Multi-Forest Exchange Consolidations

Mindset and Best Practices - Priasoft

Multi-Forest Migration?

Why discuss a Multi-Forest Migration?

In many larger organizations there exists a state where users, data, applications, and systems are split between multiple Active Directory Forests. This is often the result of growth over time by the acquisition of other existing organizations. At some point there is a desire or demand to consolidate the multiple Exchange environments into a single, unified mail system. The purpose of this white-paper is to provide information to help develop the best mindset and approach to such a project. This document will also provide guidance on some of the more common issues and some of the not-so-common and more subtle issues that occur with a multi-forest consolidation project. Priasoft has over 13 years of experience with Exchange-to-Exchange migrations and in that time has helped many customers migrate from a multi-forest environment.

Mindset

The mindset to have when approaching a migration of this type, even if it is only for a few hundred users, is to think broadly about user activities as it relates to Exchange. As a reader of this document you are likely a user of Microsoft Outlook or at least a user of Microsoft Exchange in some way. Consider the many various activities you perform as it relates to email and Exchange. Here are some common user activities:

1. Sending and receiving email
2. Browsing the Address Book
3. Using Distribution Lists
4. Using shared or conference room mailboxes
5. Opening mail or calendar folders of other users

However, these are only the common activities. There are some other subtle or less used features of Exchange to think about as well:

1. Quotas and Limits
2. Retention Policies (mailbox and deleted items)
3. Archiving
4. Applications that use Exchange like SharePoint and Dynamics or others
5. Remote user activities and setup like Outlook Anywhere and Cached Mode

When approaching a multi-forest Exchange migration, it is wise to consider as much of the possible points of user interaction as possible before starting a design. It is often easiest to start with yourself and imagine all the possible things you might do or have done with Exchange. However, this is often not enough. Exchange migrations are an infrequent occurrence in most businesses and as such there is often a state of "you don't know what you don't know". Priasoft is here to help. Beyond the topics discussed in this document, we are always willing to discuss a migration plan and to act as an advisor and sounding board for migration designs and possible solutions to technical and/or business challenges.

Multi-Forest Challenges

The subtleties increase with a multi-forest Exchange environment. There are often applications and services that exist solely to support the multi-forest environment and those same applications and services will either be dropped or will need to be changed to support a single consolidated environment. The cross-forest relationships and inter-forest activities cause most enterprises to implement some sort of solution to make life easy for both administrators and end users. Often this means third party solutions, but also just as likely are home-grown scripts and undocumented manual processes. The following sections will attempt to call out the most important ideas and will highlight the challenge and provide some guidance on how approach a solution.

Cross-Forest Representation

In a multi-forest Exchange environment, there will often be the case of a mailbox user in one forest having a related object in one or all of the other forests. The related object could be a Contact, Mail-Enabled User, or Mailbox-Enabled User. The creation of this related object could be done by many means: an automatic directory sync, a manually operated directory sync, scripts, or a completely ad-hoc manual process. It is just as important to know how the object was created as it is to know that it exists.

The importance of considering cross-forest representation has to do with one of the most common user activities: sending and receiving email. The critical factor to consider is reply-ability. Exchange is an X500 based service with the meaning that routing of email between mail related objects is done internally via X500 addresses. This value has the format of `/o=OrgName/ou=AdminGroup/cn=Recipients/cn=mailboxname`. This value is rarely seen by end users but is one of the most important values in an Exchange environment. Each object that participates with Exchange has an X500 value and is stored on Active Directory objects in the LDAP value named 'legacyExchangeDN'. Mailbox users have one, Contacts have one, Mail-Enabled Groups and Distribution Lists have one, Mail-Enabled Public folders have one, and even Exchange Servers and Databases have X500 addresses.

When a user receives an email, Exchange stores the X500 value of the sender and each recipient in the message (not an SMTP address). Later, if the same user replies to the email, that user is replying to an X500 address. The sender of the original message is actually stored as an Address Book pointer of which the X500 address is the unique identifier in AD for the Exchange related object.

Now consider again the idea of a mailbox user in one forest having a representing object in the other forests. Each of those related objects will have its own X500 value, unique to the forest within which it resides. Further, consider a case where User1 receives mail

from User2 and they are in separate forest. Each user has a Contact representing the other in their respective forest. If User1 is migrated to another forest and attempts to reply to User2, it will likely fail unless additional work was done. The reason for the reply failure is because there is no object in the new forest with X500 address of the Contact object of User2.

When consolidating multiple Exchange forests into another one, great consideration must be given to ensure that all X500 addresses representing all other objects across ALL the other forests are blended together in the target forest. Failing to do this work increases the chances of a non-delivery. The final result would mean mailboxes in the target environment with additional X500 addresses, one for each source forest. If the source environments total 4 source forests, this means that target objects will have up to 4 additional X500 addresses, provided that there is a full mesh of representing objects.

Reply-ability is not the only concern for which X500 addresses resolve. This value is used for ANY lookup used by Exchange, not just the transport. Folder delegation, Rules, and Outlook's AutoComplete cache all use the X500 value. There may also be applications or services attached to Exchange in the target environment that use this value as well. Consider common services like Archiving, SharePoint, and Help Desk applications. There may even be some custom applications that depend directly or indirectly on the X500 value.

Furthermore, there is more to consider than just a mailbox-to-contact relationship. There are likely similar relationships between contacts, DLs, and mail-enabled Public Folders. Shared Names (discussed in the next topic) create a need to merge data, which sometimes is slightly different between forests, together in the new environment. In other cases, there may be Contacts in one forest for DLs in another. Those relationships and the critical values (like X500) also must be preserved.

The solution to this challenge is to prepare the target forest with all the mail-enabled objects from each source environment before the first mailbox is migrated. This ensures that, regardless of who migrates first, migrated users will have a proper and working Address Book in the target environment. "Proper" in this context means that the target environment has mail-enabled objects representing all the external objects from all the other forests. "Working", in this context, means that if a migrated user sends or replies to mail, the item will be forwarded to the appropriate object in one of the other forests. It further means that general expectations are met with regards to the appearance of objects in the Address Book and that sending to a Distribution List properly delivers to its members.

Priasoft provides software to address this challenge in the form of a directory sync solution. Our solution is built specifically to support a migration is has capabilities and features that are specific to Exchange and email. In contrast, other more general directory sync solutions often require much scripting or complicated configuration and rules in order to accomplish the same result. Furthermore, when a business process issue is discovered, it can be difficult to reconfigure these tools to address the issue. The choice of solution is important, but Priasoft does not force any customer to use our tools and choosing not to will not impact Priasoft's ability to migrate mailboxes – but we do make it easy.

Cross-Forest Conflicts and Ambiguity

Another important consideration of a multi-forest consolidation is object conflicts. There are potentially many categories of conflict and some are business issues while some others are technical issues. Be sure to know the difference. The mindset to have with regards to conflicts is to fully understand that multiple, independent Exchange environments will need to blend together in a single target environment.

Priasoft recommends that for any case where 2 or more existing Exchange forests need to be blended together that such is done to a new, empty target forest. Doing so allows for discovery of conflicts and ambiguity without impacting any user in any forest. Sometimes it is very difficult to understand and track conflicts until it is seen in a live environment, in this case the new 3rd environment. If 2 organizations are blended into an existing 3rd environment, there is a chance of impacting the users that exist in the 3rd environment. Preparing and pre-staging a new target environment allows architects and administrators to validate the environment before any user would use it.

Shared Names

This topic is fairly easy to understand. There is a chance that 2 users exist with the exact same name, but in different forests. John Smith (in Forest A) and John Smith (in Forest B). The existing business units have likely already dealt with this in their own way, but only in the context of the forest in which they operate. For instance, John Smith in Forest A might have a contact in Forest B of John Smith (XYZ Inc.) – the descriptor in parentheses identifies him in Forest B's Address Book, but in Forest A, the same user only appears as "John Smith". A business use issue exists here with regards to how the final Address Book entry should appear. From a technical standpoint, the new object in the target forest could have any text value but it is up to the business leaders to provide some direction and possibly some rules to follow in this case. Sometimes it ends up where every user in the target Address Book should have a suffix in parentheses to identify them.

A more subtle issue with shared names are non-user objects like shared mailboxes and distribution lists. Consider a case where there is a shared mailbox named "Sales" in Forest A and a shared mailbox also named "Sales" in Forest B. It is unlikely that the same rules from above (the John Smith case) would necessarily apply here. In some cases the 2 sales mailboxes should be merged together so that the business has a single "Sales" mailbox (often for public use like sales@company.com). In other cases, it may be critically important to keep the mailboxes separated. In this scenario special handling is required to avoid ambiguity in the target environment. A common solution for such cases is a prefix or suffix in the display name and on other attributes as well a logical separation in AD by placing the 2 objects in separate OUs.

The same concern and consideration should be given to Distribution Lists as well. The previous paragraph is just as valid if all instances of "shared mailbox" were replaced with "distribution list". However, distribution lists have a few other subtleties that mailboxes do not. DLs by their very nature are a container of other objects, they are not the final destination. As such, it can be even more important to have unique display names for otherwise ambiguous DLs (consider a DL named Support). When a user receives mail by being a member of a DL, and provided the user and the DL are in the same forest, the user will see the display name of the DL in the recipient list. If the business decides to keep 2 DLs with the same display name, this can be confusing to end users and increases the chance of users sending to the wrong DL when selecting from the Address Book.

Another subtlety of ambiguous DLs relate to owner managed DLs. If the environment allows end users to manage DL membership via Outlook, a user could accidentally manage members of the wrong group or could generate support tickets claiming that his/her group has the wrong members.

Merging DLs can be done as well, but doing so means that members are now merged as well from the different source environments. For an email-only DL, versus a mail-enabled security group, this is fairly safe but users might not expect some of the new emails they receive when a user from Forest A sends to the DL and all the Forest B users (now part of the DL) get the email when previously they never did.

Shared names are not exclusive to users or groups. Consider that each source forest is a separate ecosystem. Each one likely has similar business patterns and therefore likely uses IT in similar ways. There can be shared names across any of the Exchange object types: Mailbox, Contact, Mail-User, Public Folder, Distribution List, Server, Database, etc. It is important to identify shared names

and then to identify if the issue is a technical issue or a business issue. The goal is to properly plan for reply-ability and object lookups by preserving all source X500 addresses in some way.

Cross-object Conflicts

The above scenarios all assume that the conflicting objects are of the same type. However, this is not always the case. It is not uncommon for one business to use a shared mailbox and another to use a distribution list for the same purpose. Consider a case where there is a shared mailbox in Forest A named Marketing. There may be a DL in Forest B named Marketing as well. Their purpose is the same in both companies: to provide a public email address for the marketing department. Forest A may use 'marketing@cool.com' and Forest B may use 'marketing@awesome.com'. This is technically legal since they each have a different SMTP address. However, they likely share other attributes such as the mail alias, directory object name, and NT Logon Name (aka samAccountName in LDAP). It is these other shared attributes that can and do cause issues. It is unwise to assume that any of the shared attributes can simply be "adjusted" in the new forest without consequence. Although it works from a technical point of view, the downstream impact can be unknown and possibly unknown for many days, weeks, or months. Many times applications, including Exchange itself, are periodic in nature; only performing some action every so often. It is these applications that may be depending on the mail alias or samAccountName – arbitrarily changing values just to support uniqueness can create issues that can be quite difficult to discover.

In the case of a cross-object conflict, this is almost exclusively a business issue and business leaders/owners should be informed and decisions made between them on how to handle the issue. IT can provide options to help them along, but the decision is likely not IT's to make.

Authoritative Source of Data

Blending several existing environments together also means that a discussion about attributes and metadata. Prior to consolidation – in most cases – each environment owns its own changes and is authoritative for its data. Directory synchronization tools may already exist to replicate changes, but typically each environment is the "source" for changes for which they own; like mailboxes and DLs. Even without synchronization tools, there may still be some cases of cross-forest representation but done by manual means.

When the new environment is prepared and pre-staged with the information and objects from the source environments, there is a natural separation of context to consider. In a source environment there will still be user accounts with mailboxes, contacts, DLs, and the rest of the Exchange object types. The Address Book that users see is from the source environment and users are very sensitive to changes to the way the Address Book appears. As users are migrated into the target environment their expectation is to see the same detail as they saw before the migration. If someone's name changes (marriage perhaps), this change may be initialized in the source environment from HR. From there it may manually or automatically update the source user account. This change should then flow to the target environment so that the Address Book shows the proper data. The authoritative source for the user's name appears to be the source environment.

However, there are other attributes that are very email specific such as SMTP addresses. Early decisions should be made as to whether each source environment will be authoritative for mail related attributes or if the target environment will be. This leads to the topic of Move-Add-Change (MAC) processes. It is important to understand and document the current MAC processes and how they affect directory object attributes and the Address Book. This will help identify likely authority of data. Additionally, there is often a need to develop new MAC processes: one that is used during the migration period and one that is used after all migrations have completed across all environments. It is probably unwise to continue to have new employees have new mailboxes in the

source environment only to have to migrate them. These new MAC processes may alter the authoritative source of user information.

Directory synchronization then becomes more important between the source environments and the target. Just as important however is understanding the direction of synchronization, but direction cannot be determined without first establishing the authoritative source.

Priasoft highly recommends that all parties come together to establish the rules for the MAC processes and for describing authoritative sources of data. It is further recommended to avoid support scenarios where changes to information are allowed in both the source and target environments mutually. It is much easier to troubleshoot and manage a single direction of data sync than a case where data can change in both directions and for the same attributes.

Holistic versus Partial Consolidation

Often in larger enterprises when the decision is made to consolidate multiple environments, one or more of the source environments may be restricted to only the migration some of that organization into the new environment, effectively splitting that organization into 2 parts. There might be regional or country specific reasons for this, or it may be that there is a legal separation of business units and one of the business units will remain in the current environment through and after the migration.

A partial migration of an environment imposes additional complexity and concern. As discussed earlier, there is likely already cross-forest representation of local objects (mailbox-to-contact relationships). This idea will then likely need to continue, but now between the partial environment and the new target. The users in this source environment that will not migrate are familiar with the view of the Address Book as it is now, however if additional work is not done, the address book will look different to those users as mailboxes are migrated. Priasoft has built-in support for maintaining the source Address Book in a workable state so that users are not dramatically affected. However, in a longer term view, there may be need to implement a directory sync between this partial source environment and the new target. One of the many complexities to examine is whether ALL of the target mail objects - remembering that the target environment will be a blend of all the other source environments - should be represented in this source environment, or only those source items that were originally there. However, careful consideration must be given here as there could be policy, legal, or government regulation that determines or influences whether such a sync is allowed or for how long. Priasoft cautions customers to include their respective legal and compliance teams to have a voice early in the migration plan development to keep everyone out of trouble later.

Beyond just the way the Address Book appears in a partial multi-forest migration is the potential functional impact to the users in the source environment that will not migrate. Often the only consideration made is for the users and objects being migrated, but there is an impact to those who are not migrated as well. Consider the previous topic about reply-ability and X500 addresses. The source users that don't migrate can potentially have issues replying to previous mail from migrated users the same as anyone else. If there are legal, policy, or other reasons which prevent the creation of co-existence objects in the partial source environment, well written communication should be given to those users that will not migrate of the impact: reply issues, shared mailbox issues, delegate issues, etc. Understand that there very well could be a case where users that migrate and users that do not migrate both shared a particular mailbox (like Sales, or Support). If such a shared mailbox is migrated normally, non-migrated users will not have access to the shared mailbox. If the shared mailbox is re-enabled, there is now the consideration of the content: should content created or received in the source environment be forwarded to the target mailbox also? Should content created or received in the target environment be forwarded to the source mailbox? These questions are only a couple of possibly many questions to ask, but

the decision is ultimately a business decision, and the importance of the questions are typically tied to the importance of the mailbox and its data. Shared mailboxes that influence or impact revenue streams are often very important.

Solutions

Priasoft Migration Workshop

Priasoft offers a 4-day, on-site workshop that focuses on migration design and approach. The workshop addresses technical and non-technical influencers that affect the success of a migration. The workshop helps draw out the non-obvious and subtle influencers such as end-user activities and help desk messaging. In a multi-forest consolidation, the workshop is extremely valuable since all of the topics of this article are covered but also many other topics that have not been discussed. The workshop will help develop a proper plan that includes consideration for ALL of the business units affected. One of the most valuable deliverables of the workshop is the Event Based Work Breakdown Structure. This dependency driven timeline shows how one task or event affects all other tasks and ensures that no important topic or issue is missed. The diagram also provides a tool that can be used to accurately forecast the overall project duration.

Priasoft Collaboration Suite

The Priasoft Collaboration Suite is a valuable solution for preparing a target environment. This component is easiest to describe as a directory synchronization solution. However, unlike more generic dir-sync solutions, Priasoft's tools are Exchange and migration focused. This means that the tools have built-in capability to support Exchange while more generic tools (like FIM, ILM, etc.) require coding and/or scripting to work properly with Exchange.

In a multi-forest migration the blending of source environments into a target is a very important first step. Proper tools make this work easy and ensure that as the business changes due to normal operations that those changes are reflected in the target environment. Furthermore, taking the time to prepare a target environment with a fully blended Address Book ensures that regardless of the first user that migrates, he or she will be able to use the new blended address book. Often, just having a unified address book is a milestone in such a project and is something that senior management and stakeholders can and like to promote.

Priasoft Mailbox Migration Manager

The Mailbox Migration Manager (MMM) is the component that migrates mailboxes and offers the ability to do so in a "dry-run". The dry-run feature is critical to the success of the project since it provides the mechanism that exposes issues before they affect end users. Any migration plan that does not have a dry-run component means unnecessary acceptance of risk since the only time issues will be discovered is at the time of the production migration – a terrible time to discover issues.

Priasoft Outlook Profile Update Manager

The Profile Manager is the component that updates user's Outlook settings to point it to the new migrated mailbox. The simplicity of this component allows for easy automation by any solution that an enterprise already has in place, such as Group Policies or a Desktop Automation solution. The importance of this component is undeniable as without it users will be forced to have new "profiles" created which greatly impacts the end-user experience (even the smallest change over hundreds of users is very painful). Updating an existing profile means that users leave on Friday and return on Monday without any obvious change to the Outlook experience.