

Oracle Collaboration Suite Email Migration

*An Oracle White Paper
June 2003*

Oracle Collaboration Suite Email Migration

Introduction	4
Migration Planning.....	5
Anticipated Email System Usage.....	5
Email Clients	5
Size of network connections to remote and local sites.....	5
Current configuration of deployed system.....	5
Hardware Consolidation.....	6
Number of registered mailboxes	6
Average size of registered mailboxes	6
Namespace Changes.....	6
System Configuration	7
CPU Usage.....	7
Hardware Consolidation.....	7
N-Tiered Configuration.....	7
Disk Configuration.....	7
Network Topology and Bandwidth	8
Availability Requirements.....	8
Message Transfer Agent and DNS.....	9
Message Transfer Agent	9
DNS.....	9
Directory Services Information.....	9
User Naming Conventions.....	9
Training and Communication.....	10
Administrative Training.....	10
End User Training and Communication.....	10
Email Migration Tool	10
Migration Tool Requirements.....	10
Memory.....	11
Disk Space	11
CPU Usage	11
Determine Migration Schedule.....	11
Migration Tasks	12
Migration Tool Task List.....	12
Tool Features.....	13
Migration Plugins	13
Setup Wizard.....	13
Account Creation	13

Data Migration and Pre-Migration	14
Routing	14
Notification	14
Post Migration tasks	14
Conclusion.....	15

Oracle Collaboration Suite Email Migration

INTRODUCTION

This document discusses the process and tools for migrating from an existing email system to Oracle Collaboration Suite. This includes analyzing the existing system, factors to consider in planning the new Oracle system, planning the migration steps, installation, and post-migration processes.

Email has established itself as one of the critical means of communication, particularly in corporate environments. Users typically retain much of their email for future reference and safekeeping – investigation of a typical mailbox will reveal messages with a variety of documents, presentations and pictures attached, as well as ordinary day to day email messages. Regardless of its real value, most of this will be regarded as critical information. While storage space for emails is often limited to the space allocated by the administrator of the email application, but 40MB to 100MB+ of email per user is not uncommon.

The challenge for a successful new system implementation is to ensure a smooth migration from the old system for end users, and this includes the ability to move data between the two. To preserve the working environment, this migration is not simply limited to mailbox information: it must also include other objects, such as directory information, distribution lists (private and public), and aliases.

Oracle Collaboration Suite provides a migration tool to do just that. It makes a connection to the legacy messaging application, creates and configures new accounts based on information extracted from the legacy system, notifies users of migration progress, and exports and imports all old mail and folders into Oracle Collaboration Suite. Efficiencies can be made, for example by consolidating multiple servers into one, or through space saving in the message store. Where other email applications give each individual a single copy of a message, migration to Oracle Collaboration Suite will result in only one copy of a distributed email (one message sent to many users) with a pointer for each recipient.

Migrating from one email system to another requires significant planning, however, complicated by the differing messaging needs of every installation. Migrating a large email system installation with minimal downtime is a challenging task. This paper looks at some of the common planning considerations, and the Migration Tool included with the product to facilitate the process.

MIGRATION PLANNING

It is important to understand the configuration of the current email system. Information from this system is important when planning the configuration of Collaboration Suite because it helps to understand the needs of the users, sizing the new system, and gauging its requirements. We will first look at some of the factors that will affect the final configuration.

Anticipated Email System Usage

Is email central to your business or is it a relatively lightly used means of communication? How many messages do you expect to transmit and receive each day? What is the average size of these messages? If email is heavily used or if users tend to send emails with large attachments (such as presentations or formatted papers), you may need to use larger servers to provide adequate performance. This is also important for evaluating your network bandwidth requirements, discussed later in this section.

Email Clients

Which email clients do you plan to use in your system? Oracle Collaboration Suite Email supports IMAP4, POP3, and provides a thin, web-based client that is accessible through a standard web browser. You may choose to let users access their email through any combination of these protocols. A POP3 system may permit a smaller back end configuration since messages are downloaded to the client machines. IMAP4 systems require more back end support because all messages are stored on the server rather than downloaded. The Oracle Webmail thin client also uses IMAP4, but because it is browser-based, requires a web server that accommodates the peak number of concurrent users.

Oracle Collaboration Suite allows you to separate these protocol servers into tiers apart from the message store to allow better manageability and offload resources where resources may be necessary.

Size of network connections to remote and local sites.

Obtain a map of the current network topology or work with a network administrator to create one. This will help in the determination of throughput for future email connections from local and remote sites.

Current configuration of deployed system

If the legacy systems will be deployed back into the destination Oracle Collaboration Suite configuration, they will determine part, if not all, of the platform of the new system. Be sure that the legacy platform is one that is certified.

Take an inventory of the current system configuration and match the results to the base SPEC CINT2000 website. (SPEC CINT2000 is a measure of processing power based on a system's ability to process integer calculations, and is one of the metrics used to evaluate system sizing.) This provides you with a picture of the size

of the systems supporting your current email needs and is a good guide for the processing power needed for the new system. For help understanding the performance requirements of Oracle Collaboration Suite, contact your Oracle customer support representative

Hardware Consolidation

One factor to consider when you plan to migrate your email system is the number of servers used to support the messaging system. One of the simplest ways to migrate your email system is to just move the users and associated information from each of your existing servers to corresponding servers with Oracle Collaboration Suite installed. However, email migration is a good time to shrink the amount of hardware utilized in your email system, if appropriate. Many email systems utilize many more servers than needed to support the number of users. This may be for support or administrative reasons, or because of network bandwidth issues. However, Oracle Collaboration Suite scales to support many users on a single system with centralized administration, which could lessen the number of servers you implement in your new system.

Number of registered mailboxes

Knowing the number of accounts is necessary to size the system configuration and calculate the migration duration. How many users are on the existing system? The new email system must be able to handle at least those existing users, preferably with some room to grow. Also consider the number of concurrent users on your system at peak load. If your users are spread across several time zones so that fewer users are on the system at the same time, you may not need as many servers as an organization with users that are all located in the same physical area.

Average size of registered mailboxes

This will determine how to size the database as well as the quota each user is given. A user quota is not necessary; however it does help to control the users' storage habits and force them to maintain their individual accounts.

Namespace Changes

Many companies implement a hierarchical namespace for their email domain. Some divide their email according to geographical location, for example locationa.company.com, locationb.company.com, and so on. Others set up their email systems according to organization, for example hr.company.com, finance.company.com, sales.company.com and so on. While sub domains are appropriate in certain situations, they can also create administrative difficulties, such as the necessity of changing a user's account name should she transfer from one location to another or from one organization to another. Flattening multiple sub domains into one domain (i.e. moving all users from sub domains to company.com) could ease this type of administrative task.

Oracle Collaboration Suite Email migration allows for the collapse of many domains to one during the migration. However, multi-domain installation is available for those installations wishing to maintain a hierarchical domain structure.

As an example, Oracle Corporation consolidated 97 email systems in 56 countries into one clustered system and collapsed its geographic sub domains into one oracle.com namespace when it migrated its email system.

SYSTEM CONFIGURATION

The size and configuration of the servers required for the new email system depend on the requirements above.

CPU Usage

Sizing requirements for each site depends heavily on the factors above, and are unique for every installation, but it is important to select hardware that will provide reasonable performance. Contact your local Oracle support representative to assist with analyzing your requirements for proper system configuration if needed.

Hardware Consolidation

If you plan to consolidate the number of email servers in the system, utilize servers that can handle the peak number of concurrent users that will be on the system at the same time. Servers can be used 24x7, so the same servers can be used for users across many time zones.

N-Tiered Configuration

Oracle Collaboration Suite allows you to configure your email system to have message store, protocol servers and clients on the same machine, or on separate machines. Having larger back-end messaging servers with several commodity protocol servers may be more cost efficient for certain installations.

The product is designed so that a site can add more back end servers should more message stores be required, or add more protocol servers should the number of users grow, or both

Disk Configuration

How much email do you expect to store? Looking at the size of the legacy email system should give you an idea of the disk storage required for your system. Imposing quotas on users is one way to limit growth of the message store, but remember that quotas that are too small or are too sharply reduced from previous limitations will result in reduced user productivity. Messages will still be received into the application, but users who are over their quota will not be able to retrieve their new messages until they have cleaned up their accounts.

Knowing the size of the user accounts will also help you determine the size of each batch of users to migrate.

Oracle Collaboration Suite should also have a disk configuration that meets other critical requirements such as disk mirroring, stripping, and disk partitioning, all of which will help in redundancy solution of your installation.

Network Topology and Bandwidth

Email requires good network connectivity between all sites involved. During the migration process, it is essential for all remote and local sites to have good throughput to the Oracle Collaboration Suite servers in order to move legacy information to the new system in a timely manner.

Email clients access messages in Oracle Collaboration Suite by sending requests to the product's IMAP4 or POP3 protocol servers. These IMAP4 and POP3 protocol servers use shared server processes (SPS) to service multiple clients with fewer database connections. SPS provides load balancing by using an Oracle9i Net Services listener to distribute the client load equally across multiple instances of the protocol server processes. Therefore, an analysis of network bandwidth and connectivity is necessary to provide reasonable user performance.

A number of metrics can be used to estimate bandwidth requirements. Email size is useful, but is typically highly variable. Another method is to measure peak day and hour email traffic load at several representative locations over a period of time, then divide the peak hour traffic by the total number of users. This helps determine the load capacity required of the SMTP server, and other middle tier components.

User profile	Bandwidth on Middle Tier	Bandwidth on Database Tier
ISP	2.9M/s	1.2M/s
Corporate	4M/s	1.8M/s

Table 1 Typical bandwidth requirements for email

(These numbers should be added to the values for the existing network usage.)

If you are planning to consolidate your email system onto fewer servers, remember to plan for increased throughput to fewer servers. Also consider the location of the new servers. Some areas of the world have limited network bandwidth, so you need to consider how to meet messaging needs of users there.

AVAILABILITY REQUIREMENTS

An enterprise or ISP solution may require 24x7 email accessibility. Smaller businesses may need not need to have such uptime and therefore can afford to take a system or application down for maintenance.

Determine the availability necessary for your site. Include backup and recovery as options for maintaining uptime, along with high availability solutions such as hot failover and Oracle 9i's Real Application Clusters.

MESSAGE TRANSFER AGENT AND DNS

Message Transfer Agent

Oracle Collaboration Suite includes a message transfer agent (MTA) to deliver and send messages over SMTP. Most UNIX operating systems come with an MTA called sendmail. In normal use, the Oracle MTA replaces any existing SMTP process, such as sendmail, on your server. It's also possible to use Oracle's MTA in a mixed environment alongside sendmail.

DNS

Every host in your email system uses Mail Exchanger (MX) records in DNS to determine the destination for each address.

The migration tool also makes use of the DNS records to redirect email to the new installation of Oracle Collaboration Suite. You need to plan when you will change the MX records to deliver messages to your new mail server rather than your old servers. Messages must be routed to the correct mailbox at all stages of the migration process, so that the migration process is as seamless as possible for end users. It's important to remember that not all users will be migrated at the same time, so until the process is complete some messages will be delivered to Collaboration Suite and some will still go to the old system.

Once the migration process is finished and all users are on the new Oracle Collaboration Suite system, you must configure the DNS entries so that the rest of the Internet will know to direct messages to the correct location.

DIRECTORY SERVICES INFORMATION

Oracle Collaboration Suite migrates all directory information to Oracle Internet Directory (OiD). All information is transferred to OiD, though synchronizations can be setup to integrate this with other sources of directory information.

User Naming Conventions

It is important to determine the naming convention of your migrated users. There may already be an internal standard such as user_x or user_name. The Migration Tool supports a variety of automatic naming rules for the automatic creation of user accounts, including: first.last; initial_last; first_last; etc.

This is a good time to evaluate whether to keep your old naming convention. Keeping the existing naming convention will shorten migration tasks. If you decide to change your naming conventions, consider how this will affect ongoing operation. It may be necessary to create aliases or simple rewriting rules to ensure that messages sent to the old address formats are not lost during a transitional period.

Failing to redirect messages in some way means email sent to the old addresses (e.g. via user-created aliases, distribution lists, or from customers) would bounce.

TRAINING AND COMMUNICATION

When any new application is introduced there will be a learning curve for users and administrators. Training for users, support desk personnel, and administrators eases the transition to the new system. Continued communication with end users throughout the process also improves the overall migration.

Administrative Training

Administrative training may be necessary to improve maintenance of the Oracle Collaboration Suite installation, just as with any new product. While administrators will already be familiar with email technology such as SMTP and IMAP, specific product training will be necessary. For Collaboration Suite, this may also benefit from the inclusion of database administration courses.

End User Training and Communication

One of the most important tasks during the migration process is constant communication with users. Set up training sessions for your users to get them familiar with the new application. If you have chosen to use the Webmail client, training on the features of Webmail essential for success. Since the product makes use of standard email protocols, the impact on end users may be small as they may be able to continue working with their existing mail client. Often it's helpful to create a website with an FAQ for end users, reducing the number of calls to the Support helpdesk and the administrators. Regular communication describing the changes being made, what users need to do, and when, will greatly reduce the impact of the migration on day-to-day operations.

EMAIL MIGRATION TOOL

In this section we'll describe some of the capabilities of the Migration Tool, and the tasks to be performed to move data to Collaboration Suite.

Migration Tool Requirements

The Oracle Collaboration Suite Email Migration Tool automates the movement of data from the old system. This section describes the design objectives of the tool, and the system requirements for the machine on which the Migration Tool is running.

The tool is designed specifically to handle the requirements for:

- Corporate installations with significant numbers of mailboxes and large message stores.
- Internet Service Providers that generally have a large number of users (millions) with small mailboxes (a few hundred kilobytes)

The tool provides features to manage the migration process, and minimize system downtime.

Memory

The minimum memory requirement is 512 MB of RAM.

Disk Space

The migration tool disk space requirements are:

- 300 MB for migrating 10 users concurrently from MBOX, assuming a quota of 20 MB per user (including all the intermediate store files)
- 70 MB for migrating from IMAP

Since the export of data needs disk space for downloading to file, the requirement for the disk configuration will be based on the size of legacy email accounts. By understanding the current usage of users, you can estimate how many users to batch at a time and calculate the disk space necessary during the migration process.

CPU Usage

The Migration Tool can migrate between 600Mb and 2.4Gb per hour, depending on system specification, message profile, and network performance.

When migrating from your legacy email system, it is best to run the Migration Tool on a separate box to avoid excessive impact on your production mail system.

Determine Migration Schedule

When planning the time required for your email migration, be sure to plan time for installation and testing of the Oracle Collaboration Suite installation prior to beginning the migration process.




Figure 1 Scheduling Migration Batches

The tool allows you to group users into batches of a fixed number, and schedule the execution of these batches, giving you close control of when to migrate users. Typically, a few users will be migrated initially to validate the process, and establish a small pilot group of users. Once this is successful, larger groups of users can be migrated as a background or overnight task under the control of the tool scheduling features.

This first batch of users to be migrated will validate your preparation and allow time to correct any problems or concerns that may arise before migrating the major percentage of your users.

MIGRATION TASKS

Email migration is an intricate and time sensitive process. It requires migrating accounts, preferences, rules, folders, email messages, distribution lists, shared folders, and public aliases, as well as changing the SMTP routing. Migrating a large email system installation with minimal downtime is a challenging task.

Migration Tool Task List

The *Oracle Collaboration Suite Email Migration Tool Guide* provides information about tasks performed by the Migration Tool. It's recommended that this guide be studied before attempting any of the migration tasks below.

The tasks to be performed are:

- Configure migration plugin to connect to the old system and extract data
- Run "Migration Setup" wizard from the tool
- Generate User List file from source system
- Load users into the context of the tool

- Create Accounts on Oracle
- Migrate distribution lists and aliases to Collaboration Suite
- Create Batches in the Migration Process
- Schedule the Batches to stage the migration of users' data
- Start Migration of data

Tool Features

The migration tool was designed with the challenges of large enterprise wide migrations in mind. Here are a few of the specific features included that address this:

Migration Plugins

The modular architecture of the migration tool means that the migration process is the same, regardless of the source mail system. Plugins contain the native code required to connect to the legacy system and extract information to be migrated. So, although the tool can be installed in any Java environment, plugins sometimes have platform dependencies imposed due to that mail system.

Setup Wizard

When first invoked, the tool's Wizard is invoked to configure the various parameters that control the migration process. This information is maintained within Collaboration Suite so that the status of the migration process, and the status of each user under migration are managed through the lifecycle of the migration.



Figure 2 Migration Setup Wizard

Account Creation

The tool creates all user directory objects in Oracle Internet Directory before moving email content. No accounts or mail services need to be setup before

migration. Mailboxes are created based on the rules specified during the Setup Wizard, and passwords are assigned to new accounts. This information is then sent to the end user so that they can connect to the new system.

Data Migration and Pre-Migration

The challenge faced by any mail migration process is how to cope with the large amounts of data involved, without losing any, or without causing unreasonable interruptions to service. Even with an extremely fast application, hundreds or thousands of Gigabytes of information will take time to move. During this time more email will arrive and possibly be missed by the migration. The Migration Tool works around this problem by providing a Pre-Migration feature. This task can be run repeatedly without affecting the current live system, each time it runs it identifies only those messages that have yet to be migrated to the Oracle message store. So, with repeated execution, the amount of data moved, and therefore time required reduces. This means that the final cutover when, for safety and so that no new messages are lost, the old mailbox is temporarily unavailable, is a short time that is not dependent on the total amount of data that has actually been migrated.

Routing

As we've already noted, the mail system needs to remain operational throughout the migration process. For very large organizations, this means that for a period of time some users will be working on the old system, and some will be moved to the new one. The Migration Tool facilitates the routing of messages between the 2 systems so that end users can work normally.

Notification

Communication is very important throughout a migration. The tool helps here by sending automated email notifications to users as their mailboxes are processed. New account name and password information is sent, as well as status of their migration so that users know when to move to the new system.

This notification is also customizable, so you can also add site-specific information – such as helpdesk support contact information – tailored to your environment.

Post Migration tasks

In all but the smallest of installations, old and new systems will continue to operate in parallel until migration is complete. In most cases, all external mail will continue to be delivered to the old system during this period, and so the MX record for your mail domain will be unchanged.

Once migration is complete, then mail delivery should be directed at Collaboration Suite, and this means that the DNS MX record for your domain must be updated. Since DNS is a dynamic process, this can take a few days to propagate globally, so to be safe it's probably wise to keep both systems running for a period after the

MX record is formally changed. Once you're satisfied that this change is effective, the old mail system can be decommissioned.

CONCLUSION

The migration of legacy email servers into Oracle Collaboration Suite can show many benefits – including hardware consolidation, reduced labor, and reduced maintenance costs – in addition to the benefits gained from deploying a single architecture for all collaborative applications. A simpler, centrally administered system that supports standards-based clients can ease administration without locking users into proprietary software. With careful planning and the tools provided, your Oracle Collaboration Suite migration should run successfully.



White Paper Title
June 2003
Author: Paul Nock
Contributing Authors: Ricardo Rivera

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Copyright © 2003, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.