. Peter has over 15 years experience in sales, implementation and support of content management systems. He has worked in these roles for a variety of vendors in the ECM market including FileNet, Interwoven and IBM. Peter may be contacted at pelford@au1.ibm.com.
- **Raymond Skoglund** is a Channel Manager for IBM SWG in Wellington New Zealand. Raymond has over 20 years experience in computer industry. He has worked direct sales and channel management roles for Sybase and IBM. Raymond may be contacted at raymonds@nz.ibm.com.
- **Dean Underwood** manages IBM's TotalStorage Hub and has more than 25 years of experience in sales, development and management of storage systems. He can be contacted at dcunder@us.ibm.com.
- **Glenn Korn** is a member of the Technical Support Marketing group for the IBM System Storage group where he specializes in SAN and NAS solutions. Glenn has been working with computer hardware and networks for over 15 years. He has extensive experience in IT support of both hardware and software. He holds a B.S. degree in Industrial Technology and is a Microsoft Certified Systems Engineer (MCSE) and Novell CNE.
- **Lijing Zhang** is a Content Management Specialist for IBM SWG in Dallas, TX. Lijing has over 7 years experience in developing, implementation and supporting of content management systems. Lijing may be contacted at lijing@us.ibm.com.

Prerequisites

There are no prerequisite Solution Starting Points.

However, this Solution Starting Point requires that the customer already has either a Domino or Microsoft Exchange e-mail system and either LDAP or Active Directory in their environment.

Related Solutions

The following additional solutions can be used to enhance this solution:

- *Collaborative Portal* - This Solution Starting Point illustrates a portal that centralizes access to critical information such as business applications, productivity tools, people, team workspaces, document repositories, and electronic meeting facilities.
- *How to Rapidly Assemble and Deploy Solutions* - This Technical Reference illustrates how to create, build, and update solution deployment accelerators created using the IBM Solution Assembly Toolkit
- *Managing User Accounts Inside Your IT Infrastructure* - This Solution Starting Point illustrates how to rapidly deploy and configure a user identity management solution into an existing environment.
- *Managing the Health and Availability of your IT Infrastructure* - This Solution Starting Point illustrates the creation and deployment of a monitoring solution for a mid-market company.
- *Secured External Access To Company Resources* - This Solution Starting Point illustrates how to secure access to internal company IT resources via both wired and wireless networks.

- *Providing Basic Services to Maintain Hardware Platform Health* - This Solution Starting Point enables hardware resellers and/or RSIs to provide IT implementation services and efficiently maintain the hardware health of multiple SMB customer networks.

Contents

eMail Archiving and Records Compliance Foundation Planning Guide i

Preface iii

How to Use Starting Point Assets iii

Legend iv

What's New in Version 2.2.6 of this Solution Starting Point iv

Contributors iv

Prerequisites v

Related Solutions v

Chapter 1. Understanding Required Skills 2

Chapter 2. Task Assumptions and Estimates 5

Chapter 3. Task Details 9

Requirement Discovery 9

 Define Solution Value 12

Design the Solution 16

Install and Configure Essential Software 27

Install Technical Assets 29

Customize Technical Assets for Production 31

 Enhance Compliance Administration Console 31

 Customize the Meta-data Configuration Manager 33

 Customize the SDA Wrapper (Two Machines) 37

Test and Deploy the Solution 38

Additional Services 39

Appendix. Notices 41

Trademarks 42

Chapter 1. Understanding Required Skills

To successfully implement this solution, you will need to establish a team of people with a variety of skills. The list below describes the skills your team will need in order to design, develop and deploy the solution. The time estimates suggested in the next section are based on the assumption that the personnel performing the work have at least most of these skills.

With these skills...	You can...
Solution architecture	Evaluate and analyze the customer's environment, applications, network and security requirements. Adjust the systems or application architecture and related infrastructure and evaluate solution options for customer.
Experience with Microsoft Windows 2003	Navigate the operating system; install, configure and launch applications; configure network settings; manage security and facilitate configuration changes in Active Directory
Experience with Microsoft Exchange email Administration	Install, configure, and deploy this solution using IBM DB2 CommonStore for Exchange Server
Experience with Lotus Notes email Administration	Install, configure, and deploy this solution for IBM DB2 CommonStore for Lotus Domino; manage security in the Domino Directory
Experience with IBM DB2 CommonStore for Lotus Domino	Evaluate, implement, and configure email and application archiving in a Lotus Domino email and application infrastructure.
Experience with IBM DB2 CommonStore for Exchange Server	Evaluate, implement, and configure email archiving in a Microsoft Exchange email infrastructure.
Experience with IBM DB2 Content Manager	Perform user administration and management of item types
Experience with IBM Tivoli Storage Manager	Evaluate, implement, and configure hierarchical storage management solutions to support the email archiving and general storage management needs.
DB2 DB Administration	Tune the DB2 Content Manager Library Server database in the first month of operation to optimize performance.
Blade Center Administration	Configure, update blade configurations on a blade center chassis

Depending on which scenario you choose and which technical assets your customer requires, your team may need the following additional skills to customize the solution. The scenarios are described in detail in *Systems Architecture* in the Overview Guide.

With these skills...	You can...
Experience with Lotus Domino Designer	Customize the solution to support customer compliance and records management by implementing the Meta-data Configuration Manager technical asset provided with this Solution Starting Point.
Knowledge of the IBM DB2 Content Manager API and experience with J2EE, Struts	Customize the solution to support customer compliance and records management by implementing the Compliance Administration Console technical asset provided with this Solution Starting Point.
Knowledge of IBM DB2 Content	Configure IBM DB2 Content Manager custom item types ¹ to support

With these skills...	You can...
Manager Item Types and security	enhanced compliance or records management environments. Customize the solution to support customer compliance and records management, proof-of-concept (POC) strategies, and system migration tasks. This supports the implementation and customization of the Compliance Administration Console, Meta-data Configuration Manager and enhancements tasks and customizations recommended for them. Many of these customizations exploit the Item Type Migration Model (see the <i>Customize the Solution</i> section in the <i>Implementation Guide</i> for more information)
Knowledge of IBM DB2 Content Manager Item Types and security	Install, configure, and deploy this solution for IBM DB2 CommonStore for Lotus Domino; manage security in the Domino Directory
Experience with Rational Web Developer for WebSphere Software	Develop, customize, and deploy the technology assets delivered with this Solution Starting Point that are built to leverage the IBM DB2 Content Manager API; extend Compliance Administration Console by adding enhancements required by the customer. Answer customer questions regarding basic and advanced content management related functions and the development of these functions in the environment.
WebSphere Administration and Application Development	Install, deploy, configure and customize solution components such as: <ul style="list-style-type: none"> • eClient for Content Manager (see Document Forwarding on page 33 for more information) • Compliance Administration Console (technical asset provided with this solution)

If you want to add an email search using IBM eMail Search for CommonStore, the following skills are needed:


With these skills...	You can...
Experience with IBM DB2 Net Search Extender Administration	Install, configure, maintain and tune IBM DB2 Net Search Extender (NSE). NSE works with IBM DB2 to provide text search capability to IBM DB2 Content Manager library server.
Experience with IBM WebSphere Application Server Administration	Install, configure, maintain and tune WebSphere Application Server. WebSphere Application Server application server is part of Content Manager Resource Manager and eMail Search.
Experience with IBM Lotus Domino Administration	Install, configure, maintain and tune IBM Lotus Domino server. IBM eMail Search for CommonStore (Domino) provides many optional features that requires installation and configuration of Lotus Domino.
Experience with IBM DB2 Information Integrator for Content Administration	Install IBM DB2 Information Integrator for Content. It is required by eMail Search.
Experience with IBM DB2 Records Manager Administration	Install, configure, maintain and tune IBM DB2 Records Manager. DB2 Records Manager is an optional component of eMail Search.
Experience with IBM eMail Search for CommonStore Administration	Install, configure, maintain and tune IBM eMail Search for CommonStore. eMail Search provides search capability against IBM DB2 CommonStore.

Table 3. Optional Skills - eMail Search

With these skills...	You can...
Microsoft Outlook	Install Microsoft Outlook, which is required for IBM eMail Search for CommonStore for Microsoft Exchange.

Note:

¹ Detailed instructions for the creation of attributes, item types and other models can be found in the *Creating a backend archive for CommonStore* section (p.33) of the [IBM DB2 CommonStore for Lotus Domino Administrator's and Programmer's Guide Version 8.3](#). The part number for the guide is SH12-6742-03.

 Consider using the [IBM Virtual Innovation Center](#), [IBM Redbooks](#), and [IBM Learning Services](#) as learning material to enhance the skills of your team.

Chapter 2. Task Assumptions and Estimates

This section presents a list of tasks that you may need to perform in planning, developing, deploying and maintaining your solution, along with estimated durations for each task.

Solution Assumptions

The following assumptions, or parameters, have been applied to the tasks for this solutions table. These assumptions should be revisited when you are defining the scope of your unique customer solution since, depending on your customer's actual needs, your implementation may vary from the one described here. You may therefore need to adjust tasks and/or estimates when developing your own project plan.

The assumptions for the tasks and time estimates described later in this Solution Starting Point include:

- Existence of an email infrastructure (Lotus Notes or Microsoft Exchange) in the customer's environment.
- A full email archiving implementation with a requirement to install IBM DB2 Content Manager and IBM Tivoli Storage Manager (Scenario A or B) services hours will need to be adjusted if IBM DB2 Content Manager is not included (Scenario C or D)
- The services hours estimates assume a low or no growth implementation of the solution domain in an architecture including an existing email and directory server and one new server dedicated to archiving storage. The architectural scenarios are described in detail in *Systems Architecture* in the *Overview Guide*. Additional guidance on sizing considerations may be found in [Assess Performance](#) and [Determine Physical Sizing](#) in this *Planning Guide*.
- You have the skills required to install and customize the solution and any technical assets implemented.

The range of services hours in [Table 4](#) table is based on the use of information and technical assets included with this Solution Starting Point by the business partner. The following ranges are provided:

The Solution Deployment Accelerator (**SDA**) range of services hours assumes:

- Automated installation of the basic email archiving solution software components using the Solution Deployment Accelerator (SDA). Configuration consists of manual steps.
- The customer requires email archiving capabilities only. There is no requirement for compliance.

The **LOW** range of services hours assumes:

- Manual installation of the basic email archiving solution software components using instructions provided in the *Implementation Guide* of this Solution Starting Point.
- The customer requires email archiving capabilities only. There is no requirement for compliance.

The **MEDIUM** range of services hours assumes:

- Manual installation of the basic email archiving solution software components using instructions provided in the *Implementation Guide* of this Solution Starting Point. If you choose to use the SDA, installation costs are reduced and the installation success rate is increased. Installation for solution architectures can be pre-tested and repeated in the customer's environment.
- A customer requirement for compliance exists. The technical assets (samples) provided with this Solution Starting Point will be used to implement a stepped migration to a future full records management implementation.

- The technical assets provided by this Solution Starting Point will be customized only enough to make them appropriate for the customer's production environment (e.g. interface (UI) presentation, error handling, internationalization, integration features, etc).

The **HIGH** range of services hours assumes:

- Manual installation of the basic email archiving solution software components using instructions provided in the *Implementation Guide* of this Solution Starting Point. If you choose to use the SDA, installation costs are reduced and the installation success rate is increased. Installation for solution architectures can be pre-tested and repeated in the customer's environment.
- A customer requirement for compliance exists. The technical assets (samples) provided with this Solution Starting Point will be used to implement a stepped migration to a future full records management implementation which may give the customer some level of compliance.
- The technical assets provided by this Solution Starting Point will be customized for the customer's production environment as well as enhanced with new features to satisfy additional customer compliance and records management needs.

Solution Task Estimates

The following table is provided for you to identify the tasks you plan to do. In some cases, task hour estimates will have significantly broad ranges. Unique considerations are explained in the detailed sub-sections.



This table is also provided in a spreadsheet for you to identify the tasks you plan to do and calculate project time estimates.

✓	Task	Range of Services Hours			
		SDA	Low	Med	High
	Requirement Discovery on page 9	24	24	32	40
	Customer Interview	24	24	32	40
	Define Solution Value on page 12				
	Design the Solution on page 16	52	52	84	108
	Define the Architecture	24	24	40	40
	Assess Performance				
	Determine Physical Sizing				
	Assess Return on Investment on page 25				
	Demonstrate Prototype on page 26	4	4	4	4
	Final Design and Customer Sign-off on page 27	24	24	40	64
	Define the Project Plan				
	Install and Configure Essential Software on page 27	32	46	58	78
	Install and configure IBM Tivoli Storage Manager	4	6	6	6
	Install and configure IBM DB2 Content Manager	18	28	32	48
	Install and configure IBM DB2 CommonStore	6	8	12	16
	Install and configure IBM eMail Search for CommonStore (eMS) on page 28	4	4	8	8
	Install Technical Assets on page 29	0	0	6	12
	Install Item Type Migration Model	0	0	6	12

✓	Task	Range of Services Hours			
		SDA	Low	Med	High
	Install Meta-data Configuration Manager				
	Install Compliance Administration Console				
	<u>Customize Technical Assets for Production</u> on page 31	0	0	72	96
	Customize Item Type Migration Model	0	0	24	32
	<u>Customize the Meta-data Configuration Manager</u> on page 33	0	0	24	32
	<u>Enhance Compliance Administration Console</u> on page 31	0	0	24	32
	<u>Customize the SDA Wrapper (Two Machines)</u> on page 37 ¹	0	0	0	0
	<u>Enhance Compliance Administration Console</u> on page 31	0	0	28	200
	<u>Logon and Custom Privilege</u> on page 31	0	0	8	24
	<u>Extended Search Capability</u> on page 32	0	0	4	24
	<u>Complex Data Models</u> on page 32	0	0	4	80
	<u>Delete Functionality</u>	0	0	0	16
	<u>Reclassify Record Attributes</u>	0	0	4	16
	<u>Viewing a Document in the Result List</u> on page 33	0	0	8	24
	<u>Document Forwarding</u> on page 33	0	0	0	16
	<u>Customize the Meta-data Configuration Manager</u> on page 33	0	0	16	152
	<u>Fileplan Folder Security</u>	0	0	16	24
	<u>Auto Classification</u>	0	0	0	16
	Automate Agents	0	0	0	32
	Automate Profile	0	0	0	16
	<u>Profile Type</u> on page 37	0	0	0	16
	<u>TO DO List</u> on page 37	0	0	0	16
	<u>Test and Deploy the Solution</u> on page 38	34	34	70	128
	<u>Validation</u> on page 38				
	<u>Integration Testing</u> on page 38	16	16	40	80
	<u>Certification or User Acceptance Testing</u> on page 38	6	6	10	20
	<u>Deployment</u>				
	<u>Setup Production Environment</u>	8	8	12	16
	<u>Educate the Customer</u>	4	4	8	12
	<u>Additional Services</u> on page 39 ²	0	0	0	0
	Total	142	156	366	782

Note:

1. This task has no hours associated with it in terms of production, but it will take some time to export the Solution Deployment Accelerator (SDA) image.
2. The *Additional Services* task does not have any hours associated with it; this for any ongoing maintenance or other customer-specific tasks.

For the medium and high services hours ranges, these estimates could be applied over time (6 - 12 months or more) as part of a stepped implementation plan. Some of these values are, by definition,

spread over the migration path between implementation of email archiving and the full compliance solution (see [Define Solution Value](#) on page 12 for details).

Adjusting the Task Estimates

Depending on your skills and experience, the task hour estimates presented may be too high or too low. You may want to approximate more realistic time estimates for your efforts based on whether you are a novice and new to this solution or skill area or are considered an expert. The scope of your solution will also influence the tasks and time estimates.

Chapter 3. Task Details

This section describes common tasks that are performed when designing, implementing and deploying this solution. This information has been greatly influenced by the business partners who have taken part in this project and over 50 business partners the SBE team has interviewed in preparation. Each of the technical assets (samples) and the customizations and enhancements listed for them as tasks are similar to those implemented by partners in real email archiving, compliance, and records management engagements.

Requirement Discovery

This task involves evaluating the email archiving solution requirements. You also need to have a good understanding of the customer's compliance and records management needs and what elements of the current infrastructure compliment such a strategy. Consider asking some of the questions below to obtain this information.

<i>Table 5. Questions to ask the Customer</i>	
Question	Why the question needs to be asked
Questions to gather functional requirements	
What are your email infrastructure & business process concerns in light of regulatory governance? What plans or initiatives are currently under consideration?	This allows you to understand the scope of the solution needed. Customizations and enhancements offered with this solution focus on securing the customer's compliance and records management future. This information allows you to leverage this solutions technical assets, customizations, and enhancements to define a clear functional requirement from definitions provided in this Solution Starting Point.
Are there resources in place or contracted to address: <ul style="list-style-type: none"> • Data recovery • Team to support data recovery infrastructures • Risk management in the area of data recovery contingency planning Do you have specific business related concerns?	This Solution Starting Point is written to allow the business partner to exploit an existing TSM customer base or introduce one where none exists. If TSM is installed you are leveraging/raising the value of the customer's investment and up-selling IBM DB2 Content Manager and IBM DB2 CommonStore.
Are you currently monitoring your infrastructure? Have you considered looking at employing technology to drive / self-manage business policies and needs?	Discussion points could include: security management, systems management, availability and disaster recovery, problem determination, integrated console services, workload and performance optimization, or automated storage management. These are all features that will be introduced or have a foundation laid down to support them when this solution architecture is implemented in the customer environment. These questions specifically target the Tivoli influence brought by IBM Tivoli Storage Manager and allow opportunity for identifying cross-sell in the Tivoli portfolio of products (e.g., Tivoli Identity Manager, Tivoli Access Manager).

Table 5. Questions to ask the Customer

Question	Why the question needs to be asked
Have you assessed infrastructure security risks, quantified the value of information assets and created plans to mitigate or eliminate high priority security exposures? Do you have specific concerns?	Error trapping and complex security are beyond the scope of this Solution Starting Point. This is an important topic for a compliance solution / strategy. This series of questions aims at the Item Type Migration Model and any of the customizations or enhancements for searching, viewing, or classifying documents.
(Assuming email archiving is or will be implemented) Have you considered leveraging the data your email archive represents to solve business compliance concerns?	For Microsoft Exchange or Lotus Notes email environments where email archiving is or will be implemented, this Solution Starting Point's technical assets may also be used to sell a compliance strategy.
Have you considered records management solutions to solve regulatory or business compliance issues?	If the customer has a need for records management solutions due to regulatory or business compliance concerns this Solution Starting Point's technical assets may also be used to provide a smooth migration from simple email archiving to a full compliance strategy that includes records management. This Solution Starting Point is designed so the customer may set their own watermark or level of tooling to deal with their unique set of compliance needs. They may grow the solution to just what they need, when they need it. They do not need to make a big commitment before they are ready.
What is your need for email archiving based upon?	Is there a regulatory requirement being mandated by Acts like HIPAA, Sarbanes-Oxley, etc. Is there a legal requirement to lock down data based upon legal proceedings? Is there an audit requirement for email? Is it just for mailbox management? If anything other than basic mailbox management, then eMail Search is most likely required.
Questions to gather technical requirements	
What email system does the customer have/need to use?	This solution provides value in situations where either Lotus Notes or Microsoft Exchange email systems are used. Customizations described in this solution provide an opportunity to: CommonStore for Lotus Domino may be sold as more than just an email archiving solution. For Microsoft Exchange environments that also leverage Lotus Notes for applications, partners may sell extended services to apply IBM DB2 CommonStore for archival of application data generated through Lotus business applications.
How many email users do you support and what is the average email size (special business use, for instance, large attachments)? What are the network throughputs between existing machines? Will the network support the required bandwidth if other, new systems must be introduced?	Perception and performance go hand in hand. For the users, this is a high touch application. Great care must be taken to how the business need drives data use habits among the users. See Assess Performance for insight on measurement of the throughput based on the number of users, type of traffic and frequency of

Table 5. Questions to ask the Customer

Question	Why the question needs to be asked
	traffic and peak usage requirements.
<p>How does the business need and data usage map to hierarchical storage management or the plan for it?</p> <p>How do you use your business data? What data do you use more frequently? Less frequently?</p>	<p>The business need and data usage characteristics map directly to the hierarchical storage management (HSM) solution. With this solution, you have the ability to highly optimize the automatic distribution of data for archiving and retrieval requirements driven by the business need. See Determine Physical Sizing for insight on how to size the storage requirement.</p> <p>Your customer may have no experience with a HSM solution; stress the proper use of its capability and the value it brings. High performance and reliability are not determined by just an efficient network. With this solution, the customer has the option to boost performance and reduce costs by putting low use data on low cost devices and optimizing the use of higher cost, higher speed alternatives.</p>
Questions to gather system requirements	
<p>What are the systems you have available to install the solution?</p>	<p>Determine if systems currently within the customer environment provide the minimum specifications required to install, customize and run the solution.</p>
<p>What operating system will you use for the solution?</p>	<p>This solution is available on Microsoft Windows. Although this solution includes IBM DB2 Content Manager, the Linux version of this product does not support the full-indexing features required to optimize searches.</p>
<p>What solution areas will this customer need to invest in?</p>	<p>Guidance is provided (see When Performance Matters on page 22 in this guide) to help the partner and their customer identify current and forecast future need. This is crucial to making decisions on investment in the infrastructure.</p>
<p>What infrastructure has the customer invested in?</p>	<p>This solution provides guidance assets for BladeCenter based infrastructures. This release provides an SDA that automates software installations for single and two machine architectures (see this solution's Quick Start Guide for information) for the solution available through participating channel distributors (Post-installation configuration to fit the solution to the customer requirement is still required).</p>
<p>Based on the current investment, plan, and solution requirement, what server platform makes sense?</p>	<p>This solution is certified for the BladeCenter technologies but in theory, other platforms should be fine. One suggestion for an entry level alternative could be two System x3655 servers with DS4200 storage extensions. See this URL for more on this server (http://www-03.ibm.com/systems/x/rack/x3655/index.html) and talk with your distributor about supporting this as a solution as an SDA (Solution Deployment</p>

Table 5. Questions to ask the Customer

Question	Why the question needs to be asked
	Accelerator).

Note:



The *SBE Assessment Worksheet* is provided with this Solution Starting Point. This questionnaire is designed to determine physical sizing, product selection, distill the core solution, and help point to opportunities for optional or accessory services (pattern customizations). For more information on sizing, see [Determine Physical Sizing](#) in the *Design the Solution* section of this Planning Guide.

Define Solution Value

eMail archiving is a great entry point to the customer because it is now a market tested, proven technology. Customers understand this is a must-have application. Although the customer also sees compliance as a major concern, they are not as comfortable with technologies that would address their compliance needs.

Products introduced to the customer environment by this email archiving solution also provide the foundation for more demanding content management applications. Implementing extended physical storage, storage management services, and technical solutions for implementing a compliance strategy is a natural extension of this investment.

The following sections will outline how the business partner can develop a *What's Next* vision with complimentary services strategies for their customers. This information will show you how to leverage the technical assets delivered with this Solution Starting Point to:

- **Reduce the sales cycle for a full compliance solution** - Business partners providing full compliance solutions report 12 to 24 month sales cycles. Customers require educational assistance to be successful in implementing a compliance strategy that often takes on the form of guidance, a more difficult thing to measure and to put a price on than educational units or services. Reduced sales cycles, designed to produce measurable results with their delivery of a portion of the compliance requirement, also integrate the educational or guidance value the customer needs. This design allows the business partner to finance the kind of solution and support model their customers must have to be successful.
- **Speed up ROI** - Historically, full compliance solutions are orchestrated across the entire company and target the end goal (full compliance through records management) and education all at once. This solution's technical assets will provide a way to 'step' implementation and education at a department level. This approach allows compliance concerns to be targeted and shows the business partner how to deliver results to their customers in the short term.
- **Reduce the customer's Cost of Ownership** - Business partners report that the high cost of full compliance solutions is due more to implementation costs related to services rather than to licensing costs. This solution's technical assets are designed to move some of the more expensive services (e.g. storage management, fileplan analytics and development, rules development, declare and classify records - usually data normalization, avoided using this solution's methods) and position them as migration and education based services for the customer. By concentrating the services in front of the full compliance implementation, a more intelligent budget model, stepped to a more realistic educational and adoption model, is achieved. In the end, the final step to full compliance (e.g., records management) is less painful and less expensive.

The sections that follow will reflect the values characterized in the bullets above. Each technical asset will be discussed in the context of how it affects delivery of these values. Also described is how to leverage

the customer investment in email archiving by providing a risk mitigating and cost saving migration plan to a compliance solution, and doing so in a way that fits the customer's needs.

Developing and Understanding a What's Next Vision

A *What's Next* vision is crucial to the customer as it provides a picture of what they should be concerned with next and how you, the business partner, can provide what it takes to address those concerns. Without a clear vision there is no value for the customer to repeat business after the sale of a simple email archiving solution. Email archiving can be just a starting point for your relationship with a customer.

Three key components must be identified and aligned with the customer's business needs:

- The Starting Point
- The End Goal
- The Migration Path

Simple eMail Archiving: A Starting Point for the Relationship

To gain the confidence of a customer and begin a lasting relationship, you should look for a way to show you understand the things they value. To win a customer looking for a dependable consulting and support organization or to supplement their own, you need to bring them a high value application at a low (or at least a reasonable) price, and show that you are competent to deliver that solution and the next.

eMail archiving and the assets delivered with this solution are that starting point and provide the following values:

- **A set of enterprise scalable products at a mid-market price.** This solution domain includes content management and storage management solutions at some of the lowest prices ever. The customer can leverage this low cost acquisition of enterprise scalable product, educate users and administrators in their use, and develop more advanced strategies before being forced to make a larger investment.
- **The skills and experience to perform a quality deployment, quickly.** The technical asset, Solution Deployment Accelerator (SDA), provided with this Solution Starting Point allows you to establish the base quickly. Your ability to repeat this process for a complex solution domain and make it look routine shows a level of competence. This asset comes with a complete set of instructions for learning how to modify the SDA to any customer's requirement.
- **The depth of knowledge in the solution domain required to tailor a solution.** This Solution Starting Point provides guidance on how to size and optimize the solution to get the best configuration and performance for the customer.

The End Goal: A Level of Compliance

If you want to grow the customer relationship and repeat business, you must know where the customer is going. A customer may know the general direction (e.g., compliance) but they usually have no idea on how far they need to travel to get there. This is value the business partner should bring, judging that distance and providing some clarity to the direction and decisions.

Business partners agree that any compliance solution is a long engagement and a very expensive one for both the customer and the business partner. This Solution Starting Point helps bring clarity by providing a system of small steps, focused on immediate compliance needs and the tooling required to achieve them, to educate the customer and involve them in the process of designing their own compliance solution and strategies early. This is accomplished by providing the partner a way to 'step' a customer to a compliance goal in the following ways:

- The Solution Starting Point provides a way to develop a department level approach to developing the fileplan for managing a record's life cycle that can be shared by the entire organization.

- Tools are provided that make it easy manage change through this process. This allows the immediate return on the investment of the customer by providing records management functionality through the design and change process in smaller (functional) increments.

The Migration Path: Learning While Getting Results

The task of making the determination that a full compliance solution is a customer's goal is dwarfed by the challenge of devising a plan on how to get there as there is a lot to accomplish. This Solution Starting Point recommends that walking, not running, is the best approach. Walking is the approach of using a 'stepped' implementation. This Solution Starting Point proposes a component strategy and these components should be balanced and complimentary to one another. The migration path should always contain the components described below.

Education for both the Business Partner and the Customer

In a complex up-sell situation like compliance, it is possible to learn everything about the requirement before starting. However, this is not advisable as it is easy to over-analyze. As discussed in [Allies in a New World](#), there are stakeholders to be considered and each will have their own needs and possibly a budget and agenda to go along with them. They will teach you what the requirement is and you will teach them how to support each other's compliance goals.

Customization and enhancement ideas identified for the technical assets delivered with this Solution Starting Point (see sections on installation and customization of technical assets beginning with [Install Technical Assets](#) on page 29) are designed to help the business partner identify the features the customer needs to facilitate their compliance requirement. Implementing the customizations and enhancements for the customer allows the customer to define their fileplan and compliance rules. The business partner should leverage their knowledge of the customer's entire organization to provide continuity across departments within the business where fileplan and rules must provide a cohesive company-wide strategy.

Services that make Sense for the Customer

Services that make sense are those that produce a return on the investment the customer makes. This Solution Starting Point delivers technical assets and describes customizations and enhancements that have been recommended by business partners as common requirements in a compliance engagement. These customizations and enhancements should be used as a reference when you gather the requirements from your customer.

The obvious services are those that implement the customer requirement of those customizations and enhancements that fulfill their compliance needs. There is much more to consider than simple features though. The educational and analytical processes behind developing a plan for compliance are an expensive and long process. Although this process is long, it can be made productive by combining the education of the customer and the business partner to a stepped implementation that produces real compliance capability as the plan for compliance is developed (e.g., fileplans, life cycle rules).

Customization and enhancement ideas for the technical assets include features that allow the business partner to provide services for managing the collection of fileplan and rules information. Some departments will be more savvy about compliance requirements than others. This solution can be delivered as either an administrative function where the customer has a more active role or one where the business partner leads the fileplan and rules development process. Delivery could be made more flexible by splitting responsibility. Some departments may take more of a lead role than others.

Note:

Departments who have a deeper knowledge of compliance and the company's liabilities for non-compliance should be targeted first in the applications tooling process. It is important to include the customer in these decisions to avoid liability issues. See [Allies in a New World](#) in the *Design the Solution* section of this Planning Guide for more detailed information.

Methods for Reducing the Cost of Reaching the End Goal

The technical assets (and the customizations and enhancements ideas that will compliment them) enable the business partner to seek and find the correct level of compliance for a specific customer. In this solution, there is nothing implicit about the end goal. The technical assets are designed to make the end goal for the customer's compliance requirement a watermark that can be what the customer needs when they need it.

Demands change and not just within the context of a legal or business compliance requirement. The stepped implementation, using the technical assets recommended in this Solution Starting Point, allows the customer to seek their own watermark. This allows the compliance requirement to grow as the customer's level of understanding improves, driving a demand for improvement beginning with the customer.

Compliance is a journey made much easier if the customer is driving the sale. Repeat business with the business partner should be born of a need the customer perceives. This perception comes from conclusions drawn from the leadership provided by the business partner using this solution.

Justifying the Lotus Up-sell

If you are a Lotus Domino professional, you may think that many of the functions offered in this solution that leverage the IBM DB2 CommonStore, IBM DB2 Content Manager, and IBM Tivoli Storage Manager products look very similar to features already available in Lotus Domino. Therefore, you might be tempted to avoid the skill investment of learning the products recommended for this Solution Starting Point if you feel you can cover the customer's business requirement with the products you are already familiar with.

The information below identifies the key characteristics that may help to contrast the two solution options and define a threshold for adoption of one over the other, or a combination of both:

- While Lotus Domino's native archive capability manages the data within a Domino messaging environment, IBM DB2 CommonStore allows it to be managed in an external repository. This is important if configured information must be kept for years and maintained on cheaper storage devices. IBM DB2 CommonStore can also maintain email content (or other documents) in their original Lotus format or convert to a more portable format for you.
- Domino creates separate NSF files for each user, so no single, central repository and no cross-mailbox search is available. IBM DB2 CommonStore can retain, manage, and search across large volumes of archived email content, a feature that is important for discovery searches.
- Domino archives can reside only on a Domino server's disks or local user's DASD. Lotus Domino is not designed to manage Lotus Notes content in other formats. For this reason, data is not accessible directly from a hierarchical storage management system and cannot take advantage of the benefits of hierarchical storage management (migrating data from disk to jukebox to tape for example) that this solution domain provides through IBM DB2 Content Manager and IBM Tivoli Storage Manager.
- For long-term retention, some regulations are requiring that WORM (Write Once Read Many) devices to be used. Native Domino archives cannot manage this type of storage device. Device management features of the IBM Tivoli Storage Manager (which is recommended by this Solution Starting Point) can support these devices. All the scenarios in this Solution Starting Point support this option.
- The IBM DB2 CommonStore environment enables archive data compression as a feature of IBM Tivoli Storage Manager. This improves ROI for long-term retention and volume, because storage requirements are less.
- Starting with Domino R6, IBM DB2 CommonStore and Domino's native archive rules can be synchronized. As a result, Domino administrators may still define Domino native archive rules, which are then passed to IBM DB2 CommonStore, which manages the rest. This feature allows Lotus Domino administrators to manage their archives in the same, familiar way they are accustomed to but rely on IBM DB2 CommonStore to manage the archive seamlessly.

This solution is a natural extension for Lotus business partners who have been exploring ways to extend their skills and product and services strategies to other IBM products. This provides a great opportunity for business partners to leverage their existing Lotus customer install base.

They can introduce enterprise-class content management environments with a delivery model similar to the rapid development and deployment models Lotus Notes developers are accustomed to using.

Design the Solution

Designing the solution involves identifying and analyzing the customer requirements, designing and architecting the solution, and scheduling and documenting all of the activities needed to create the solution. Proper planning in the first step of the design process ensures that your final design is user friendly, flexible and extensible, and it eliminates expensive and unnecessary redesigns.

Define the Architecture

It is important to have a high-level view of the interaction between the many processes that enable an email archiving environment. This view allows you to consider the system architecture from the context of the function and the business need driving usage requirements. Before deciding on the best approach for your situation, quickly review the diagrams and information below to determine the best physical architecture for your purpose.

The default architecture described and implemented in this solution (as shown in the diagrams below) addresses the very smallest of production implementations where no or low growth is expected in the environment. The information below addresses the process and product integration considerations required for a higher volume environment.

Note:

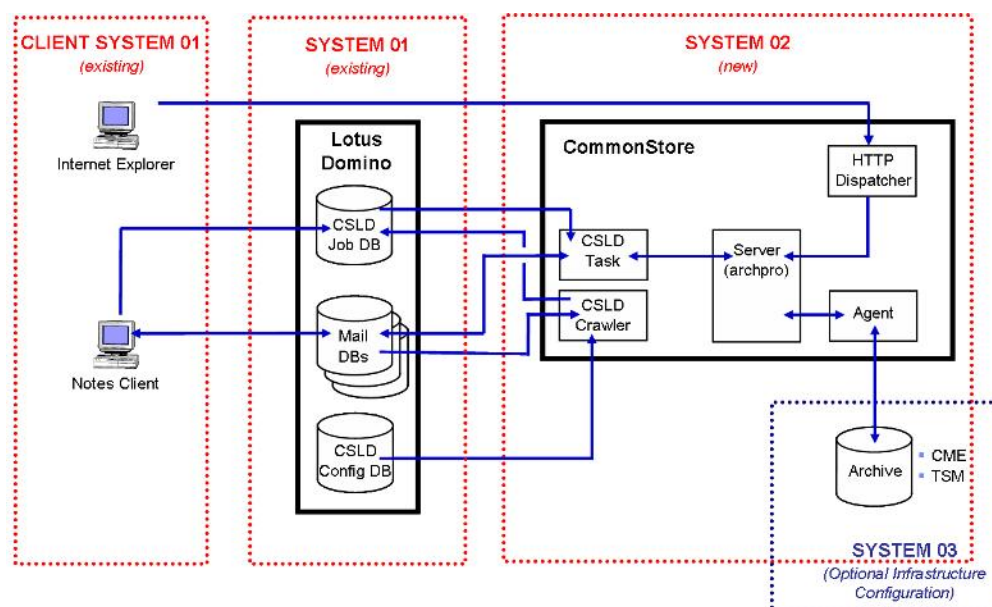
See *Systems Architecture* in the *Overview Guide* for a detailed description and pictures describing architectural scenarios.

The IBM DB2 CommonStore for Lotus Domino (CSLD) Architecture and Process Interaction

In the world of Lotus Domino administration there are some general, best practice rules that are typically followed. Lotus administrators are very selective in what they add to their Domino Server installations. As a rule they do not install other products (within reason) onto a company-wide production Domino server.

Administrators are especially sensitive to adding anything to an email server that might affect performance and affect many users. The figure below illustrates how the IBM DB2 CommonStore solution components are distributed as implemented in this Solution Starting Point for scenario A.

Figure 1. Suggested Solution Component Architecture for IBM DB2 CommonStore for Lotus Domino



Unlike in the diagram above, administrators recommend installing the IBM DB2 CommonStore for Lotus Domino task, crawler and server on a separate machine (along with a Lotus Notes Client). If the volume of user archival and retrieval traffic is low, this could be the IBM DB2 Content Manager server (no optional SYSTEM 03 included). A less powerful machine and/or higher traffic could overwhelm this design. If this occurs, adding an additional machine for the archive (Content Manager) would be required (where optional SYSTEM 03 is included).

In addition, implementing IBM Tivoli Storage Manager Extended Edition for IBM DB2 Content Manager Enterprise Edition is typically done in situations where a high volume of data is expected. It is common to see lots of content driving a solid hierarchical storage management requirement and/or a need to use TSM in other places. Implementing TSM on the same server as IBM DB2 Content Manager does not facilitate growth and should be limited to situations where there are low data volumes, low or no growth and where TSM is needed only for use by IBM DB2 Content Manager in a single use application.

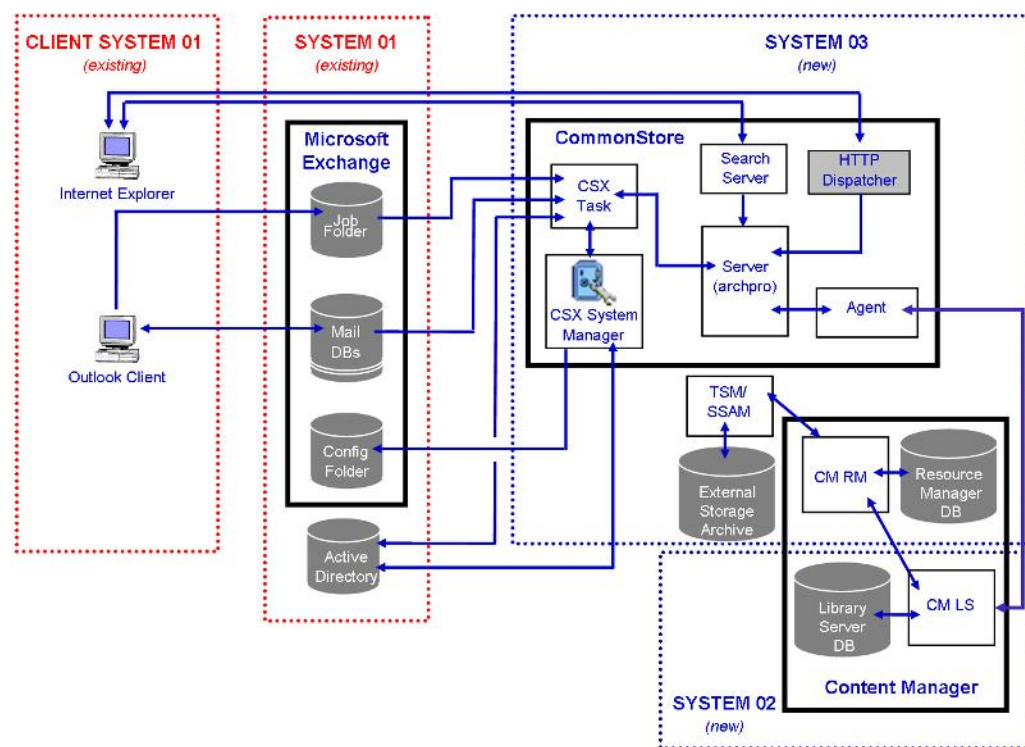
IBM DB2 CommonStore for Exchange Server (CSX) Architecture and Process Interaction

Much of the same information provided for IBM DB2 CommonStore for Lotus Domino applies here also. The CommonStore Task and archpro Server would probably be on the same new system. The agent must be physically located on the same machine as the IBM DB2 CommonStore for Exchange Server (archpro).

Active Directory would probably not be on the same machine as the Microsoft Exchange server, nor would IBM DB2 CommonStore. The primary reason for this is that directory and email servers are critical to normal operations of the business. For this reason directory and email servers are generally given dedicated system resources. Adding extension software for applications that do not support these functions is not a recommended practice.

The figure below illustrates how the IBM DB2 CommonStore for Exchange Server solution components are distributed as implemented in this Solution Starting Point for scenario B.

Figure 2. Tested Solution Component Architecture for IBM DB2 CommonStore for Exchange Server



Planning for Growth

Proper consideration and planning is required to align the business requirement to the technical requirement of moving or growing from one technology base to the next. Not having a plan for growth or not knowing *What's Next* can cause architectural issues. Use the following information to help your customers avoid costly mistakes later.

- With small business customers, it is not unusual to be without the detailed statistics required to fully assess the storage infrastructure required. Sometimes the permanent storage solution will follow the initial implementation of the CommonStore pre-load. For a short time after, assessments can be made to fully assess the customer need and implement what is appropriate later. Consider the following when making these choices:
 - In addition to the requirement for data backup and growth, a plan for backup and restore of the customer environment should be considered. This topic is touched on later with guidance provided in [Determine Physical Sizing](#).
 - The plan for backup and restore should include the planning of policies and the supporting infrastructure. IBM Tivoli Storage Manager is provided as part of the pre-load license bundle for this reason. Without this plan in place, the customer can not achieve a permanent, reliable and stable environment
- It is typical to see a year or more between the customer's adoption of eMail Archiving and the realization that they need records management if they are left on their own and given no line-of-business story/strategy for migration.
- Failure to communicate the NEED and VALUE of a future compliance and records management strategy clearly to the customer early leads to problems later (sometimes including the loss of that customer).
 - After a year has passed, major taxonomy changes (item types) would be required to move from simple email archiving to a compliance and / or records management strategy. This causes a cascading effect that impacts numerous integration points in the environment.

- Delaying the implementation of a basic compliance data structure means volumes of email now needs to be normalized. Indexing records to be managed by a records management solution (classification of documents) after a year or more is a huge expense and often leads to a business decision to take the loss and gain the risk associated with not knowing what they have in their archives

This Solution Starting Point addresses these issues by providing a line-of-business (LOB) or stepped migration story that mitigates the customer's risk while saving them money. Combining a *What's Next* vision with a services strategy allows the business partner to talk about predefined assets and a modest, but stepped migration to a compliance strategy that gets immediate results for their customer without a huge investment up-front.

Assess Performance

For this solution, performance and throughput is really focused on the processing of email data by IBM DB2 CommonStore. It greatly influences the perception the customer's users have of the solution. They are also a key measurement of the success or failure of the engagement. Making the correct decision with and for the customer requires the proper application of technical measurements and assumptions the business partner must distill from an analysis of the business need, application requirement, and the 'as is' or 'to be' architecture of the customer's environment.

The considerations and guidelines below will help you establish the proper level of performance and throughput for the customer.

Assumptions

The following assumptions are made for the information that follows. This will help you form a baseline for conclusions you draw from each customer's unique requirements

- Average email size is 80KB
- 100MB connection between Microsoft Exchange, IBM DB2 CommonStore for Exchange Server, and IBM DB2 Content Manager
- IBM DB2 Content Manager is configured to handle the email volumes sent by IBM DB2 CommonStore for Exchange Server. IBM DB2 Content Manager itself is not a bottleneck (see [Determine Physical Sizing](#)).

Overview of Components during Archiving

When an email is archived it is processed by the following IBM DB2 CommonStore for Exchange Server (CSX) components. See [Define the Architecture](#) for diagrams of these processes and the CommonStore *Administrator's and Programmer's Guide* (PN: SH12-6742-03) for a more detailed description of each.

- CSX task, worker thread
- archpro
- Content Manager agent
- CSX task, committer thread (post-processing of message)
- Resource management (object storage is transparent to CommonStore function of Content Manager Resource Server and Tivoli Storage Manager or System Storage Archive Manager)

Knowing the basic sizing elements enables you to perform estimates of performance and throughput in emails per second. An understanding of the architectural components helps to understand optimization of the configuration (discussed below).

Single vs. Multi-threaded Archiving



IBM DB2 CommonStore allows single or multi-threaded configurations. For better throughput, the multi-threaded configuration is recommended.

NOTE: The following data assumes a dedicated Resource Management System (the optional SYSTEM 03 is added). See *Systems Architecture of the Overview Guide* for more information.

For a single-threaded configuration (1 worker, 1 archpro, 1 Content Manager agent), expect to see a little less than 1 email per second archived (or approximately 3,000 per hour). When considering what configuration is right for your customer, consider the following:

- Just adding more CSX worker threads will not help - you need to add more Content Manager agents also.
- Just adding more Content Manager agents will not help - you need to add more CSX worker threads also.
- 10 Content Manager agents won't give you 10 times the throughput on 1 agent - expect that 1 archpro instance can give you up to 4 times the throughput of 1 Content Manager agent. Therefore, configure the number of Content Manager agents in the range of 4-8 (and about the same number of CSX worker tasks).
- 1 archpro instance gives you a throughput of ~3 emails per second (or 12,000 per hour).
- There is no data available showing that 1 archpro instance running on a multi-processor would provide better throughput.
- Tests have shown that a second archpro instance, configured with 4-8 Content Manager agents and running on the same 2-CPU box can give a slightly smaller throughput. Therefore, you get to 8 emails per second throughput.
- For a four CPU machine, though, you should not assume that four archpro instances will give you four times the throughput of a single instance. You should only assume a threefold throughput (12 emails per second, or 36,000 emails per hour)

IBM DB2 CommonStore for Exchange Server Configuration

The following restrictions exist when configuring a CSX server:

- Each Exchange server can be serviced by just 1 CSX task. 1 CSX task, though, could connect to several Exchange servers. For performance reasons, this is not recommended.
- The number of mailbox stores on the Exchange server is irrelevant for CSX configuration.
- Each CSX task can connect to just 1 archpro instance. Therefore, CSX can offload a maximum of ~12,000 emails per Exchange server. In a single Exchange server configuration, you cannot increase the maximum number of parallel archiving threads beyond the 4-8 mentioned above.
- Is that going to be enough? Yes! Assume an Exchange server with 2,000 users where each mail user is receiving and sending 50 mails per day. This allows 100,000 emails to be archived per day. The connected CSX server needs to run roughly 8 hours to support this volume.

When Performance and Throughput Are Critical

Multi-threaded configurations are possible with IBM DB2 CommonStore for Lotus Domino (CSLD) but are more difficult to configure. In spite of this additional difficulty, do not avoid implementing a CSLD solution. A performance improvement of approximately 50% can be achieved over an implementation using IBM DB2 CommonStore for Exchange Server.

If the customer currently has a Microsoft Exchange email infrastructure and strong performance of their email archiving solution is crucial, consider recommending migration to a Lotus Notes email infrastructure. Please refer to the IBM DB2 CommonStore *Administrator's and Programmer's Guide* (PN: SH12-6742-03) for a more detailed description of this configuration.



For additional performance information for IBM DB2 Content Manager, see the IBM Redbook *Performance Tuning for Content Manager* (SG-246949), available at <http://publib-b.boulder.ibm.com/abstracts/sg246949.html>.

Determine Physical Sizing



Physical sizing is the estimation of Disk Access Storage Device (DASD), other machine attributes, and taxonomy considerations to support the performance, throughput and relative data volumes. One of the assumptions for performance and throughput was that IBM DB2 Content Manager (on the eMail Archiving Storage System) is NOT a bottleneck. This implies that planning has occurred to properly perform a physical sizing of the business application and the systems that support it.

Physical sizing is more difficult because the parameters are always dependent on the business application's taxonomy as well as elements of the performance and throughput measurement. This makes the formula for calculation complicated and radically different from one customer to the next. If external storage will be added, then some planning is required to size the amount of disk. Typical questions include:

- How long will the e-mail be archived?
- How many e-mails will be archived per day?
- What is the average size of the e-mails?



Included with this Solution Starting Point is the *SBE Assessment Worksheet* that includes information developed at workshops with partners currently offering this solution practice to their customer base. This assessment questionnaire provides a comprehensive listing of questions for you to choose from based on the customer opportunity you are faced with. Guidance is included on what questions to use on initial interviews, definitions, and typical answers you might expect from the customer.

Once this information is known, then you can calculate the size of the repository that is needed for your customer. The RedBook and tooling described below can help you dramatically in understanding this sizing exercise.



Fortunately, IBM has sponsored development of the Redbook *Best Practices for Setting Up an IBM CommonStore Solution For Mailbox Management, e-Mail Retention and Discovery* (SG24-7325-00), available from <http://www.redbooks.ibm.com/abstracts/sg247325.html>. This Redbook includes significant detail on how to approach assessment, sizing, and more. It includes both a documented case study and a 'Sizer Tool' spreadsheet. While sizing a content management implementation is a topic unto itself, this reference and tool does a nice job of helping you to estimate what your customer needs. An alternate link for this Redbook is <http://w3.itso.ibm.com/abstracts/sg247325.html?Open>.

Note:



A copy of the 'Sizer Tool' spreadsheet is provided as-is, optimized by the practice team working with this Solution Starting Point to help answer some of your sizing questions. This spreadsheet can be modified as needed to meet your customer's requirements. The Redbook website is updated occasionally to provide updates to the tool.

In addition to sizing for data storage, some consideration should be given to Backup and Restore policies and infrastructure. IBM Tivoli Storage Manager is bundled with this solution because it provides an extremely cost effective way to manage the infrastructure and its growth, greatly reducing administration overhead. Two references are available for business partners on this topic:

- the *Backup and Recovery and Maintaining CommonStore Solution Integrity* section of the RedBook *Best Practices for Setting Up an IBM CommonStore Solution For Mailbox Management, e-Mail Retention and Discovery* (SG24-7325-00).
- the *Backup requirements* section of the Redbook *Content Manager Backup/Recovery and High Availability: Strategies, Options, and Procedures* (SG24-7063).

The customer review and assessment should include mention to the customer of this as a requirement and next step for establishing a permanent, reliable and stable environment.

When Performance Matters

This solution's SDA (Solution Deployment Accelerator) assets are certified for both the System x 3655 and the IBM LS21 BladeCenter Solution. For middle market customers who have not invested in the Blade infrastructure this will present budget obstacles unless the investment can be justified. To avoid these obstacles the business partner must guide customers as they consider solution areas (current and future) they must be prepared to tackle.

The selection of these two platforms was designed to provide coverage for both an entry and growth investment decision the middle market customer could be faced with. More information on strategies partners should consider for positioning or contrasting these platform decisions is included in the next section.

Below are some common applications areas partners will want to identify and consider to assist for the customer with this decision process:

Business Continuity

- Features to look for
 - I/O & Storage to rapidly execute data backup/restore activities and expedite system recovery
 - High Availability to enable 24/7 access to mission critical applications and data
 - Systems Management to monitor, maintain, and maximize server availability and utilization
 - Security for compliance, audit, and access control of sensitive content
- Other considerations
 - Density
 - Scalability
- Attributes & sample applications
 - Easy to install and administer email and content backup, archiving, & retrieval

Email / Collaboration

- Features to look for
 - High Availability to enable 24/7 access to mission critical applications and data
 - Systems Management to monitor, maintain, and maximize server availability and utilization
 - Integer Performance for the horsepower to accelerate processing and enable rapid response
- Other considerations
 - Scalability
 - Memory Throughput
 - I/O & Storage
 - Density

- Distributed Deployment
- Attributes & sample applications
 - Email and collaboration solutions including Lotus Notes / Domino and MS Exchange to connect employees and content with ability for quick search and retrieval

Content & Document Management

- Features to look for
 - I/O & Storage to rapidly capture, store, and retrieve content
 - Memory Throughput to manipulate large amount of data
 - High Availability to enable 24/7 access to mission critical applications and data
 - Scalability for investment protection to meet the future needs of a growing business
- Other considerations
 - Floating Point & Integer Performance
 - Systems Management
 - Distributed Deployment
- Attributes & sample applications
 - Content management and workflow solutions including DB2 Content Management to handle core business content as well as legal and compliance records management

Decisions on investment in the infrastructure can be made easier once all aspects of a customer's solution requirement (current and future) are considered. Although in this release of the solution, both the BladeCenter and x3655 have been certified with DS4200 external storage technologies, in theory, the SDAs may be used on other platforms. Partners should work with their distributors to identify demand and focus resources to implement deployment on specific platforms (see [Requirement Discovery](#) on page 9 for more on hardware alternatives).

Assessing the Customer Solution: ENTRY and GROWTH Strategies

In a solution sense, partners should be sure to position price and market share (how many customers have BladeCenter infrastructure or not) on an equal footing with scalability as dimensions for assessment of the hardware solution. At first glance, one might be tempted to describe the LS21 Blade as being less scalable than the x 3655, especially if considered independent of the BladeCenter. When considered alongside the BladeCenter, it is clear the customer can get into the x 3655 solution at a lower cost considering less infrastructure (no BladeCenter) is required. To effectively assess this solution for customers, partners must also consider the dimension of growth.

From a solution perspective, we need to understand the value behind the infrastructure a BladeCenter brings and the growth dimension is a key to this understanding. It is fair to envision the entry (x 3655-based) and growth (LS21 BladeCenter-based) solution strategies as being separated by what could be termed an *integration threshold*. An environment requiring growth of up to 3 and possibly 4 servers would be best suited with the entry strategy, while an environment planning to grow beyond 4 servers, possibly including a requirement for greater investments in storage, should plan for investment in the growth strategy.

This section is provided to contrast the two options and allow the partner to help their customers assess the hardware need with an eye to the future. Growth is a pivotal factor in this decision process so accurate sizing is important. There is a tendency to underestimate future growth and the cost of transition from one platform to another.

Contrasting the Server Options

The selection of the x 3655 and the LS21 Blade servers as the two platform options highlighted in this solution was made after careful consideration by IBM hardware professionals. This requirement integrated the feedback from partners and distributors who shared their experiences and understanding of the market based on customer interaction. The goal was to achieve broad a market share coverage as possible by delivering the value needed to cover both entry and growth opportunities.

It is important to know a bit about the value of each of the servers before examining how they fit into an entry or growth strategy.

- *Value Proposition for x 3655, the ENTRY Strategy* - Offering great I/O performance (e.g., PCI-E and HTx options), up to 64 Gb of memory, and option for up to 6, 2.5" drives the x3655 is a great value. Add to this the fact it can be integrated into more common rack infrastructures and it becomes a great entry strategy for customers exploring options for implementing applications that can place high demands on an environment's infrastructure (like email archiving).
- *Value Proposition for the LS21 and BladeCenter, the GROWTH Strategy* - Offering great performance and density for memory intensive applications, high level of server and switch integration, 32 Gb of memory and option for up to 3, 2.5" drives. The BladeCenter offers a shared, modular infrastructure, more efficient power and cooling, and helps customers sustain a high ROI by reducing costly infrastructure (e.g., switch, cabling, etc.) upgrades as growth occurs.

A number of factors related to consumability and extensibility were considered at the component level also. These included but were not limited to things like integration with other systems and devices (e.g., SAN, DASD, network). Additionally, thought was given to manufacturing of the components with respect to availability and scalability independent of and also when integrated to the entry and growth strategies discussed in this document.

Understanding the Integration Threshold

The integration threshold can be defined as a point in time a customer crosses when their operational environment grows beyond an entry strategy and reaches a point where implementing a growth strategy is required to avoid a diminishing return on their investment. This change to a growth strategy allows the customer to sustain a much higher rate of return on a continued investment in their infrastructure. When considering a customer growth requirement, if changes exceed the definition of an entry solution too soon a growth strategy should be considered from the start.

Below are a few simple discussion points the partner should consider when reviewing the hardware assessment they supply to their customers.

- The BladeCenter provides a high degree of switch integration. Many times, costly switches must be upgraded after the addition of more servers (e.g., x 3655) where the BladeCenter allows for seamless integration of up to 14 servers. The BladeCenter raises the integration threshold by allowing less cabling (i.e., points of failure and cost), better server integration (i.e., integrated chassis, administration), and switch integration between a greater number of devices.
- Storage provided in the x 3655 is more than adequate for many customers who have an email archiving requirement. Higher volumes for data storage (especially true if legal or regulatory retention requirements are a concern) can quickly drive growth and demand for a more comprehensive solution. Both the BladeCenter and the x 3655 allow connections to external storage but the BladeCenter may be the answer when managing a Storage Area Network (SAN) becomes part of the requirement.
- The partner's most important job must be helping customers consider their future goals and, managing growth heads the list. The BladeCenter can help the customer manage their future by controlling growth through a single 14U chassis (see Note below) that conserves expensive real estate, managing power and heat through the integrated BladeCenter mid-plane, and by providing one touch administration through the integrated management module.

Note:

The new BladeCenter "S" Chassis (just announced) is designed to offer an entry strategy for the Blade technology in terms of both initial investment and cost of ownership. It includes integrated storage options and a standard (110) power requirement.

For customers choosing an entry strategy, the partners' challenge is to help them understand when their need will reach this integration threshold (at about 5 to 6 servers or other devices), When this point is reached server integration becomes much more costly from both an operational and equipment perspective. The strategies discussed in this document provide a means to assess and grow the infrastructure in a stepped implementation. Changes will be required but with proper planning can be managed in a way that provides a kinder and gentler customer budget strategy.

Assess Return on Investment

As part of assessing the customer's requirements, you will identify potential ways to increase and enhance the company economics in terms of return on investment (ROI) and associated costs. Here are examples of soft dollar (areas whose values are difficult to quantify) and hard dollar (easily measurable) benefits that might apply to your customer.

Soft Dollar Benefits

- More efficient communications for the users are provided through improvements in email organization. Users are provided tooling to organize, search for, and retrieve their business email efficiently.
- The new email archiving solution will allow size limits on email systems to be removed since they are managed more efficiently. High volumes will no longer be an issue on those high-use high-performance email servers. This performance bottleneck is managed by the email archiving solution.
- Legal actions may require documents to be placed on 'hold'. This benefits legal discovery by enabling searches of the IBM DB2 Content Manager repository to locate these email documents. These documents could be managed without resorting to full records management.
- Becoming compliant with regulatory issues (avoiding possible fines) and achieving a better rating (e.g., Basel II - Banks).
- Improved employee productivity due to an optimized email experience with faster search, better organization, and fidelity of archived content.

Hard Dollar Benefits

- Reduce license costs by offering search capability from the Lotus Notes or Microsoft Outlook email client, thus enabling a customer to avoid the cost of an IBM DB2 Content Manager license for every user license.
- There is no need to maintain any alternate methods for receiving large files (e.g., FTP sites established and managed to support uploads for receipt of large inbound files).
- Redirect administrative resources that were previously providing user services and administrative services for email management (e.g., load balancing, administrator directed archiving).
- Leverage cheaper storage - granular data organization allows content to be assigned to storage that compliments data use (e.g., low demand/use content can be directed to tape while only high demand / user content can be directed to more expensive but faster storage options)
- Reduce cost of future Records Manager deployment effort by classifying documents up-front and avoiding the cost of future data normalization and migration issues (e.g., classification of documents long after they are stored is a huge expense).

Allies in a New World

Since measuring ROI for a compliance solution is difficult, making a convincing case for compliance to a customer's IT department without allies they can trust is a bad idea for gaining traction within the mid-market business. While the IT department may have a sense of how expensive a compliance solution can be, many times it has absolutely no understanding of the value the solution may bring to the business.

The people who know compliance and understand the risk often control large portions of the IT budget in these companies. They can demand and get an equal portion of respect from the company's IT department. Subject matter experts representing these departments, enlisted as allies, can provide the support an IT organization trusts (or at least respects) during discussions on the topic of compliance.

Before moving forward with proposals for this solution, solicit individuals related to one of the departments, professions, or of the rank listed below as champions first. Professionals in these areas or with this expertise would understand the value of compliance and records management better than an IT professional because of the rules and regulations imposed on them in their area of specialization.

- Finance - CFO, CEO
- Banking - Accounting, HR Retirement Management
- Insurance - HR Benefits Administrator, Product Liability Specialists
- Healthcare / Pharmaceutical - professionals unique to these industries (e.g., chemists, doctors, lab technicians, researchers)

It is important to remember the phrase "a level of compliance" should always be used, not the definitive like, "in compliance" or "compliant". Simple wording could have the effect of transferring liability for not meeting some legal or regulatory obligation the customer has to the business partner. The business partner can and should assist in writing the *functional* requirement describing "a level of compliance" their customer wishes to achieve but the requirement should remain the responsibility of the customer.

Avoid definitive or declarative language pronouncing the solution or a solution feature makes the customer compliant; instead focus on fulfillment of the customer's functional requirement. At all turns the customer should be reminded in verbal and written communication that they are the expert on what "level of compliance" they need and the business partner acts as the expert for implementing solution features to accomplish goals the customer has set forth for the business partner in their function requirement.

Demonstrate Prototype

A prototype of the concept should be set up to validate your understanding of the customer requirements as well as your general architecture. Based on the customer feedback to this prototype, you may find areas in your requirements analysis and architecture that need to be reassessed or reorganized.

You can create one from scratch, but cost estimate vary widely based on your skill level and the customer requirement. For a measurable (possibly faster) result, follow the *Implementation Guide* of this Solution Starting Point to install and deploy the middleware and technical assets. Estimates for creating a prototype using this Solution Starting Point are provided as the SDA and Low (email archiving alone) and as the Medium and High (including a compliance strategy) ranges provided for in [Table 4](#).

If architectural concerns are not on the table for discussion, IBM Solutions Builder Express provides a new asset called SAT FastStart, a single machine software stack implementation that can be auto-loaded to a partner's laptop that allows this solution to be demonstrated. Contact your SBE representative for more information on how to acquire SAT FastStart.

There are two SAT FastStart assets built for demonstration, technical reference, and technical training called SAT FastStart for both CSX and CSLD. A special Use Case, *Preparing for Demonstration of Archiving for Customer Email Data* may be found in the Implementation Guide for Scenario A (CSLD) that provides guidance to load a customer's data for demonstration of the solution.

To demonstrate the solution to your customer using SAT FastStart, you can run through other use cases provided to show how the prototype illustrates the typical user scenarios, or, if the samples from this Solution Starting Point were used, you can use the instructions provided for them in the Implementation Guide.

Final Design and Customer Sign-off

During the demonstration of the prototype to the customer, additional requirements will become apparent, or better understood, leading to addition of select customizations and enhancements to the finished design. In this task, you are looking for customer acceptance of the prototype design prior to full development of the solution.

Following the demonstration, you should meet with the customer to revisit the list of requirements developed earlier, in addition to any new or redefined requirements. The customer should understand that the prototype should not be expected to address all of their requirements, but is a representative sample of the overall solution. Customer signoff indicates the customer's agreement that a certain number of requirements are addressed by the solution and an acknowledgement that the solution provides a good foundation for addressing the remaining requirements.

The services hours required for this task will vary depending on the relative complexity and customization of the prototype. See [Install Technical Assets](#) on page 29 for strategies on leveraging the solution technical assets (samples) to gather and complete the customer requirement and gain their sign-off. The estimates include creation of a deliverable document.

Define the Project Plan

Your project plan should describe how and when the project will be completed. It will detail your deliverables and your timelines and the plan to deliver the solution.

Consider including the following information:

- All dependencies, risks and assumptions
- A comprehensive schedule listing major checkpoints for yourself and the customer
- A clear and substantial project scope definition, for you and for your customer
- A detailed list of deliverables the customer will provide (such as organization charts, data, formats, and so forth)
- A detailed list of deliverables you will provide (including documentation, training material, and the solution itself)
- A staffing plan that describes the required skills and responsibilities for each team member
- A description of the testing to be done prior to and during deployment, along with quality requirements that need to be met
- A standard list of project management and quality assurance procedures
- An official definition of exceptions, including scope changes and adding or removing project team members

Install and Configure Essential Software

These tasks describe what is necessary to install and configure the products required to support this solution. The implementation details for these tasks can be found in the Implementation Guide for this Solution Starting Point. Not all tasks listed will be applicable in every implementation and will vary dependant on the individual customer requirement.

- **Install and configure IBM Tivoli Storage Manager (TSM)** - A thorough review of the business use of the data (emails) to be archived is important. Real cost savings can be achieved by mapping content to appropriate devices (e.g., is the data needed often in the business process so it must be ready and retrievable quickly - on-demand?). The physical storage nodes should be chosen based on this mapping and lower cost alternatives may be leveraged when possible without loosing efficiency.

Remember that implementation of backup and restore policies and infrastructure can be crucial to support a reliable solution environment. TSM takes incremental backups one step further by providing a method that eliminates any further "full" file backups after the first full backup. Additionally, through "collocation" and "reclamation" processes, TSM automates allocation of storage media moving data to make maximum use of space across multiple tapes or DASD. This is explained in more detail in the *Tivoli Storage Manager Administrator's Guide*, GC32-0782.

Installation of this product takes little time but to take full advantage of these efficiencies and cost savings allowances must be made in the project plan. As mentioned above, time should be allowed for the study of data use characteristics for the business. This knowledge should be applied when configuring the product.

Note:

IBM Tivoli Storage Manager (TSM) is included with the license for IBM DB2 Content Manager (CM). Customers wanting to take advantage of Tivoli for full backup and restore of the environment would need to consider additional TSM licenses. If the customer wishes to take advantage of external storage solutions discussed in this Solution Starting Point, and compliance is important, the IBM System Storage Archive Manager (SSAM), an extension of TSM, may be used. License for SSAM is not included in the cost of CM. Pricing is set by the terabyte of data stored and additional service hour aspects for the engagement should be considered to cover SSAM configuration.

- **Install and configure IBM DB2 Content Manager** - Business use of the data plays a role with this product too. Security and consumability starts with thorough planning of the item-type taxonomy behind the email archiving and compliance solution. Testing and tuning the taxonomy may be facilitated by using (and customizing, if necessary) the Item Type Migration Model technical asset.

IBM DB2 Content Manager relies on other products included with it. This task includes installation of these products. They are: WebSphere Application Server (plus fixpack and GSKit) and IBM DB2 (plus fixpack). Installation for these products accounts for approximately 40% of the total installation of IBM DB2 Content Manager.

The medium and high ranges of services hours are slightly higher because of additional steps needed for a compliance requirement. This additional time is spent on configuration for the most part. Representing the item type and meta-data model relationships dedicated to the compliance solution in the IBM DB2 CommonStore and IBM DB2 Content Manager configurations, respectively, account for most of this configuration effort, about four hours.

Note:

It is important to understand that the Implementation Guide is designed to provide easy installation of the four solution scenarios. Two do not include IBM DB2 Content Manager (scenarios C and D). Without IBM DB2 Content Manager the compliance portion of the solution (and related samples) may not be leveraged.

The solution task estimates, provided earlier in this guide, reflect a simple installation of email archiving (scenarios C or D, with and without SDA), and an implementation of the solution with a compliance requirement considered (scenarios A or B, with and without compliance technical asset enhancements).

- **Install and configure IBM DB2 CommonStore** - Once again, business use of the data is important. These solutions are highly customizable for a reason. Most of these types of solutions require some mix or balance of user-initiated and automated processing. The customer's data use requirements will drive these optimizations.

Note:

Installation of IBM DB2 CommonStore takes approximately 30 minutes. The balance of the time dedicated to this task is configuration.

- **Install and configure IBM eMail Search for CommonStore (eMS)** - IBM eMail Search for CommonStore delivers discovery capability to email archiving server. A system with IBM DB2

Content Manager, IBM Tivoli Storage Manager, IBM DB2 CommonStore, and IBM eMail Search for CommonStore delivers a more comprehensive eMail archiving, compliance and discovery solution. At the time of this writing, if eMail Search is added to the solution, the item types must be set to a resource item type in Content Manager in order to be able to be searched. The regulatory or legal requirements on how long data needs to be saved, or discarded needs to be updated in the configuration. Based on the number of different types and requirements for the data are the basis for the time spent in this task.

A Solution Deployment Accelerator (SDA) has been provided with this Solution Starting Point to significantly reduce the time required to install these products. This will provide more time for you to focus on business data use requirements and sizing tasks.

The SDA includes a Quick Start Guide that explains how to quickly deploy the solution environment.

There is also a related IBM Solutions Builder Express Technical Reference named *How to Rapidly Assemble and Deploy Solutions* (available from <http://www.ibm.com/partnerworld/solutionsbuilder>), that describes how to adjust the SDA to support a customized solution. Ensure you allow time for development of this important component if you wish to use it to deliver your solutions to the customer.

Install Technical Assets

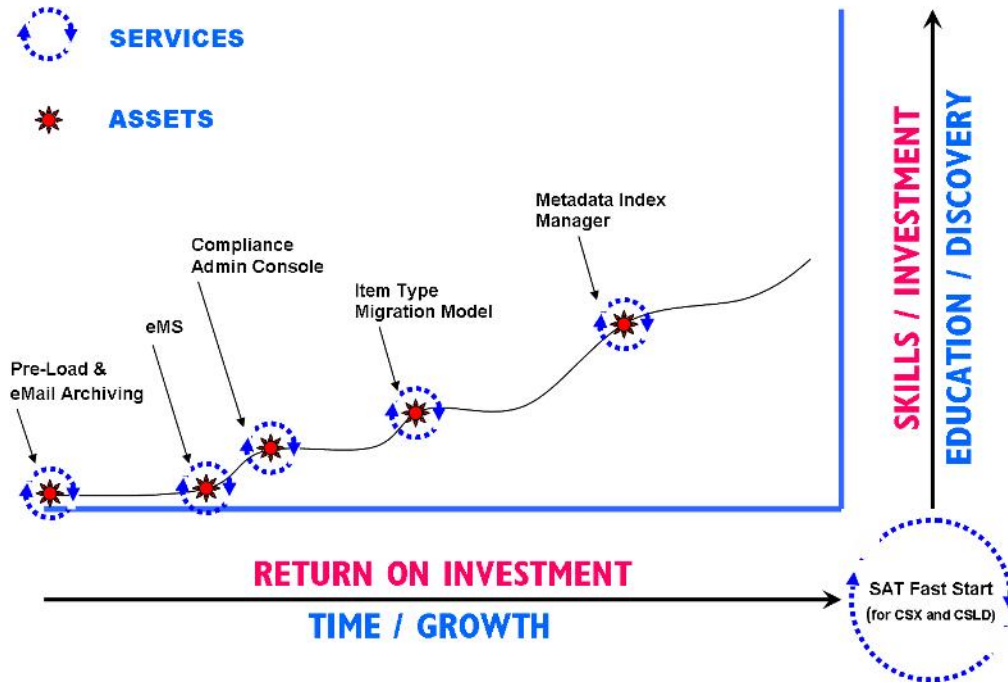
The following list describes how the technical assets are used in this solution:

- **Item Type Migration Model** - This is a multipurpose asset that supports a number of goals unique to each customer engagement:
 - **Testing of the item type model for the customer** - Used during solution development and deployment to allow easy testing and manipulation of the item type model without tedious manual manipulation.
 - **Deploying the solution** - For deployment of the solution to a customer's environment from the business partner's lab or to develop and quickly deploy any number of industries specific versions of a solution as part of a proof of concept library.
- **Meta-data Configuration Manager** - Customized and then deployed in the customer's runtime environment to provide classification of email documents (to support compliance and records management strategies).
- **Compliance Administration Console** - Customized and then deployed in the customer's runtime environment to provide an interface for a compliance administrator to view the organization's emails that have been archived.

By now, the reasons for leveraging the technical assets provided with this Solution Starting Point should be clear. Determining the timing and sequence of deployment of these assets depends on the needs of the customer's stakeholders. Customizations and enhancements to the assets are driven more by customer need in the beginning. Most enhancements will be required but some will require more customization due to the unique requirements of the customer.

Starting with a diverse product base brought with the pre-load and adding code assets brought by this solution, the business partner can begin to develop other opportunities to expand their business with the customer. By wrapping value added services around these assets the partner can provide quick, measurable ROI that makes sense to the customer. One possible approach is pictured below.

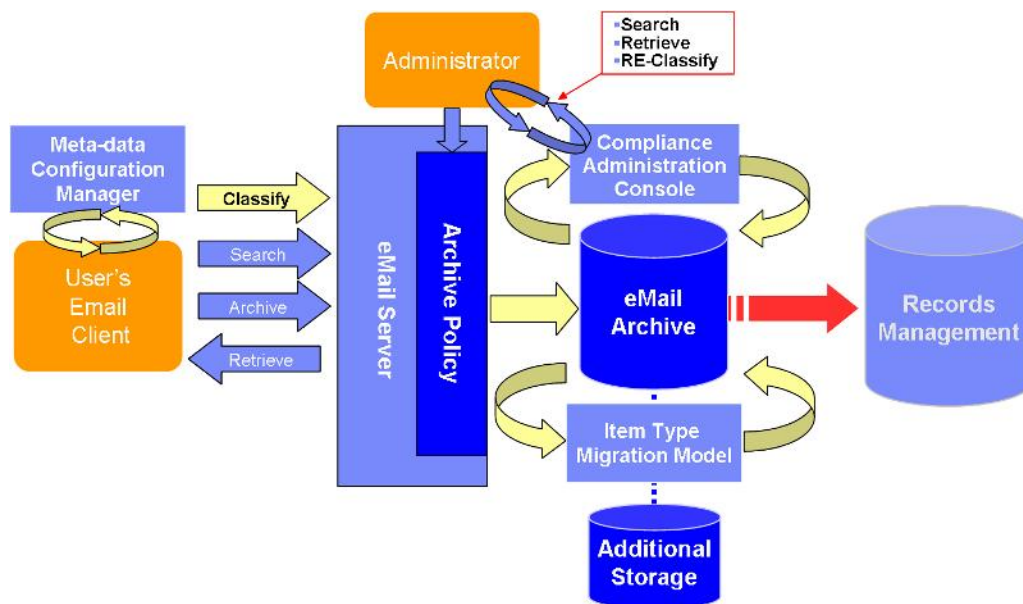
Figure 3. Assets for Growth



Batch processing, security, and special data use modeling should come first but will be iterative with a department level approach. Once the decision is made to begin development of a compliance strategy for the customer, special customer requirements could be implemented through work on the Compliance Administration Console technical asset. In a Lotus environment, the Meta-data Configuration Manager asset can be modified to implement and then expand fileplan and rules models through the requirement development and discovery process.

The figure below shows the relationships (functionally and at a high level) between the key solution elements and significant business roles (users and administrators). It also positions the technical assets relative to these solution elements.

Figure 4. eMail Archiving Components



This list maps the key solution elements to their related products:

- User's Email Client - Lotus Notes or Outlook client
- Archive Policy - component of IBM DB2 CommonStore
- eMail Archive - IBM DB2 Content Manager
- Additional Storage - IBM Tivoli Storage Manager
- Records Management - the full compliance solution (possibly IBM DB2 Records Manager)

Customize Technical Assets for Production

User interface enhancements, the addition of error handling, and internationalization are just some of the elements that should be considered to customize the technical assets for a production environment. The assets (listed below) are meant as the foundation for a products and services strategy provided in this Solution Starting Point. They all require modification before applying enhancements to facilitate discussion and definition of the compliance requirement.

- Item Type Migration Model
- Compliance Administration Console
- Meta-data Configuration Manager

The services hours estimate for this task assumes the addition of error handling, look and feel, customizations, and internationalization. Implementing the Item Type Migration Model varies since "customization" is difficult to fully scope out.

Enhance Compliance Administration Console

It is important to understand that enhancements are actually strategic assets that should be selected and supplied as a method of fulfilling a customer's compliance needs. Ideas for enhancements are designed to support the cost model of both the business partner and their client.

Some enhancements have been implemented as proof of concepts within the sample. They will require customizations or optimizations (sometimes extensive) depending on the customer specification. The task estimates provided in this guide do not include documentation.

The following enhancements are related to administration of the business compliance rules. A compliance administrator would be the primary role associated with the use of this technical assets and the enhancement discussed below.

Logon and Custom Privilege

The logon validation that is implemented in the sample allows all registered IBM DB2 Content Manager users to access the administrative console. Depending on a customer's need and usage this could be undesirable as it introduces risk and privacy issues. There are several ways to implement security to only allow a select group of users to access the interface. Samples include via LDAP call, additional group password, authorized user file, J2EE security, etc. However, IBM DB2 Content Manager already contains a mechanism that fits quite well with this type of need. Here are examples of the need for this feature:

- A compliance requirement that administrators view to only indexing or classification data (e.g., an administrator cannot view a CEO's email correspondence but can place these items on a legal hold).
- A compliance requirement that an administrator role passes documents for destruction. Only destruction administrators may actually delete the records. There are multiple administrators. A 'destruction' administration role should be formed and extended the correct privilege.

Extended Search Capability

The sample Compliance Administration Console only contains a limited set of the IBM DB2 Content Manager searching capability. More importantly, this sample provides a study in the Content manager API. Other advanced features customers ask for might be:

- Support for typical "operators" such as like, between, greater than, less than, not equal, etc.
- Support for other search string conjunction operators such as **AND** and **OR**.
- Support for the interface to perform checking, validation, or apply size limits on the data to be returned from a search request.
- Support of unique features in a narrow customer requirement that falls outside the scope of the feature set offered in IBM eMail Search for CommonStore (eMS)

Estimates are based upon privileges being defined in and accessed via the IBM DB2 Content Manager Administration Console. Be sure to include costs for any tasks required for integrating directory systems (e.g., LDAP, Active Directory, and Domino).

Complex Data Models

Each customer will require a set of complex data types. IBM DB2 Content Manager provides support for many complex data models. The Compliance Administration Console assumes a simple single tier data model; however, there is no restriction on extending it to support a more complex data model.

The capability of the interface implemented in this sample is sufficient for many administrator functions, but would probably need to be enhanced if complex relationships are required.

Since exposing some complex relationships can require knowledge about a specific environment, a business partner may expect complex relationships such as:

- Child components - a component that provides multi-valued detail information (e.g., a human resources record may contain an insurance policy with the 'insured' listed, a child component of 'insured' may contain related birth dates, employee number, sex and other information)
- References - a single direction join designed to relate a child component to more than one root component (e.g., an insurance application with a customer root and underwriter root component may share the same claim child component).

Use caution when preparing these estimates. These tasks can become complicated due to the flexibility provided by the DB2 Content Manager system.

Delete Functionality

A compliance administrator may need the ability to delete data. The ability to directly delete data is never recommended and is therefore not represented in the current sample. There are normally two or more roles associated with this function (deletion), to satisfy the compliance requirement. It may be more appropriate to:

- Create a "disposition" document validation that enforces rules requiring the document be deleted according to business policy.
- Upon "disposition", allow a decision or enforce a routing process for approval that meets with compliance guidelines.
- If the asset is being used for environment setup and testing a delete function may provide a very important and time saving utility. This function should be removed once production levels are achieved.

Reclassify Record Attributes

The Compliance Administration Console has the generic ability to update the attributes of an item to reclassify an email included in the current sample. However, the Compliance Administration Console is designed to be easily extended to allow **all** objects stored under an object type to be updated.

For compliance reasons the sender, recipients, subject and date should not be editable for a variety of reasons. An appropriate enhancement to the console when used in a production context would be to restrict the console so that only the records management attributes are editable within the interface.

An important feature to remember is that 'resource types' can be managed through this asset. Resource types are defined for environments where IBM eMail Search for CommonStore (eMS) is or will be implemented or as required to support a customer compliance requirement. The Content Manager Windows Client, normally used for environment setup, can not be used with resource types. Compliance Administration Console could be used to convert items with like attributes from type to type (e.g., 'document' to 'resource' types) and back or to manage test data during setup of the environment where resource types must be used.

Viewing a Document in the Result List

Although not included in the current sample, there may be a need to examine the document once the search result is returned.

In general, careful consideration is needed when determining who can see email - primarily because the repository may contain all email from an organization. For example, the ability of an administrator to view an email may have security implications. If there is any conflict of interest about a document, an administrator should not be able to view that document.

Document Forwarding

While viewing of an email is not implemented in the current sample, it may be necessary for an administrator to forward an original email to another individual. For example, an administrator may have a question about the lifecycle classification of a document and may require additional information or confirmation of a classification.

Although using email is probably the most common form of document distribution, document forwarding can be implemented in a variety of ways including - email, ftp, writing to a CD, or eClient API functions, etc.

Batch Functionality

Although not implemented in the current sample, batch features such as editing attributes, re-indexing and deleting whole groups of emails are commonly encountered requirements that many content repository administrator interfaces support. The capability allowing an administrator to alter large batches of documents at a single time rather than individually can single-handedly affect the outcome of a successful project.

Within the scope of email compliance, the considerations for these types of batch functions still apply but can be complicated if the search results contain multiple document types. Consider these things when implementing batch utilities.

Customize the Meta-data Configuration Manager

Customizations to the Meta-data Configuration Manager are strategic assets that should be selected and supplied as a method to help define and fulfill a customer's compliance requirement. Ideas for customizations are also designed to support the cost model of both the business partner and their client. Services should combine an educational component for the customer and allow the business partner to finance educational needs as well as deliver the application specific needs of the customer.

Most or all of the features offered by the Meta-data Configuration Manager, provide an opportunity to extend these functions to archive the data of other Lotus Notes applications in addition to just email archiving. If your customer's environment includes Lotus Notes business applications, this option should be considered and suggested.

Before covering further details about the Meta-data Configuration Manager asset, some records management concepts are required.

Understanding Records Management

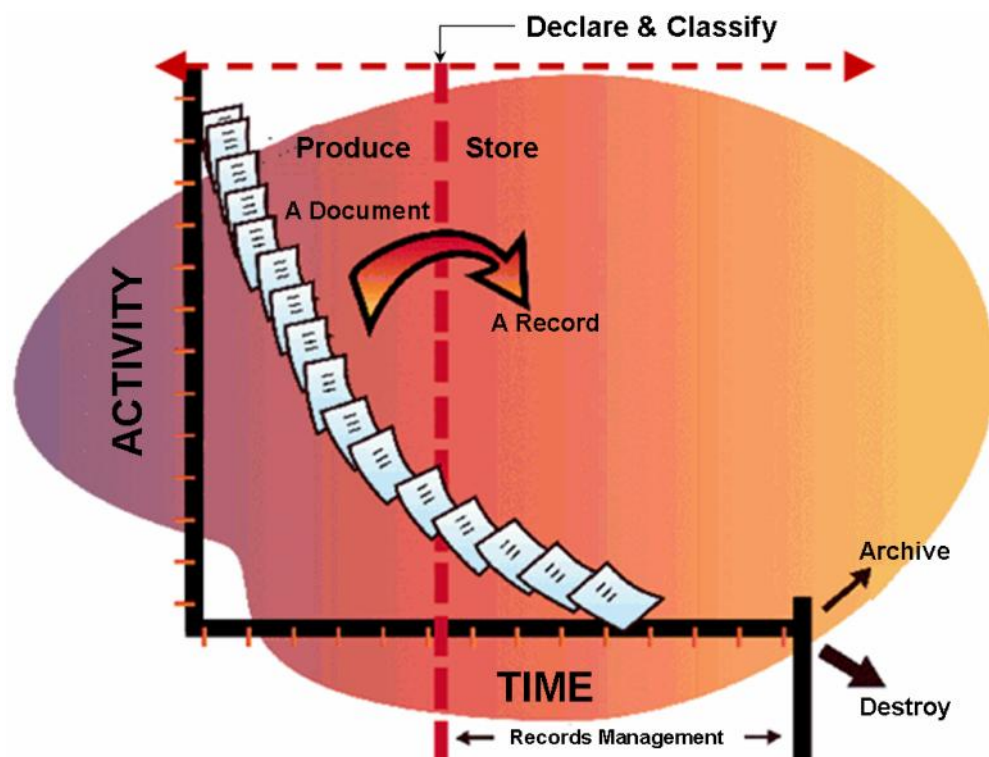


This section provides a little contextual information on how records management applies to email archiving. This is especially important if you are new to compliance and records management solutions. See the REDBOOK, *E-mail Archiving and Records Management Integrated Solution Guide Using IBM DB2 CommonStore and DB2 Records Manager* at the URL: <http://www.redbooks.ibm.com/abstracts/sg246795.html> for additional detail.

A record is when?

A record is a record when it stops changing. It can change again but then it becomes a version of its old self. The figure below depicts a simple view of how work on a document (Activity axis) tapers off as the document comes closer to completion. At some point in the life cycle it is declared a record and classified (note the vertical dashed line labeled "Declare and Classify" in the picture), so it may be found in the records management system later.

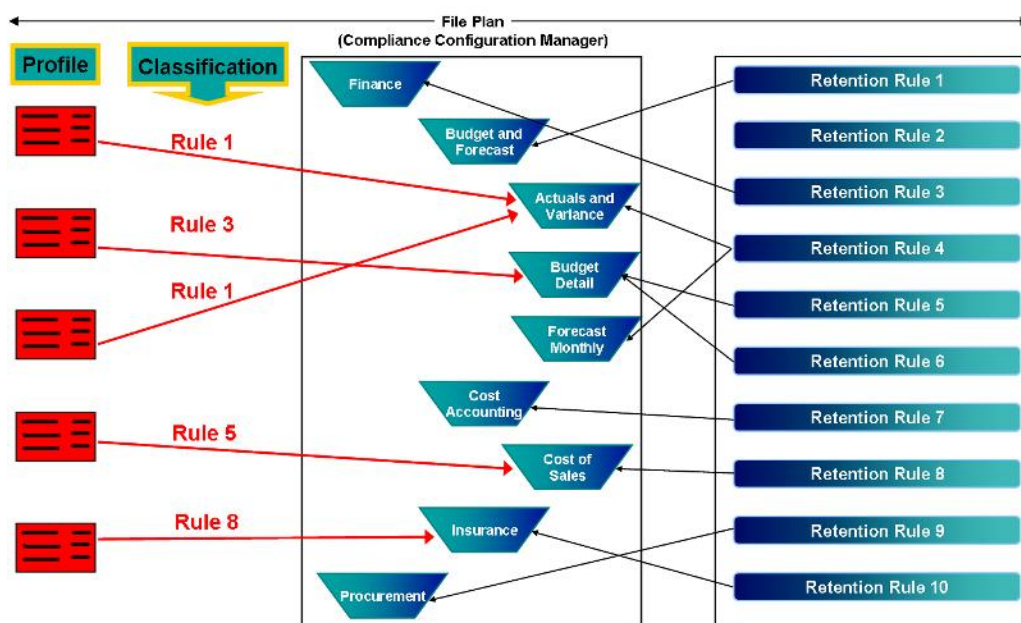
Figure 5. A Record's Life Cycle



My record's Life Cycle

After it is declared and classified, the record begins a trip through a life cycle defined by associating retention rules, profiles and a fileplan. The figure below shows an example of these relationships.

Figure 6. Compliance Information - A Record's Classification



The following terms will help you understand the life cycle of a record:

- **Profile** - associates users, groups of users, or categories of users to a fileplan node.
- **Classification** - associates the record to a business use classification (e.g., Financial, Communications, Legal Contracts, Insurance, Personnel Information, others).
- **Fileplan** - represents the logical organization of the business (e.g., Departments, Regions, Markets, others). Individual fileplan nodes can represent a category level of the plan or actually contain records.
- **Retention Rules** - determines the path a document takes through its life cycle. Retention rules include things like:
 - Times related to various conditions (e.g., hold for X years in this form or with this access).
 - Who approves moving and how the document moves through the life cycle (e.g. storage administrator must approve move A, storage administrator and destruction Administrator must approve DESTRUCTION Requests).
 - How or where a record is destroyed or stored.

A full compliance solution (e.g., IBM DB2 Records Manager) provides a much greater level of control and granularity for the records management process than this Solution Starting Point illustrates. However, this Solution Starting Point provides:

- a low cost way to define this environment
- a less dramatic and less costly way to transition the business and its users to a compliance strategy
- An approach that provides a method to insure a smooth transition to full compliance (when it occurs) that will more closely fit the need and budget of the business.

This way, you deliver compliance features and functions to the customer now as they go, instead of waiting until they are willing and able to invest in a full records management solution.

Fileplan Folder Security

The records management fileplan organization offered by the Meta-data Configuration Manager allows planning of the fileplan and assignments of the node to email records. Simple security is provided to prevent selection of nodes under development or nodes that have been withdrawn due to organizational changes but which still have meta-data stored in email document items in the archive. Potential customizations include:

- Develop an access list that controls security at the user, user group, or agent level (to use in conjunction with auto-classification).
- Look at leveraging fileplan folder security to 'reserve' fileplan nodes for special legal, regulatory, administrative or other uses to classify and organize email document items. (e.g., use the security feature to allow only testers to manage records under this fileplan container or category node in production)

Auto Classification

There are a number of customizations that could be implemented to support auto-classification of email as records:

- Development of *Automation Agents* in conjunction with modifications to the Fileplan Item form. The Fileplan Items form is part of the Meta-data Configuration Manager design (see the Implementation Guide for details).
- Extending the design of the Meta-data Configuration Manager Fileplan Item form to develop the *Automation Profile* in support of an automatic category of functions the agents provide.

The following components could be developed to implement these customizations :

Automation Agents - These two Lotus Notes agents use information from the Meta-data Configuration Manager (MCM) database to assign an email to an item in the fileplan. Automation Agents should run on the server and may have the following capabilities or characteristics:

- They would be modular and could be included as an add-on to the optimized mail template.
- They are flexible enough to be moved and used in the Meta-data Configuration Manager (not the email template) if the customer requirement demands that organization, performance improvements, or load balancing be imposed.
- Implementation as a set of two modular agents should be considered. One would manage all INBOUND emails for users and the other to manage all OUTBOUND emails. Optionally, these agents could reference an Automation Profile (details below). The importance of managing INBOUND and OUTBOUND emails is a function of both a technical and a business requirement. From the perspective of the technical requirement, a user must manage the automation using different functions and agent types (INBOUND to OUTBOUND). From the perspective of the business requirement, the agents must reference different profiles and profile types. This organization simplifies development of the agents and management of the configuration profiles.

Automation Profile - The *Automation Agents* would reference *Automation Profile* information stored in the Fileplan Item documents defined in Meta-data Configuration Manager to manage the classification meta-data. This configuration information, referred to here as the Automation Profile, would also act as control information to direct the agent's operation based on the relationship between the email content, the email user, and the business organization.

The Automation Profile design may have the following capabilities or characteristics:

- Supports members listing (like a user Group) that allows the email's addressee to be used to match the email to the profile by the Automation Agent.

- Supports a keyword listing matched to a Fileplan Node's optional classification meta-data. Selection, based upon keyword matches from the email content, would determine what meta-data would be assigned to (written to) the email record.

Profile Type

A *Profile Type* switch (e.g., STANDARD, CRM, etc...) would allow the administrator to set an *Automation Profile Type*. *Automation Agents* would use the *Profile Type* as a decision support switch to support specialized legacy integration without having to rewrite the application's agents.

This approach could be used if the compliance systems had to manage content related (e.g. email) to legacy systems. This would allow creation of instances of a profile that might apply to different systems (e.g., INBOUND or OUTBOUND messages to or from CRM systems to associate or auto-forward (email or other) records to the archive for compliance management).

- Include a "senders listing". The senders could be a collection of contact persons in a CRM system (e.g., a profile category by group of users).
- Incoming mail could automatically be dropped into the CRM system and / or classification information for the document could be applied by its ranking in the CRM system.

TO DO List

This feature could be implemented as a new decision point option provided to the users that allows more control of automated batch classification. For example, instead of processing all INBOUND & OUTBOUND email automatically, a dialog interface (TO DO List) would be provided to the email user and give them an option to decide on the disposition of emails listed in the dialog. For example, they could select a fileplan node or place a hold status for all or some of the emails presented in the TO DO List.

The interface design may have the following capabilities or characteristics:

- Checkboxes for options that can be selected. Possibly a SELECT ALL control, if appropriate.
- A description of the records being assessed. This might include the Date, Subject, and possibly any Keyword list matches used to classify it. It would also include the Fileplan Number.
- A hotlink (as part of the description) that allows the user to quickly open and scan the item.
- A Hold Status Designation to allow users or administrators to place items into a hold status (e.g., I don't want to decide now, I can't decide until I find out if I have a business reason - email is part of an email thread that ultimately may not fit the compliance requirement).

Customize the SDA Wrapper (Two Machines)

Customizations to the Solution Deployment Accelerator for Option Two: Two Machine setup bring the opportunity for a much more scalable solution. Unlike Option One: Single Machine Setup, this customization provides a much more scalable solution for the customers who have specified a solid growth requirement.

This customization also brings an important change in the way IBM distributes software. Under this new deployment model, IBM distributors can offer the business partners an option to order the customer's hardware with the solution products pre-loaded. This allows the business partner to bring this highly scalable solution to the customer quickly and concentrate on design and delivery of the business solution.

For more on this customization, see *Customize SDA Wrappers* in the *Customize the Solution* section of the appropriate scenario Implementation Guide for your architecture.

Test and Deploy the Solution

The technical assets delivered in this Solution Starting Point can be leveraged for quick setup of the test environment as well as the quick deployment of the completed application. They can also deliver a strategy for providing a stepped compliance implementation strategy that mitigates risk and saves that customer money.

Validation

The technical assets delivered with this Solution Starting Point are designed to support a rapid development and deployment environment and to achieve short-term compliance goals quickly. The design is meant to involve the customer in the design and development because efficient compliance tooling integration is about the day-to-day business process and improving the use characteristics, not hampering the user and hurting productivity. For this reason the users have validated the tooling and data user characteristic at each milestone and in many cases have been part of the design team.

Normally, this task can take between three days to two weeks with the business partner effort falling around 25% of the effort. If the business partner decides to follow the recommendations for a stepped implementation and discovery with the customer, the effort (perceived by the customer) may be less.

Integration Testing

This Solution Starting Point is designed to facilitate a higher level of customer involvement (from the beginning and throughout the project). It is also focused on returning a value in the form of facilitating the compliance functions and solving the related business needs quickly. Customers involved in the process can experience the new tooling providing a risk mitigation effect, avoiding data normalization later (e.g., email is classified automatically as it is generated), and see this as development progresses. Using this approach, integration testing happens in increments, is closely coupled with the user's production environment, and shows results almost immediately.

Certification or User Testing

Following the recommendations of this Solution Starting Point the customer is involved in identifying their archiving and compliance strategy. The customer should play an interactive role in the design and testing of user acceptance (or certification). This strategy can have the effect of transforming this task to one less formal, reducing or removing cycles more common to the user acceptance testing (or certification) task.

Deployment

As stated throughout this Solution Starting Point, all the technical assets are tuned for fast deployment. Many different deployment models may be followed and will be since this Solution Starting Point is also designed to scale slowly and with the budget of the middle market customer. Many deployments are simply a matter of switching a test environment to full access by the users and others will be more intensive and regulated. The technical assets provided accommodate any scale deployment.

Typical deployment tasks might include:

- Configuration of the taxonomy to support integration to existing systems (e.g., export/import of item types and IBM DB2 CommonStore mappings)
- Facilitate security setup to customer directory (e.g., LDAP, Lotus Domino, and Microsoft Active Directory)
- Make or implement modifications to templates for email or other applications (e.g., develop a Meta-data Configuration Manager for Microsoft environments, apply template changes in the Lotus Domino environment, and/or add enhancements to support customer legacy applications)
- Assessment and configuration of fileplan security and selection criteria based on the customer's organizational hierarchy

Setup Production Environment

Business partners report that it is common for production and testing environments to be closely coupled. What this means is that closed test environments with simulations to support integration of legacy systems are the exception. Normally, test user accounts and security are the methods used for separating the operational environment from the test environment. It is common for data to be transmitted from and to the production environment and the testing environment. The result is that except for tasks similar to the deployment tasks listed above, moving new systems to production is relatively painless.

Educate the Customer

Although not all of the customer educational requirement can be satisfied in this manner, most comes as the result of direct involvement in the design and implementation process. For the most part, estimates for this task represents those issues that require one on one. Mostly these are administrative tasks.

Since the strategy being proposed in this Solution Starting Point is at the department level, much will depend on the level of experience your customer has at each department in their organization. There are some things the business partner can do avoid educational roadblocks:

- An effort should be made to identify technology champions and to organize internal support among the user community. Compliance will have a home (a competency) related to the function certain departments or centers of expertise drive throughout the company (see [Allies in a New World](#)). Technology champions are people who are knowledgeable of the business or business compliance issues the solution solves and who understand the value to the company.
- Take care to include these technology champions and those they nominate in the design and testing of the tooling. The strategy of taking a departmental level approach with short sales cycles is meant to avoid an extensive or prolonged educational requirement. Education on this type of project is closely linked to a successful deployment. Education on the tooling is important but it must also include a cultural component to drive change. eMail Archiving and/or compliance tooling (no matter how automatic or well done) will change the way people work. Without this additional educational element, the risk of failure due to user acceptance issues is much higher.

Additional Services

Your customers may want you to provide ongoing maintenance services or additional services to enhance this solution. These might include

- **Content and storage management** - Implementing any part of this Solution Starting Point introduces great content and storage management tools and related opportunities. Be sure and think about the following topics when discussing the new systems with you customer:
 - Content management of application data, not just email
 - Opportunities to migrate data from less reliable systems to IBM DB2 Content Manager
 - Compliance management for application data. Be sure you are looking to include systems the customer has not yet integrated to the new solution.
 - Storage management including: integration of existing storage systems to be managed by the hierarchical storage management system (IBM Tivoli Storage Manager).
- **Records management** - Remember the following when discussing this with your customer:
 - If IBM Records Manager has not been discussed, this Solution Starting Point provides the information you need to assess and produce a convincing case
 - This solution is scalable but this work should be measured. If the full-compliance strategy that is right for you customer is IBM Records Manager, a plan should be developed to leverage the assets with that strategy in mind.

- Monitoring - The customer environment is being optimized along the businesses data use characteristics. More and more the availability of this critical data resource will become necessary for the businesses successful operation. Consider introduction of Tivoli products for developing a strategy for providing a reliable infrastructure (see *Related Solutions* for ideas on how the IBM Solutions Builder Express Portfolio can assist you in your next engagement with the customer).

Appendix. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM makes no representations, warranties, or other commitments whatsoever about any non-IBM Web sites or third-party resources that may be referenced, accessible from, or linked to this document.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
TL3B/503
3039 Cornwallis Road
Research Triangle Park, NC 27709
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX 5L ^(R)	e-business on	Netfinity ^(R)	System i ^(TM)
AIX ^(R)	demand ^(R)	NetServer ^(TM)	System i5 ^(TM)
alphaWorks ^(R)	e (logo) server	Notes ^(R)	System p ^(TM)
AS/400 ^(R)	e (logo) business	OmniFind ^(TM)	System x ^(TM)
Asset ID ^(TM)	eServer ^(TM)	OS/400 ^(R)	System z ^(TM)
BladeCenter ^(R)	FlowMark ^(R)	PartnerWorld ^(R)	Tivoli ^(R)
Candle ^(R)	Holosofx ^(R)	POWER ^(TM)	Tivoli ^(R)
Chipkill ^(TM)	IBM ^(R)	Predictive Failure	Enterprise ^(TM)
CICS ^(R)	ibm.com ^(R)	Analysis ^(R)	TotalStorage ^(R)
ClearCase ^(R)	IBM Virtual	pSeries ^(R)	WebSphere ^(R)
Cloudscape ^(TM)	Innovation	QMF ^(TM)	Workplace ^(TM)
DB2 ^(R)	Center	QuickPlace ^(R)	xSeries ^(R)
DB2 Universal	i5/OS ^(R)	Rational ^(R)	zSeries ^(R)
Database ^(TM)	Informix ^(R)	Redbooks ^(R)	
developerWorks ^(R)	iNotes ^(TM)	RS/6000 ^(R)	
Domino ^(R)	Intelligent	Sametime ^(R)	
Domino Designer ^(R)	Miner ^(TM)	ServerAID ^(TM)	
DRDA ^(R)	iSeries ^(R)	ServerGuide ^(TM)	

DS4000^(TM)

Lotus^(R) SP^(TM)
Lotus Notes^(R)
Lotus Workflow^(TM)
MQSeries^(R)

IBM, the IBM logo, ibm.com and DB2 are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

The e-business logo is a trademark of IBM Corporation.

Rational is a trademark of International Business Machines Corporation and Rational Software Corporation in the United States, other countries, or both.

Adobe^(R), the Adobe logo, PostScript^(R), and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel^(R), Intel logo, Intel Inside^(R), Intel Inside logo, Intel Centrino^(R), Intel Centrino logo, Celeron^(R), Intel Xeon^(R), Intel SpeedStep^(R), Itanium^(R), and Pentium^(R) are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java^(TM) and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux^(R) is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft^(R), Windows^(R), Windows NT^(R), Windows Server^(R), PowerPoint^(R), Excel^(TM), Microsoft Internet Explorer 6, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX^(R) is a registered trademark of The Open Group in the United States and other countries.

Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries.

Oracle is a registered trademark of Oracle Corporation.

Red Hat is a registered trademark of Red Hat, Inc.

SUSE is a trademark of SUSE LINUX Products GmbH, a Novell business.

Other company, product, or service names may be trademarks or service marks of others.