The **Experts** in
Exchange Migration™

THE TOP 10 OFFICE 365 MIGRATION CHALLENGES

Explore the top challenges and risks you will face during an Exchange migration to Office 365. This is a must read for all IT professionals that are planning or conducting a migration project to Office 365.

PRIASOFT™

Introduction

There are many benefits to moving to Office 365 but it is critically important that you evaluate both the current and future needs of your business. Switching is never easy. The experts at Priasoft have developed this helpful guidebook to discuss the main challenges that you will be presented when migrating your Exchange email to Office 365. It will help you understand risks, potential issues and technology impact of navigating a complex migration project while providing insight on designing a successful migration process for your business.

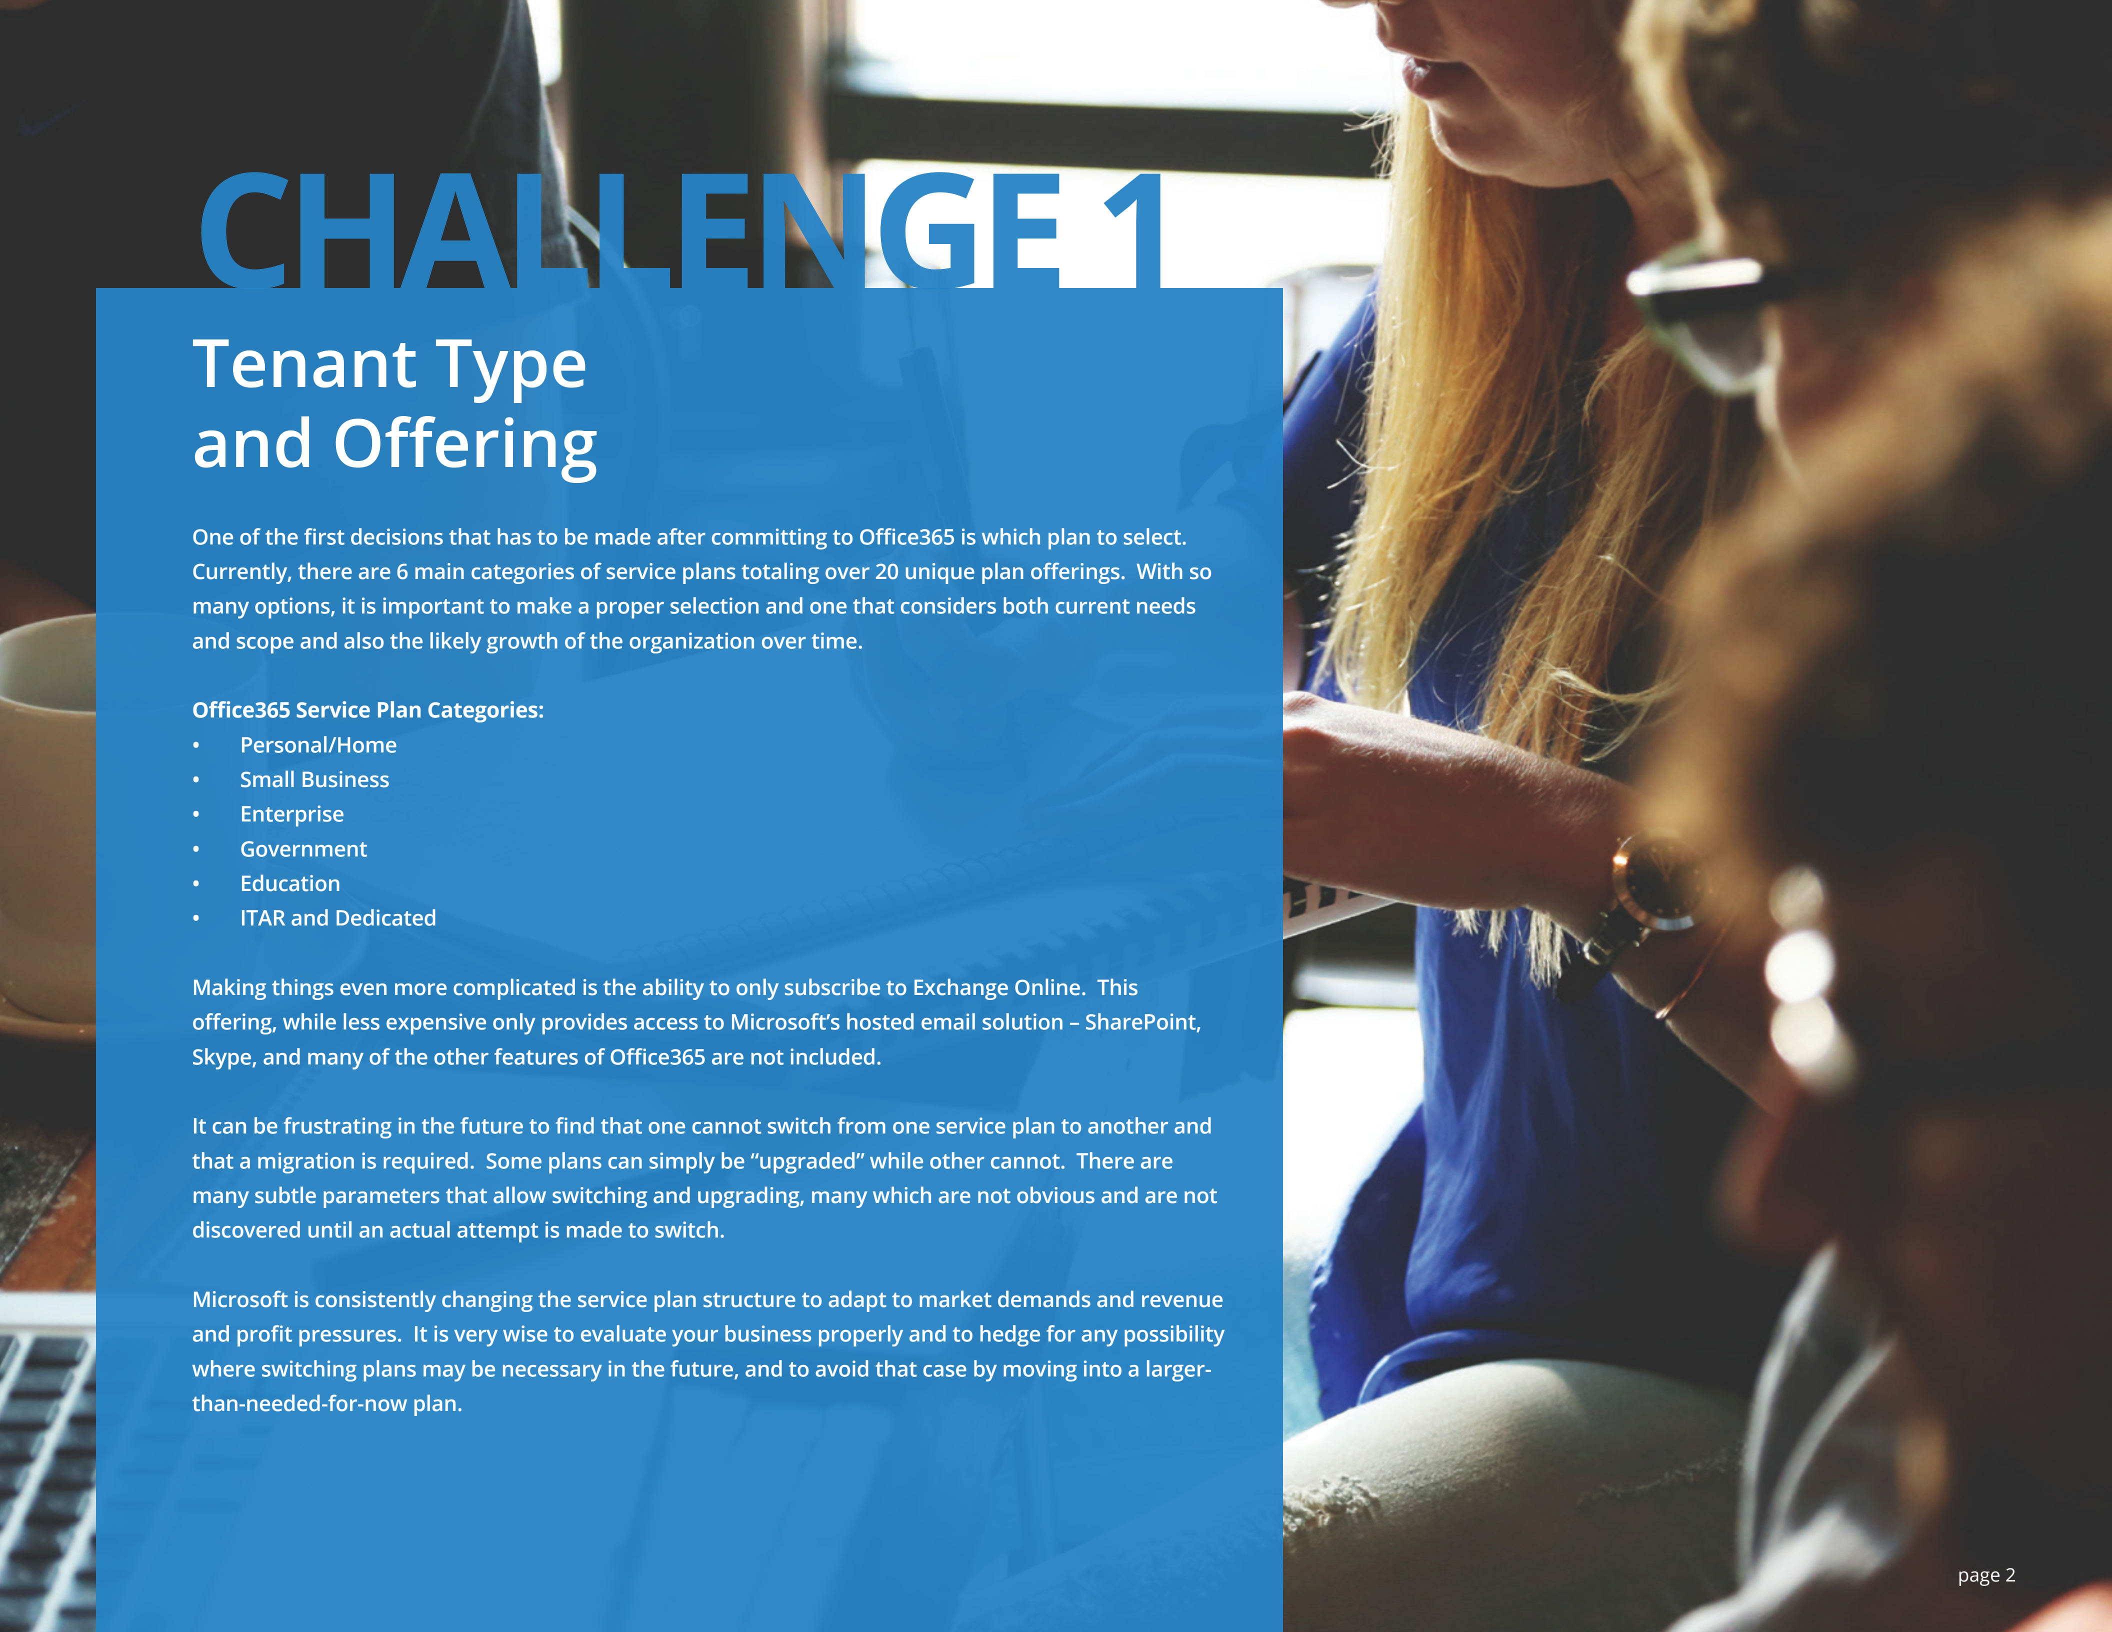

CHALLENGE 1

Tenant Type and Offering

One of the first decisions that has to be made after committing to Office365 is which plan to select. Currently, there are 6 main categories of service plans totaling over 20 unique plan offerings. With so many options, it is important to make a proper selection and one that considers both current needs and scope and also the likely growth of the organization over time.

Office365 Service Plan Categories:

- Personal/Home
- Small Business
- Enterprise
- Government
- Education
- ITAR and Dedicated

Making things even more complicated is the ability to only subscribe to Exchange Online. This offering, while less expensive only provides access to Microsoft's hosted email solution – SharePoint, Skype, and many of the other features of Office365 are not included.

It can be frustrating in the future to find that one cannot switch from one service plan to another and that a migration is required. Some plans can simply be "upgraded" while other cannot. There are many subtle parameters that allow switching and upgrading, many which are not obvious and are not discovered until an actual attempt is made to switch.

Microsoft is consistently changing the service plan structure to adapt to market demands and revenue and profit pressures. It is very wise to evaluate your business properly and to hedge for any possibility where switching plans may be necessary in the future, and to avoid that case by moving into a larger-than-needed-for-now plan.

CHALLENGE 1

Tenant Type and Offering

Personal and Home Plans

These plans are for individuals, independent professionals, and small businesses that consist of only a few people. There are versions of these plans that can include the full Office product suite (Outlook, Word, Excel, etc.) or can be "web only" versions. There is no ability to simply "switch" from these plans to any business plan.

Small Business Plans

These plans are for small businesses ranging from a few dozen employees up to 300 employees (as of 2015). There are several plans offerings in this category with different software and service options. Not all offerings in this category include Microsoft Exchange.

Enterprise Plans

These plans are for mid and large size organizations. They support an unlimited number of users, hybrid-mode deployments, and most enterprise-class features that would be expected.

Government Plans

There also exist several offerings specifically tailored for Federal, State, and Local Government institutions and agencies. These plans offer certified compliance with specific rules and laws and ensure that data is not stored out of the country in any way. There are sometimes the need to make specific requests for certain compliance, like CJIS or HIPPA in order to get a certificate that proves the entity is compliant. The data for government entities is segmented away from the normal multi-tenant data of the general Office365 environment which ensures that government data is in no way commingled with non-government data.

Education Plans

Office365 has separate and distinct offerings from education institutions and organizations. The value of some of these plans is reduced cost, free mailboxes for students and alumni, and education focused tools and support. As with all the other categories, this plan has several offerings for different size entities.

ITAR and Dedicated Plans

These plans are not discussed very often and Microsoft does not actively promote these plans as they are very specialized and niche. However, for some organizations, especially ITAR organizations, this may be the only option if you are looking for a trusted Cloud platform. Microsoft has had dedicated hosting plans for many years (it used to be called BPOS-D). Dedicated plans are often more expensive as they entail the use of completely segregated computing resources such that data in a dedicated platform does not commingle with the multi-tenant offerings. Dedicated plans have better support for customization and approved 3rd party integrations. Due to the costs, these plans are often only justifiable for very large organizations.

Unseen Challenges

As an organization grows, or shrinks if that is the case, there may be a need to change the Office365 plan to better accommodate the size of the organization or to better control the spending on the service. The challenge is that, while there is some ability to mix different offerings within a category, there are partitions between the different categories. These partitions create a situation that requires a migration in order to change plans. Microsoft doesn't currently provide support for cross-tenant migration tasks. Changing categories, for example from a small business plan to an enterprise plan, is much like changing to a new hosting provider and all the complexities that go with it.

This is especially important for the small business plans that limit user counts to 300 or less. If an organization, at the time of selection of the plan, choose the small business plan and later acquires another company that would cause them to exceed that limit, 2 migrations would be necessary: one from the small business plan to an enterprise plan, and another migration from the acquired company to the enterprise plan. This challenge is exacerbated if the acquired company is also on Office365 due to lack of tools in the industry for tenant-to-tenant migrations.

CHALLENGE 2

Geo-location of Data

The next topic of importance with considering a migration to Office365 is “where will the data be located?” This can have subtle and obvious consequences. By the nature of Microsoft’s Cloud offering, a customer’s primary datacenter will be based on information provide during the signup process.

The challenge here can be that if your organization is highly distributed, across many time zones or countries, some users may get a better experience than others. Microsoft does not currently support any model that lets an administrator have a tenant span multiple geographic locations. While it is true that data protection exists in Office365 by replicating to 2 or more other datacenters, that only exists for failover cases. A tenant admin doesn’t have the ability to dictate or influence the redundancy and users will not connect to those replicated datacenters.

The issue is most particularly felt when users span countries and the physical distance alone causes high latency. Outlook in particular, even with cached-mode enabled, behaves differently in a highly latent setting than it does when on LAN or low latent networks.

This issue is further complicated if your organization has any governmental requirement produce and keep data “in country”. A distributed organization would then likely require multiple Office365 tenants, one for each country that has the requirement. This setup immediately leads to complexities and collaboration issues that are not apparent. Imagine needing to provide access to a shared mailbox between tenants.

The determination of how a client connects to Office365 is a bit of fuzzy science. When a request is made, for example to ‘outlook.office365.com’, a DNS query is made to get an IP address for the name. However, since the name ‘outlook.office365.com’ is a generic name used world-wide, it is important that the IP address returned for the name be a close to the requester as possible. Microsoft has implemented a complex system that, in most cases, the results of a DNS query are returned with the most likely data center IP addresses.

CHALLENGE 2

Geo-location of Data

This is not a strict science because there is not a central registry of IP addresses and their physical location on the earth. One could buy or lease a few public IPs from an ISP and the "record of ownership" might list the location, but there's nothing that requires that those IP addresses be used in at that location in the record – they can be used anywhere in the world. In the end, admins and architects should spend time to see what IPs and CNAME records are returned by DNS in their local region for the many common Office365 hostnames. Some common names are:

- Outlook.office365.com (used by Outlook)
- Autodiscover.outlook.com and autodiscover-s.outlook.com (used by many tools and applications)
- Portal.office.com (used by web browsers)
- Ps.outlook.com (used by PowerShell)
- ProvisioningApi.microsoftonline.com (used by AzureAD powershell)

The exercise may show that inconsistent results are returned, such that each new DNS query responds with a different regional endpoint. If this is not corrected, this can lead to sporadic performance results where in one operation things are responsive and smooth, but in another are very slow and faced with timeouts. Priasoft maintains list of world-wide, dynamically discovered regional endpoints for 'outlook.office365.com' that you can view here at Worldwide Office365 Endpoints.

When a lookup is performed for 'outlook.office365.com', the public DNS server on the Internet that makes the query is used as a hint to determine a regional specific CNAME record for the generic 'outlook.office365.com' name. Ideally this should place the connection to the nearest physical Microsoft datacenter. However, if the public DNS server is unknown by Microsoft, or if the exit to Internet for a company is through a private connection to another location, the CNAME may not be the closest. For example, a query of 'outlook.office365.com' may return 'outlook-emeasouth.office365.com'. However, if the physical location of the client or server making the connection is in the USA, this is a horrible result. These statements are being made here because testing and field results from other customers have shown that this exact scenario can happen.

Each regional name, like 'outlook-emeasouth.office365.com', is served by Round-Robin DNS – a feature of DNS that, upon each new request, will reorder the IP addresses for the name. This provides a simple way to provide connection balancing across several IP addresses. However, and especially when performance is important, not all of the IP addresses for a regional name will have the same responsiveness. It is wise to ping and develop latency numbers for each IP address in the likely nearest datacenter hostname. One can then lock in the use of one of the IP addresses by the use of a "hosts" file (works for a single computer) or by adding the name in the organization's local DNS, possibly providing Round-Robin features for only a few of the IPs; the ones with the lowest consistent latency.

Lastly, once a tenant is created, the location cannot be changed. If the location chosen during setup was done incorrectly, there is no choice except to create a new tenant, with a different name, migrate from the old to the new, and remove the old tenant. Then finally rename the tenant. However, as simple as that may seem, one cannot have the same email domain active on 2 tenants – so that means that temporarily, the new tenant will need a different email domain until the migration completes and the old tenant is removed.

This situation can also occur in a more natural case. Consider a situation where a business decides to relocate to a different region. There is no ability to simply change the tenant details with regards to that change. Users will change from being near the datacenter to possibly being very far away from it, increasing the network latency in the process.

CHALLENGE 3

Email Clients

Office365's email service is based, currently, on Microsoft Exchange 2013. As such, only certain mail clients will work with this service. It is important then to understand the full distribution of types and versions of the various email clients that exist in the organization. Failure to fully catalog how the business will connect with Office365 can mean that some users cannot get mail, or that some applications – that depend on a specific version of Outlook, for example – may not work with Office365.

Microsoft Outlook, being the primary client most organizations use to access email and to collaborate with other users, is of special importance. While Office365's Exchange platform is based on Exchange 2013, that will not be true forever. Work has already started to move tenants to the next wave of Exchange – Exchange 2016. Exchange 2013, and therefore Office365 as well, supports Outlook client connections from Office 2007 and later, with appropriate service packs and roll ups. However, it is likely to be a case where tenants that are created or transitioned to the next wave of Exchange find that Outlook 2007 will no longer connect.

Microsoft attempts to make this an easy decision by offering in most plans the inclusion of the Office suite, which includes Outlook. Currently the office suite is Office 2013, but in the near future Office 2016 will be available and nearly forced as the version to use. However, there are many differences between prior versions of Outlook and Office and Office 2013. Some of the differences are obvious while many are subtle or nearly invisible. In a larger organization, training then becomes a potentially unseen cost with a transition to Office365. The visual differences alone between Office 2007 and 2013 are stark enough that many users feel lost as to how to do certain things. Failure to provide training to end users means that help desk calls are likely to go up and user productivity to go down.

CHALLENGE

3 Email Clients

The subtle differences are also many, the most disrupting being the way that cached-mode works in Outlook 2013. Microsoft added a new feature – albeit a very good one – to control the size of the offline mailbox copy based on a date range. One can specify that only the last 12 months of data be stored locally and any other data is to be retrieved from the server directly. This great feature though required a completely new file format and one that is incompatible with the previous versions. Deployment of Office 2013 then must take into consideration the fact that users will be required to download their mail from the server once again. If users are already at a distance from the current Exchange platform (whether that is an on-premises or cloud platform is irrelevant), the download of mail may take minutes or hours, depending upon how many items are to be synchronized. This issue is made worse if the change to Office 2013 is made after migrating to Office365 – the download of the users mail would then be done over the internet and a highly latent network.

Outlook Web Access is another dramatic change from prior versions. So much so that training is quite necessary for many of the collaborative features. There has also been a recent trend to shift users of certain criteria – mail consumers, for example, versus mail producers – to convert from using Outlook to using OWA. This trend is based on cost savings as there are plans in Office365 that only provide OWA access for less per user than a full suite plan. Companies can actually mix plans together in one tenant – as long as they are in the same category.

The use of Outlook also presents another issue that is even more subtle, but extremely important. The issue is the fact that Outlook does not have code built-in that would be called "migrated code". The meaning here is that when an organization transitions to Office365, such is very much like a cross-forest migration scenario. Outlook does have the ability to detect when a mailbox moves to another Exchange database, but only within the same Exchange Org deployment – meaning within the same AD Forest.

Outlook, especially older versions, store detail in its configuration settings – called an Outlook Profile – that point at Domain Controllers, Global Catalogs, and Exchange servers from the environment in which it was first created. There is a large misleading idea that a service – AutoDiscover – provided by Exchange (including Office365) will somehow fix this. The AutoDiscover service is nothing more than an information service; it provides information about a user, nothing more.

Outlook uses this service to get current connection points and URLs for the various services that are provided by Microsoft Exchange. However, when a pre-existing Outlook profile is simply "repointed" to another environment, it does not attempt to detect if the change was simply a database change, or a cross-forest/cross-premises change. As such, Outlook will simply try to use the new connection points returned from AutoDiscover – it does no cleanup of the old, on-premises settings that exist in the profile.

The results then are mixed and intermittent. For some users, there appears to be no issue. For others, Outlook locks up intermittently or takes an unusual amount of time to start up, when first launched. When the source environment resources, for which the Outlook profile has pointers, goes offline or is finally removed from use, Outlook will start to lock up. The reason is that while those settings still exist in the profile, it will attempt to make a connection with the server, even if it has no need for its services. TCP times on the network can range from a few seconds to 1 minute or more, depending upon settings.

Customers and consultants may claim that Hybrid-Mode solves this issue, but such is not actually true. Hybrid mode, by design, forces the on-premises resources to remain alive and active. However, as the business progresses in time, the older on-premises resource will likely be replaced or removed, even if Hybrid-Mode will be maintained. When those events occur, Outlook will begin having problems.

Others may say that upgrading to Office 2013 will somehow take care of this. This is also not exactly true. If Outlook 2007, for example, is replaced with Outlook 2013 through an upgrade, Outlook will continue to use the profile that Outlook 2007 was using, including all the pointers to legacy resources.

It is imperative then to consider how the Outlook profiles for end users will be managed. Failure to use a tool or process that specifically addresses profiles means that issues are simply waiting in the dark and will arise at some later date, in a very uncontrolled fashion.

CHALLENGE 4

Applications, Devices and Services (A/D/S)

Office365, by the nature of the service, is restrictive with regards to A/D/S and how those might interact with Exchange and email. It is important then to analyze the current environment to try and discover what applications, devices, or services might be integrated or dependent upon the email system.

For cases where these things are deeply integrated, it may not be possible for them to work with Office365. Some examples of deep integration are:

The use of an application mailbox.
In this case, the application likely logs on directly to this mailbox and uses the same to send and receive mail. Help desk applications very commonly are implemented in this way.The challenge then is configuring the application or service to use a mailbox in Office365. Currently, and for the foreseeable future, access to mailboxes in Office365 require the use of Outlook Anywhere (an HTTPS connection mechanism). Unless the application or service is up to date, it may not be able to make such a connection. For example, there are several applications that depend upon Outlook 2003 – such applications will not work with Office365 while that dependency exists. Also remember that when the next wave of Exchange is deployed internal to Microsoft, Outlook 2007 may drop from support – any application depending specifically on Outlook 2007 may be an issue in the future.

System Privileges and Admin Access
In an on-premises deployment, it is possible to setup a service account that has broad access to all mailboxes in a database, all databases on a server, or all databases for the entire organization. Blackberry Enterprise Server (BES) was well known for this approach.

There is no ability to have a permissions model like this in Office365. Applications that depend upon a single account having broad access will likely not work and cannot be forced to work in this way. System Privilege is a special access mechanism in which an application can access Exchange data with an elevated privilege level that ignores user permissions. For example, system privilege allows an account to access the entire Public Folder tree in an organization, regardless of individual folder permissions set by users or administrators – this is the same level of access that most Exchange services

Applications, Devices and Services

use like Hub Transport, System Attendant, and many others.

Direct access to Active Directory

Many A/D/S are configured to work via LDAP with the local Active Directory environment. There is no LDAP access to information in Office365. A/D/S that depend on LDAP simply won't work with Office365 directly and this may force the organization to use Hybrid-Mode.

Transport Rules

Some A/D/S are setup and configured to use one or more Exchange Transport Rules, or are dependent upon the same to direct mail towards itself. Office365 does support transport rules, but not necessarily with the same level of feature as on-premises. This is especially important if the on-premises transport rule depends upon some local service, API, or fileset on the Exchange server – there is no ability to customize the Exchange server implementation in Office365.

At a minimum, there will be need to recreate the transport rules and depending upon the complexity of the rule, may or may not be able to support the A/D/S that uses it. Lastly, not all service plans include the ability to create and manage transport rules. Be sure to pick the right service plan if a discovery is made that requires the use of transport rules.

POP3 and IMAP

Some A/D/S are only able to work with less rich Internet protocols. In Office365 POP3 and IMAP are supported, but have to be set on a per-mailbox basis. If the choice of migration tools do not carry forward the source mailbox's setting it will need to be set manually aftewards. Or, if these settings were set by a policy in the on-premises environment, manual work will likely be needed to ensure that the new Office365 mailboxes have the settings needed.

Custom Address Schemes

Some A/D/S use custom address types and either a transport rule, custom gateway, or custom receive connector to facilitate the routing of such emails. For example, faxing solutions of use a custom address type (non-SMTP) to handle routing. However, Office365 does not support custom address type send connectors like an on-premises deployment.

Additionally, it is important to select a migration solution that WILL migrate all the source email addresses, even if some are not SMTP (again, like a faxing solution), however in order to have routing work based on those custom address types, a transport rule may be necessary.

Locally Accessible Data or Services

Some A/D/S require the install of a service, application, or libraries directly on an Exchange server, or on a member server in the same domain as Exchange. Office365 provides no ability to modify the servers hosting exchange. For applications like this, it may be necessary to look for a newer version of the product that specifically supports Office365, or a different product entirely.

For products that require a host in the same domain as Exchange, the same issue exists in that there is no ability to integrate with the hosted environment in Office365 in such a way as to add member servers or the like.

Outside of A/D/S that have deep integration, there may be products used by an organization that are indirectly dependent upon the on-premises organization. In those cases, it may be as simple as reconfiguring the product to point to services at Office365 (POP, IMAP, etc.). Check with the vendor's support team for details on Office365 support, and more importantly, how to transition from on-premises to Office365. This latter question is important because a product may support Office365 for new installations, but may not have any tools or feature for transitioning, possibly meaning much manual work to reconfigure.

Lastly, there may be client add-ins for Outlook that depend on things that are expected to be available in an on-premises deployment but which don't exist with Office365 – for example, Active Directory. A survey from users and department leaders is prudent to capture how users interact with Exchange. It may be found that there are add-ins that are either 3rd party or are created internally to support some LOB application, service, or process. Failure to discover and apply a transitional path to such things can cause much frustration the day after a migration to Office365 or, potentially worse, failure to collect revenue if an add-in is specifically related to such an activity.

CHALLENGE 5

Migration Performance

Due to the fact that most on-premises deployments of Exchange will be at some physical distance from Microsoft's Office365 datacenters, there is an inherent limiter to speed at which a migration can occur. Network latency is a frustrating reality that affects all migrations – on-premises and cloud alike. The issue is due to the nature of how data is stored in Exchange and how it is accessed.

The core of the issue is not how much data there is, but how many items there are to be migrated. Consider 2 mailboxes of the same exact size: 500 MB for each mailbox. However, Mailbox-1 has 5000 items in total, but Mailbox-2 has 15,000 items – the second mailbox will take nearly 3x longer to migrate due to the item count. This issue is compounded by network latency.

Network Latency

Latency is simply a delay on the network waiting for electrons or photons to flow down a cable. The greater the distance the longer the delay. If the organization's on-premises datacenter is 2000 miles away from the nearest Office365 datacenter, there will be high latency over that distance. Adding to this issue is the possibility for poor network routing that artificially increases latency by routing data over a much greater distance than necessary. Consider a case where, in order to reach the nearest MS datacenter, data is sent from the on-premises datacenter, to another state, timezone, or even a country. Looking at how data is routed may show that it would be better to bypass that route, for specific destinations (like the MS datacenter) and have it go directly from the on-premises datacenter to the MS datacenter.

The impact of latency is quite important to understand as it affects ALL TOOLS and applications, including Outlook clients. In any given application that needs to request and receive data from a remote server, there is the network protocol underneath that makes it all work, primarily TCP. Stateful connections like TCP (of which HTTP is built on) have a series of short back and forth conversations before real data starts to flow, a sort of handshake that allows each side to know what is about to happen and when it ends. Latency causes the delay to be felt before and after each transmission of data.

In the case of Exchange, this delay is felt on each side of an item copy. There's a delay just before the item is copied, and then at the end when the item copy completes. If the latency between the on-premises mailbox and the Office365 mailbox is 50ms, and the mailbox has 10,000 items to be migrated, at a

Migration Performance

minimum there will be 500,000ms of just waiting around, which equals over 8 minutes of doing nothing. However, accessing data in Exchange is not so simple. Items in exchange, whether they be a folder, message, contact, calendar, or some other object type, are not files – they are not stored as a row-of-bytes in a database. Each item is a table of properties, possibly with sub-tables of other properties (like the recipient list and attachment list on an item). There is more back and forth accessing an Exchange item than would be if it was just a file. This means that the 50ms latency could be over 100ms for each item.

Large Mailboxes by Item Count

Given the above accounting of how latency affect the performance of a migration, item counts become the most important metric to understand in an organization. It is prudent and valuable to generate a per-mailbox accounting of this information. Identifying those mailboxes with exceptional items counts and then providing a strategy to reduce those items counts will help control the schedule and duration of a migration effort.

Large item count mailboxes have a secondary and subtle concern. When large item counts are to be migrated, this means that the activity for that mailbox will take considerable time – so much time that the likelihood of network interruptions and environment changes increases. Most processes that migrate mailboxes do not have durable mechanisms to handle unreliable network connections. It has been reported many times that even the Microsoft tools will fail, or just “lock up” on large mailboxes, likely for this reason. It is then important to ask for details about how large mailbox are handled if a 3rd party solution or scripts will be used.

Another subtle consideration of very large item counts is the impact they have on Exchange servers, especially older versions. When a request to access a folder’s contents is made, Exchange must reserve memory to hold the table of items as a list for the requester. This creates a sort of “table session” for clients and tools. However, when the list is very large, it can consume considerable RAM on the Exchange server. Older versions of Exchange had recommendations to not exceed 5000 items in a single folder for this reason. While there should not be much concern about large item counts in Office365, it can affect the performance of a migration on the source side, especially if the RAM allocation of the host is not properly scoped.

Throttling

Microsoft Exchange 2013, of which is currently used by Office365, has the ability to throttle and control many different types of activities. This throttling is key to providing stabilized and consistent availability and up-time of a large scale email platform – whether on-premises or in a hosted environment. However, this same throttling then quickly becomes a bottleneck for large scale operations.

Unlike on-premises deployments of Exchange where the throttling values and assignments can be controlled and manipulated, Office365 does not allow changes to these values directly - there is the possibility of requesting a temporary relaxation of throttling for a short period, if your organization is of the right size and by Microsoft’s discretion.

Throttling affect the performance of many migration tools and processes, including Microsoft’s own tools. There are both measurements of concurrency and time-on-feature that are throttled. However, all throttling is based on authentication. If the same account is used to do multiple, concurrent operations, throttling is quickly found and operations are either block or delayed (aka tar-pitting). This effect should be analyzed both with tool selection for migration AND any other 3rd party A/S/D for which performance or concurrency play a role. One side effect that can happen is that a 3rd party LOB application that works with Exchange does many operations, operating as a user. That same account is used for several other tools, services, or applications and all end up suffering with the effects of throttling.

Good Tools

Good tools will take these things into consideration and will have methodologies to be durable and performance oriented. Multi-threading and multi-process designs can combat the effects of latency by running multiple channels of operations concurrently. Better tools will know the impact of throttling and will have strategies and designs to work with or around it. Tools that suggest the use of export/import routines wouldn’t be considered a good solution. While it may seem like the latency issue is eliminated, there are many other risks that are worse, and in truth, if one considers the full effort of an export/import **routine, a direct migration of data will still be faster if, for no other reason, the data is only copied one time.**

Priasoft Migration Suite for Exchange™ 6.5 is the leading Microsoft Exchange migration solution that provides class leading flexibility, ease of use and performance.

Learn more at priasoft.com

CHALLENGE 6

Duration and Scheduling

Projects consume time – there's no escaping that fact. However, the duration of a project is variable and can have many influencers. Migration projects are no exception to this fact either. Pressure to start or complete a migration can range from little to overwhelming and stress on the team members is often proportional to the level of confidence that proposed timelines can be met.

Confidence then becomes a very important aspect of a migration project. Without it, much second-guessing and retesting occurs, further impacting the schedule. How does one gain confidence then? The answer is best achieved through metrics and measurements of actual actions. It is far too easy to make a bad assumption about quantity, quality, and performance. While some testing may lead to better assumptions and be based off of some level of calculation, such are still assumptions with a high margin of error.

Dry-runs and simulated user tests are key components to providing the metrics that begin to remove assumptions and thereby gain confidence. Dry-run processes that perform actual migrations are best as they show not only how long a process will take, but more easily surface bottlenecks and influencers in addition to showing success and failure. Without a dry-run process, one cannot know the quality of the process until the production migration occurs – which is the worst time to discover issues. A proper dry-run process will expose issues across all of the influencers: network, storage, content, time, etc. By the nature of a dry-run, discovery of issues is then without stress since they do not directly affect end users. In addition, the removal of such stress allows a more measured approach when analyzing the issues and provides an ability to discover patterns and root causes. This is in direct contrast to a production migration with issues where there is no time to wait for a pattern to develop.

Once the actual duration of a migration can be determined, better decision making can occur with regards to the actual migration strategy. Without this measured approach, one cannot know how long things will take which in turn influences the design of the migration process to be one with long coexistence period, migrating small batches of mailboxes over many days or weeks but without any real understanding of when the project will complete. A dry-run will provide exact metrics for time and performance and positively influences the scheduling efforts. Knowing that a migration can be completed in 40 hours lets

decision makers look for weekend opportunities to migrate while an 80 dry-run result shows that at least 2 events is necessary. The quality of information derived from dry-run exercises prevents decisions that like: “I think we just have too much to do in a single event. We should then do this in 8 events”. The best approach is to have a strategy that works towards having the fewest migration events possible, with a single event migration being the best.

Migrations to Office 365 are slower than on-premises migrations, not only due to latency (as described earlier) but also due to the mechanisms used to manage and manipulate objects in the cloud, like PowerShell and SOAP. Good metrics and timings will help with scheduling and with how many migration events are actually necessary.

Lastly, it should be understood that a long migration period – one that takes many weeks to complete – are the most common cause for disruption, dissatisfaction, and negative perceptions by end users. The coexistence between on-premises and Office365, while good, is not 100% transparent to end users. Public Folders, shared mailboxes, and cross-premises delegation are some of the issues found during coexistence. For more details about how coexistence impacts migrations, please see our documents about the topic: Co-existence vs. One-time Migrations and Detailed Challenges of Coexistence.

CHALLENGE 7

Email Address Namespaces and Domains

Office365's mail service allows for a tenant to accept mail for one or more custom email namespaces. This is important for many organizations and is an easy feature of the Exchange product. The existence of multiple domains happens for many reasons: acquisitions, applications, regional use, etc. Regardless of the reason, it is important to survey and understand ALL of the email domains that an organization has, prior to migration.

The complexity introduced with Office365 is due to how one sets up the tenant to accept mail for a domain. There is a process of setting a record in DNS, which proves that the organization owns the domain – the idea being that only the domain owner would likely be able to modify public DNS in such a way. Services in Office365 will then periodically check public DNS for the domain and when it is able to query for the specific record, it will then allow mailbox to be received for the domain.

Imagine now if an organization has 10 different email domains, likely due to several acquisitions over time. Furthermore, those 10 domains are managed by 5 different public DNS hosting services (like Network Solutions, GoDaddy, VeriSign, your Internet ISP, and many others). In order to be able to receive mail for those 10 domains, work must be performed across those 5 DNS hosting providers. Rarely are these tasks scriptable and often involve cumbersome web page portals and the like. Making things even worse is the commonality that Exchange administrators may not have the credentials to manage Public DNS for an organization as this is often under the scope of the network or Active Directory team.

Faced with this burden, the easy choice may be to only setup for one or 2 of the "primarily used" domains. Leaving out the other 8 or 9 for a setup at a later time. There is a subtle but frustrating side effect of this idea in that one CANNOT set an SMTP address on a user in Office365 to a domain that is NOT a validated and confirmed accepted domain. If a migration is to occur, any mailboxes that have addresses for the 8 or 9 other domains, depending upon the tools used, will NOT have those addresses, or will be adjusted to be something other than SMTP (Priasoft's tools, for example, will set such addresses as "LSMTP", the 'L' designating 'legacy').

CHALLENGE 7

Email Address Namespaces and Domains

There may not be any immediate issue with this approach right after the migration completes. However, some days or weeks later the issues may start to appear, but now are more difficult to understand due to the distance from the event. Users may have periodic email from suppliers, partners, sibling organizations, etc., that only come in once a month or quarter. A user might not have been able to voice this during project discovery and surveys and experiences a case of "I think I'm missing emails". The supplier or customer, that has an automatic system that still uses an older email domain, doesn't yet know that their emails are bouncing and it can be many days or weeks before it is all sorted out.

Adding to the complexity is internal user activity and company applications and services. If it was determined that a particular email namespace was to be deprecated, or could NOT be reused (perhaps due to ownership of the domain), there can be many cases where an internal reply is done, initially, by SMTP. Consider a case where an add-in for Outlook, that integrates in some way with the company's CRM solution, only knows how to store and work with SMTP addresses. After a migration completes, the add-in works as it always has and addresses a mail to name@unmigrated-domain.com. Since it is not possible to set the domain on the accounts in Office365 unless the domain is verified, this either means the address is missing, or is modified. The add-in will create the mail, but the transport server will reject the message because it is unable to find an object with the email address.

So, listing out all the accepted domains from the current on-premises environment is a good start, but is possibly not complete. In an on-premises deployment of Exchange, an administrator can set additional SMTP addresses on a user regardless of whether the domains are listed in the Accepted Domains of the Exchange Org. This is sometimes used for custom email relaying – a namespace specific send connector is created, perhaps like 'mail.sibling-corp.com', and any internal user that attempts to send to an email address @mail.sibling-corp.com will be routed through that send connector. Thus, in addition to capturing all the accepted domains in use by an organization, evaluating all the send connectors is equally as valuable.

Lastly, and also in an on-premises Exchange deployment, users in the same Exchange environment can send mail using ANY address that is found in the 'proxyAddresses' list of another user. There does not have to be a receive connector nor an accepted domain. Normally this does not occur by design, but mergers and acquisitions can create this situation over time. For example, consider a company "the-big-co.com" acquired "cool-stuff.net" many years ago. Assimilation over the years makes it seem like the "cool-stuff.net" address space is no longer in use and it is removed from the accepted domains list and is removed from Public DNS's MX records. For all intents and purposes, the namespace is dead and users external to the mail system would not be able to send to addresses with that domain. However, when the accepted domain was removed, no one went thru an removed all the "@cool-stuff.net" addresses from all the user accounts. A long time user, by habit, has been using the "@cool-stuff.net" namespace ever since the acquisition to email some of his co-workers, even after the namespace was taken out of use. Exchange will happily do lookup for SMTP addresses using LDAP and a Global Catalog in Active Directory to try and find an object to which to deliver.

If the "@cool-stuff.net" is not properly considered or preserved in Office365, the user will have issues and may drive up tickets at the help desk. If this particular user is the former CEO of Cool-Stuff and still has a high-ranking and influential position, things can get worse. Perception is reality in cases of migration projects. Even if this person accepts that he can no longer use the namespace, at a minimum, his perception is likely to be "They should have at least told me before the migration. It seems they knew it would be an issue." If "cool-stuff.net" is used by applications, services, devices, or other users, the importance of preserving or properly communicating about its final deprecation is important.

CHALLENGE 8

Public Folders

Public Folders are an Exchange core component and a feature that has been around almost since the product was first created. Due to the nature of how Public Folders work and the fact that they are “Public” causes them to often grow in size, depth, and complexity. For many organizations, Public Folders represent a unique challenge in that they consist of company data but for which most of it is simply a data repository and not in heavy active use. Compliance and retention regulations often dictate that this data be maintained and preserved, not only from a “the data exists” point of view but often from a “is searchable” standpoint as well, meaning that the data must be brought into Office365.

For nearly all organizations that have Public Folders, there exists some percentage of them that are in use and for which users are dependent. A migration to Office365 presents a unique challenge in that there is not a Microsoft supported (or even unsupported) option to synchronize public folders. While Microsoft does provide a script to migrate the folders, this often is a complete failure for several reasons, the primary being that in order to migrate the folders, cloud users cannot use the Remote Public Folders feature, resulting in users not being able to see or work with their folders (because they do not yet exist in Office365).

In cases where the migration pattern is one that calls for a migration over multiple events (days or weeks), synchronization of Public Folder data becomes an important topic very quickly. The Remote Public Folder option from Microsoft attempts to address this by directing all Public Folder requests to the on-premises data, but this just delays the issue as described in the previous paragraph. What is actually necessary is to have both on-premises and cloud users seeing the same data, but in the respective environment relating to where their mailbox resides. Simply put, synchronization is necessary.

Even if the number of mailboxes in the organization is relatively small and the mailbox migration effort can occur in a single event, it can and often happens that Public Folder data is larger. In this pattern, mailboxes migrate first, users are unable to use the Office365 public folders – either because they have not yet been migrated. When a Public Folder structure is very large, either in data size or item count, or both, the consideration of speed and performance become extremely important. Users might not be able to wait for an extended period of time to get to Public Folder data in Office365.

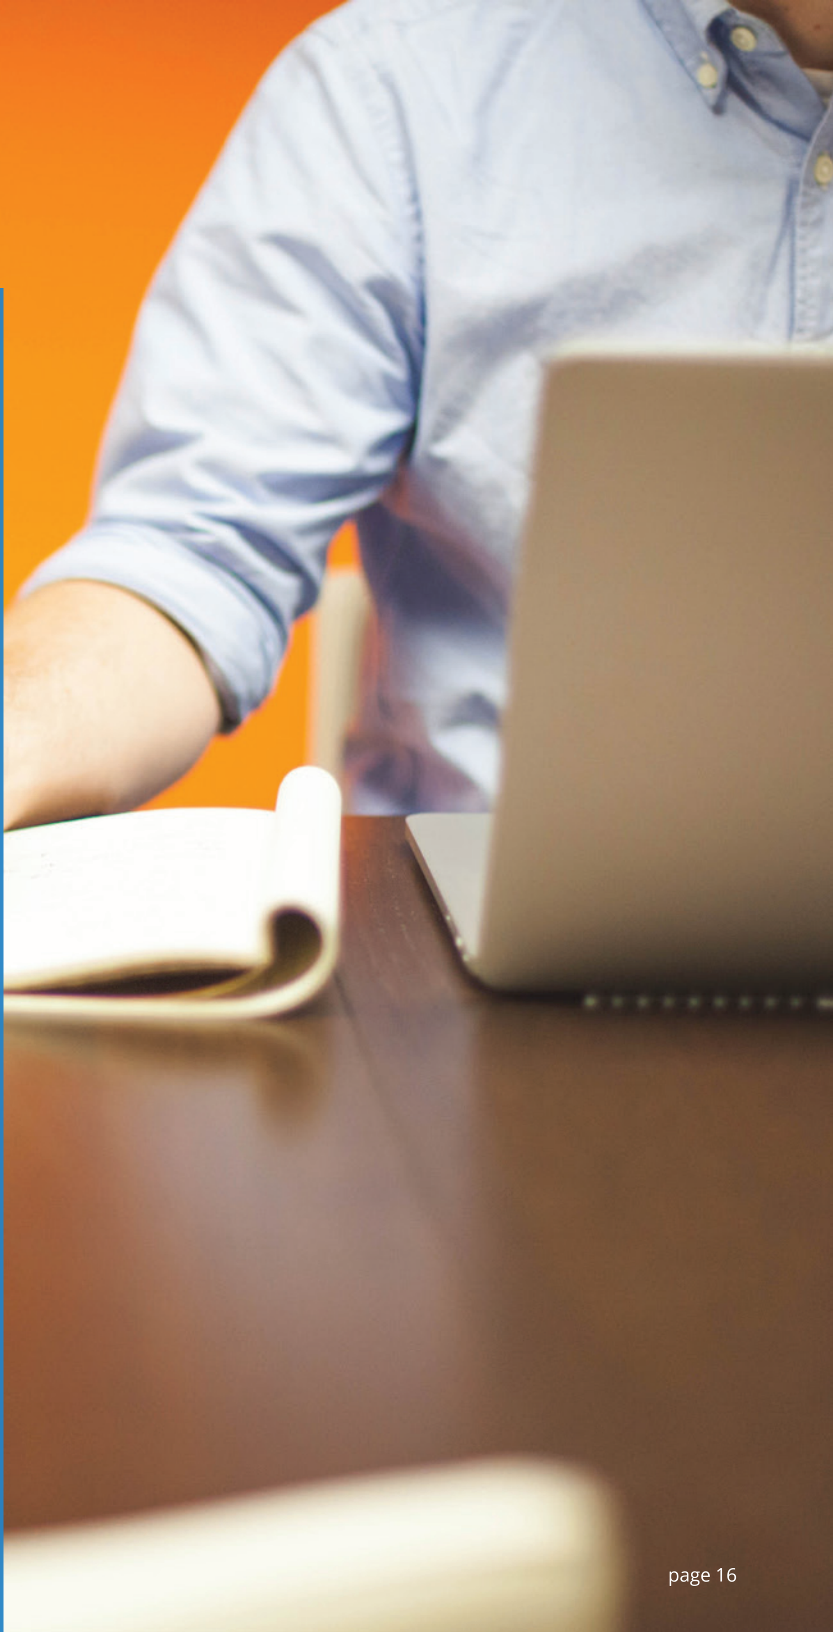

CHALLENGE

8 Public Folders

Consider the following likely scenario and one that has been seen before:

In an organization of 1500 users, a senior level decision is made to migrate to Office365. Much testing ensues and tools are engaged to solve some specific problems and, for the most part, the mailbox migration effort appears to be well managed. During the analysis and discovery phase of the project, it is found that there exist 3200 total Public Folders. The total data size is approximately 600 GB, but is not evenly distributed across the folders; there a few folders that constitute 90% of the data while most of the other folders have little or no data.

Due to the size of the Public Folders, guidance is given to use Remote Public Folders so that there is no delay in starting the mailbox migration efforts. Testing and business tolerances influenced the scheduling and there will be 5 migration events, 1 each week. The first groups of users that migrate continue to access the on-premises Public Folders. Work to determine the effort required for Public Folders is easily delayed because users are not affected by the current implementation. The Remote Public Folders makes it so that users are not impacted. As such, the challenge and issue is hidden from the migration project and a poor assumption is made that it should be a non-issue.

The last migration event occurs and now all users are using mailboxes from Office365. At this point, work is started to migrate the Public folders to Office365. It is discovered, unfortunately late in the project now, that in order to migrate the folders, Remote Public Folders must be disabled. For whatever duration exists to migrate the Public Folders, users cannot access the data needed in the system because it does not yet exist. This is a double-edged issue in that not only can users not access prior data, but they also cannot create new data, at least not without creating a conflict for the Public Folder migration task.

Due to the poor performance of the Public Folder migration (Microsoft's scripts, or PST exports, or the like), the forecasted duration becomes too much for users to tolerate and department leaders complain and the effort is put on hold so that Remote Public Folders can be re-enabled to allow users to do their work. At this point, things work, but there's almost a sense of dragging around a boat's anchor in that there seems no productive way to exit from the on-premises deployment. Users continue to make changes in the Public Folders (on-premises) and they continue to grow. Cost savings that were supposed to be realized by a move to the Cloud begin to dry up a bit since there is still a spend on maintenance, support, and management of this final on-premises component of the Exchange email system.

At some point, hands throw up and decisions are made to "just do it" and migrate the Public Folders forward. Help desk teams are told to be ready for more calls until the Public Folder migration can complete. This even assumes that there's no issues in the actual migration effort, but reality sneaks in and shows that the task has its own set of delays and resets due to the size of the Public Folder deployment. What was supposed to be an effort of a few days now becomes a week. A few of the folders are then found to influence revenue into the organization and rushed and rash decisions are made to deal with those issues, further frustrating the migration effort by causing duplication and conflicts.

If only there was a way to have either pre-loaded the Public Folder data prior to mailbox migrations or, even better, to have had a synchronization job to keep the on-premises data and the Cloud data in sync.

The above scenario is very real and is hidden in migration projects as one of many "you don't know what you don't know" situations. Given the above, scoping in Public Folders early in the project plan is important and necessary and then calls for appropriate tools and processes. Delaying the problem only makes it worse. The above is further complicated by the fact that Public Folders can have specific user permissions to allow access that can be set by user or group. If the migration process does not handle folder permissions, it is likely a wasted effort since the data would exist but users would be unable to access it. Public Folders can also be mail-enabled. However, unlike permissions which can be seen, reviewed, and managed with Outlook clients, mail-enabled folders are only seen with Exchange management tools. As such, mail-enabled folders are difficult to visually see and many users may not know if a folder is mail-enabled. Migration patterns then must include this feature otherwise post-migration issues may appear of which are likely very subtle – users notice only after some days or weeks that a folder is "missing" data they expected to be there.

In the end, properly scoping in the Public Folder deployment is important not only for the data that comprises it, but for how they are used and accessed by **users. Tool choice is then important and the best tools would be those that have coverage for such things.**

CHALLENGE 9

Archiving

In larger organizations, or those that have strict rules or regulations for data retention, often have Exchange archiving systems. Migrating out of an environment that has Exchange archiving means that there is a second set of user data that is often disconnected in some way from the user's primary access (mailbox or public folder). Failure to include archived data in a migration plan to Office 365 can not only cause user issues and complaints, but could impose legal issues or even fines if it is determined that access to archived data is lost or cannot be retrieved within the boundaries of acceptance for a rule or regulation.

In addition, archived data is often of a size that is of a multiple of the current "live" data meaning that possibly several terabytes of data may be sitting in one or more archive platforms. Given the close relationship between an archive system and an on-premises Exchange deployment, there are often features provided by archiving vendors for accessing or rendering such data in a convenient and transparent way to end users. Often such is facilitated by access to the on-premises Active Directory resources and the inherent security framework it provides. However, a migration to Office 365, while an Exchange deployment, does not allow direct access to Active Directory in the same way as an on-premises deployment.

Many 3rd party archive solutions do support Office365, but only for new data and with some limitations. Furthermore, many organizations are vacating their 3rd party archive platforms in favor of Microsoft allowance for archive data to live in the Exchange environment directly, and for some of the Enterprise plans, with no cap on data size. The challenge then can be due to the sheer size of the archived data. While Microsoft does have plenty of scripts and simply utilities for mailboxes and such, they do so because they are supporting their own product. Archive data migrations however are out of scope for Microsoft and many Microsoft advisors and some consultants will side-step the discussion as it is seen as out-of-scope for their required tasks

Further frustrating a migration plan are cases where the archiving platform leave "stubs" or "shortcuts" to real data in users' mailboxes or in Public Folders. Migrating these lightweight items rarely works well, if at all. Older versions of this idea involved the use of a "Organizational Form" that contained code that dealt with the rendering of the data, authentication, and access controls to the archived data. Office365 doesn't

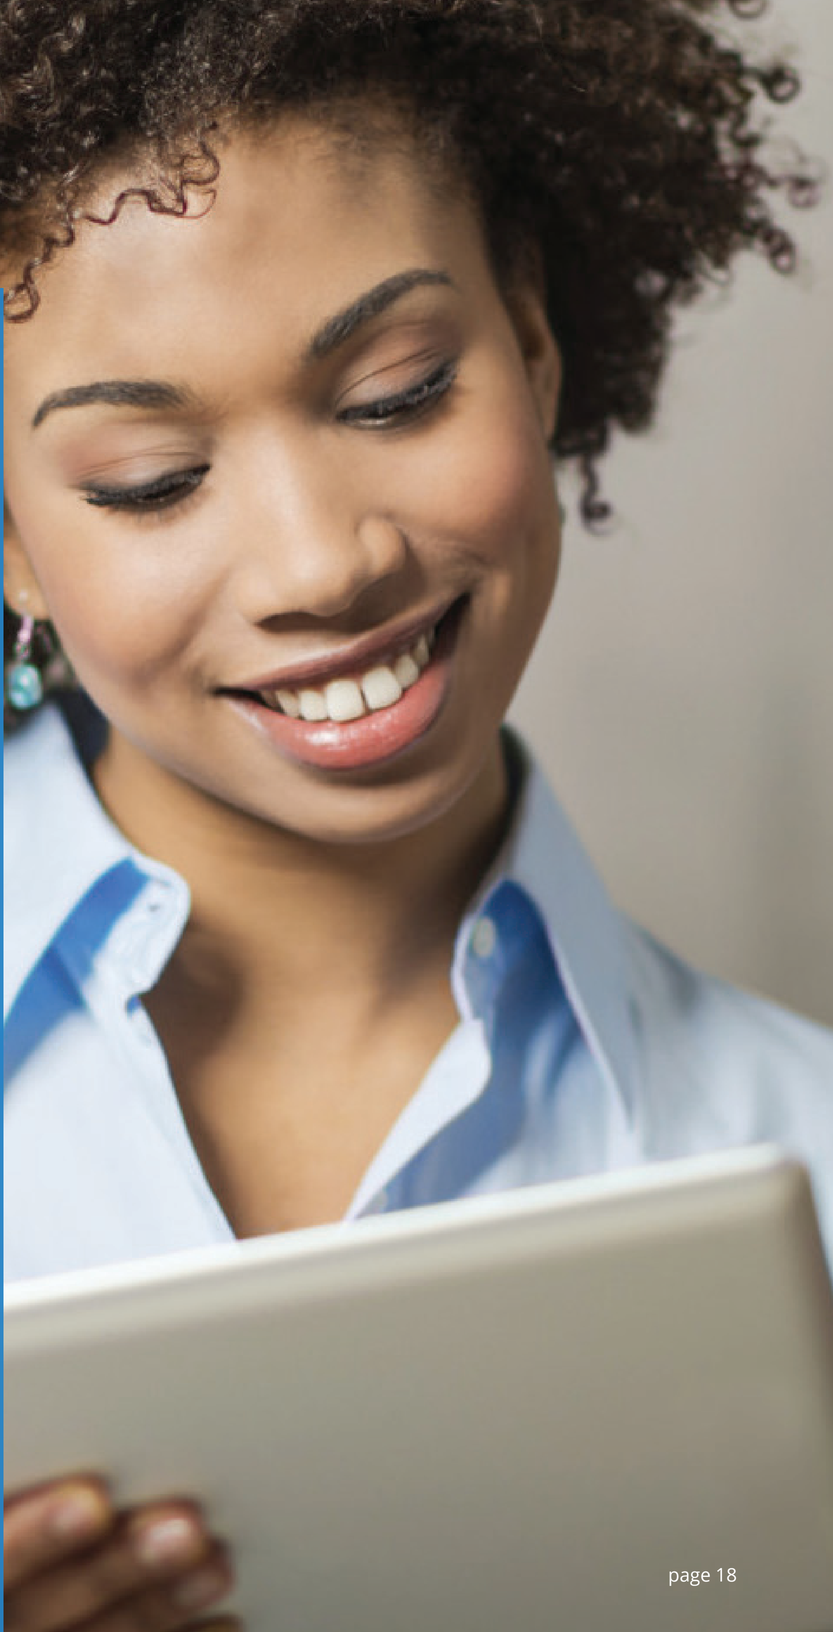

have broad support for Organization Forms and it is unlikely when or if such will be available. However, even if such was available, the forms often are custom to the environment and rely on the fact that the archive servers and the Exchange servers are in the same windows domain, relying on domain features for securing access to data. The same doesn't work with Office365.

When stubs do exist, and when it is seen that they should not be migrated, then new work appears to handle and manage them in the on-premises mailboxes before migration. Current ideas are to either "rehydrate" the stubs, turning them back into real items, or to remove them. Both cases are affect end users in the production environment prior to migration, with the latter possibly create substantial negative feedback from users. Rehydrating the stubs, while initially appearing as an appropriate avenue suffers from data size issues – there may not be enough free space available in the Exchange platform to house all of the archived data. This causes the project to be stalled or a horrible "shell game" to be played where a few mailboxes are de-archived, then migrated, then deleted in the source to free up space for the next batch. This means that rollback scenarios are not possible and adds unnecessary risk. A good solution is one that encompasses the impact of archiving and likely means the use of a 3rd party archive migration solution. However, don't expect to get much support from the archive vendor itself – nearly all will offer an ability to export archived data, but not to migrate it. Once the realization is made that 3rd party, non-vendor tools are needed to migrate the archived data, the next task is how to integrate that into the overarching migration process. Unless the mailbox migration tools support some level of direct integration with the archive migration tools or provide some scriptable events, it can be quite challenging to provide an atomic process for migration. This can then involve several manual steps – of which one or more can be forgotten, done in the wrong order, or mismatched.

An even better solution is one where the migration tools for mailboxes and the migration tools for archives can work together in such a way to reduce the effort overall. One of the immediate ways is if the mailbox migration tools allow for the exclusion of stubs during migration. With this idea, the efforts of either rehydrating/dearchiving or removing stubs is eliminated as a task. This provides a better rollback story, does not cause concerns about free space, and does not make changes in front of end users. Excluding stubs means the migration of mailboxes is quicker, does not create items that users cannot use, and shortens the overall timeline since pre-work is not necessary on the archiving side.

CHALLENGE 10

Language Issues

Search for the phrase outlook wrong “folder language” office365 in Google or Bing and you will get millions of results. The issue stems from a change Microsoft made starting in Exchange 2010 where the standard default folder names could no longer be modified directly – in prior versions of Exchange, the folder names could be modified just as any other folder (by a developer or tool writer). With this change, Microsoft decided to become authoritative for the names of the standard folders and put the responsibility of the creation of the standard folder on Exchange, taking away the responsibility from the MS Office team.

The language of the default folders is determined by only a few factors, and once set, does not change automatically in reaction to any client activity.

The ‘msExchUserCulture’ value on a user object in Active Directory

- If a value is set here, then at the time of the FIRST LOGON to the mailbox, it will receive folders based on the language value.
- This attribute only appeared in the Exchange AD Schema starting with Exchange 2007 and no much was written about the feature and it is not required to have a value.
- This means that upgrades from Exchange 2003 would never have had this value set.This value does not have a corresponding user interface element in any of the management dialogs for a user or a mailbox. It would only ever be set by a script or by direct edits using an LDAP editor (like ADSIEDIT).

The language of the host that does the first logon to the mailbox

- The subtlety of this statement should not be missed. The first logon, during a migration is not going to be the user but will be the tools that perform the migration. If that host is ‘en-US’ but the user’s mailbox folders are in French, the standard folders will be English afterwards.
- Even outside of a migration, there can be cases where the first logon is not the user. Consider where shortly after mailbox create completes, but before a user logs on, a new piece of mail is sent to the mailbox. At that point, the transport server responsible for mail delivery will be the first logon. The language of that host will be used to set the folder language.

CHALLENGE 10

Language Issues

- Even if the first logon, after migration, is the user's computer, the language on that host may still not be a reflection of the language of their folders. The user could have received a new PC between the time of the creation of their source mailbox and now.

Changing the language in OWA

- If a user changes the language option in the OWA logon page, the folders will be set to that language.
- This is a common way that users with English PCs can have non-English folders in their mailbox

Running a powershell cmdlet

- An administrator of Exchange can run a powershell cmdlet to force the language of the default folders.

As the world has become "smaller" and workers have become more cross cultural, even a small organization may have some users that have folder names that are language specific. It is wise and prudent to have a discovery step in the migration project to identify if there are mixtures of folder languages.

Compounding this issue is the fact that there is no setting or marker on a mailbox that identifies the language used for the default folders! The implication of this statement is that at the time of migration, there is no simple attribute that can be read to determine the folder language – such simply does not exist. Therefore, the Microsoft tools and many 3rd party tools do not take folder language into consideration – the target mailboxes will simply have the language of the first logon. Users will complain after migration and the help desk will not know how to solve the issue. Only when it reaches the Exchange admins will a resolution occur by the use of the aforementioned powershell command, but the admin will need to know the language to use.

Adding to this mess is the fact that prior to Exchange 2010, is was possible for a mailbox to not have all the default folders – because the responsibility was Outlook's. If the mailbox were a shared or application mailbox for which Outlook never did a direct logon, the mailbox might only have Exchange's original 4 default folders: Inbox, Outbox, Sent Items, and Deleted Items. The names of those folders could be in one language while the rest of the Outlook folders could be another.

The only way to handle this issue then is to have some sort of database of all the different possible folder names. With such a table, a tool or script could compare the names of the source folders and perform a lookup to see what language matches. Then, that language value could be used along with the powershell command to set the folder appropriately. However, there are currently over 300 different supported language and culture values in Exchange.

Complicating this issue further is the fact that, over the years, Microsoft has changed the names of some of the default folders in some languages. For example, in Outlook 2007 and earlier, the Dutch name for the Calendar folder was "Agenda". In Exchange 2010 the folder was changed to "Kalendar". This means that even between versions of Exchange there is some consistency issues. Having the table of folder names mentioned above might not have that subtle difference.

However, as important and troublesome this issue is, the discovery of the issue has nearly equal importance. Producing a report of mailboxes that do not have all the default folders, have different names, or have a mixture of languages is not a trivial task. If the source version of Exchange is 2003 or earlier it is extremely difficult to analyze as there is no powershell scripting available. This is a key value of dry-run migration activities – to discover, in real time, which mailboxes would present a challenge with regards to folder languages.

Priasoft Migration Suite for Exchange™ 6.5 is the leading Microsoft Exchange migration
solution that provides class leading flexibility, ease of use and performance.

**Learn more at
priasoft.com**

Conclusion

While the challenges outlined in this guide can seem a bit overwhelming even with the proper planning, it is important to understand that you are not alone in the process. Utilizing an experienced partner, like Priasoft, and a migration solution with advanced capabilities and features will set you on a path to a successful migration for your business.   With the proper planning, scope, partner and technology platform, you will be able to successful navigate the complexity of an Exchange migration and deliver cost-effective and on-time results.

PRIASOFT™

About Priasoft

As a trusted Microsoft Partner, we bring the expertise, software, and support to help you successfully transition your infrastructure safely, securely, and with reduced risk.

Copyright © 2015, Priasoft Inc. and/or its affiliates and partners. All rights reserved. Office365 logo is a registered trademark of Microsoft Corporation. This document is for informational purposes only and the contents hereof are subject to change without notice. This information should not be considered a substitute for independent, professional advice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. 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. Absent fraud, we disclaim all liability for any reliance you choose to place on any of the information provided here.

Corporate Headquarters
60 E. Rio Salado Parkway | 9th Floor | Tempe, AZ 85281
TEL: 1.480.656.7402
FAX: 1.480.366.5801
EMAIL: sales@priasoft.com

www.priasoft.com