WHITE PAPER

Supplementing the Limitations in Office 365

An Osterman Research White Paper
Published March 2018

Osterman Research, Inc.
P.O. Box 1058 • Black Diamond • Washington • 98010-1058 • USA
+1 206 683 5683 • info@ostermanresearch.com
www.ostermanresearch.com • @mosterman

PROLOGUE

We normally don't begin white papers with an "opening statement", but we chose to do so for this paper to ensure that we make an important point right up-front: while the title of this paper may imply that we are dismissing Office 365 as an inadequate offering, nothing could be further from the truth. On the contrary, Microsoft set out in Office 365 to provide a robust set of communications, collaboration, security, archiving and other capabilities at a range of reasonable price points – they have succeeded and they continue to build on that success. However, because Microsoft never set out to include every possible feature, function and capability in Office 365 – instead offering only a strong foundation of capabilities – third-party solutions are necessary for mid-sized and large organizations (and some smaller ones) that have requirements that go beyond the intended scope of the various Office 365 plans. Consequently, our focus in this white paper is to discuss objectively what Office 365 does and does not do, and to suggest areas in which third party offerings will supplement its native capabilities.

EXECUTIVE SUMMARY

- **Office 365 offers a significant and useful set of productivity, collaboration and other services, but it is not the only solution that most organizations will need to satisfy their archiving, data security, encryption, eDiscovery, backup/recovery and other requirements.**

- **Instead, Office 365 should be considered as a starting point to deploy services from Microsoft and third party vendors. These include offerings like Azure Active Directory, Azure Information Protection, and the specific features and functions that are added by Microsoft to the individual Office 365 plans on a regular basis; as well as the growing array of third-party solutions that can supplement or replace the native capabilities within Office 365.**

- **The capabilities of Office 365 are evolving rapidly, making it challenging to know when a particular capability offered in the various platforms will be adequate to meet specific organizational requirements. In short, the speed with which new features, functions and capabilities are introduced and modified makes it difficult for corporate decision makers to keep up with what Office 365 can do at any given point in time.**

- **Microsoft offers customers higher-priced plans and add-on services across its range of cloud services portfolio so that they can gain more advanced capabilities. Osterman Research recommends that decision makers evaluate these offerings, but also the third-party offerings that compete with them.**

- **While Osterman Research recommends that organizations seriously consider the native capabilities of Office 365 and deploy them it where it makes sense to do so, many third party offerings provide better functionality, often at lower cost. The result is that improved functionality and lower total cost of ownership can be achieved through a combination of lower cost Office 365 plans and third party tools to replace or supplement the native Office 365 functionality.**

ABOUT THIS WHITE PAPER

This white paper was sponsored by ZL Technologies; information about the company is included at the end of this paper.

OVERVIEW

Microsoft Office 365 has taken the world by storm, including the corporate and enterprise sectors that were expected to reject cloud services only half a decade ago. Microsoft claims

Microsoft set out in Office 365 to provide a robust set of commun-ications, collaboration, security, archiving and other capab-ilities…they have succeeded and they continue to build on that success.

to support more than 120 million active users in commercial organizations with Office 365 (at the of end October 2017), and sometime during the next 12 months, expects 70 percent of its customers to be using Exchange Online in Office 365, rather than Exchange on-premises. Osterman Research's surveys, as shown in Figure 1, clearly demonstrate the validity of Microsoft's claims.

Figure 1
Percentage of Corporate Users in Mid-Sized and Large Organizations Served by Office 365/Exchange Online

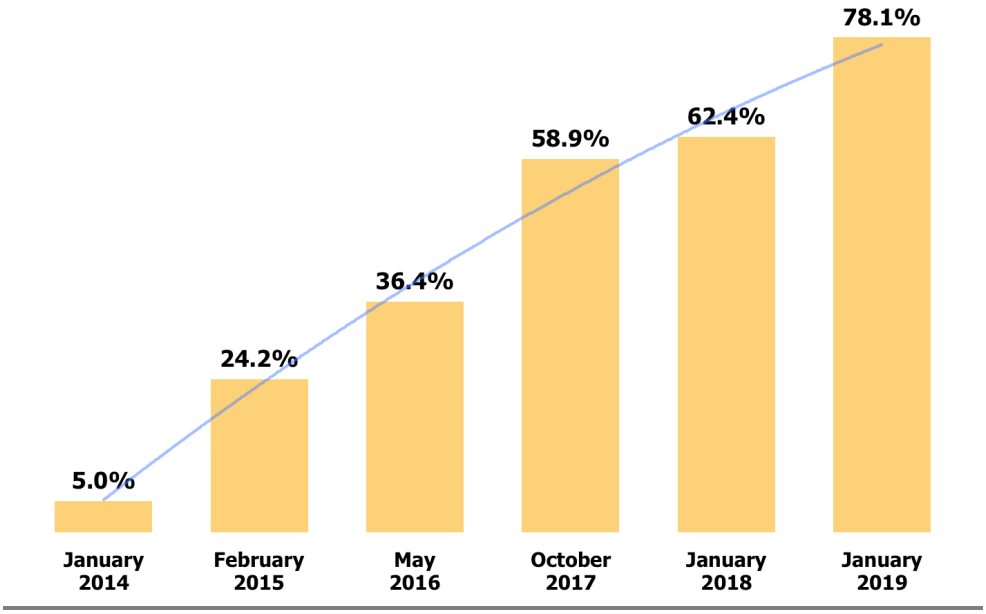

Source: Osterman Research, Inc.
Note: Yellow bars are actual survey results; blue line is a trend line

However, despite high usage numbers for Exchange Online and Microsoft's traditional Office productivity suite licensed and delivered as a cloud service (Office ProPlus), customers embracing Office 365 must make some important decisions about many of its features and functions compared to those offered by third parties. That's not to say that organizations should not consider and deploy Office 365 (we believe that in most cases they should). But decision makers must be fully aware of the limitations inherent in the native capabilities offered with Office 365 and how third party solutions can often better satisfy their requirements.

In this white paper, we evaluate what's available in Office 365 in 2018 in the areas of archiving, compliance, encryption, backup/recovery and eDiscovery, highlighting areas of concern for customers adopting Exchange Online, SharePoint Online, OneDrive for Business, Skype for Business and Azure Activity Directory.

KEY TAKEAWAYS

- Office 365 is not a single offering from Microsoft. Instead, it's a starting point for licensing a range of higher-priced and additional services from Microsoft's cloud portfolio, such as Azure Information Protection, Azure Active Directory, and higher priced plans that offer more advanced capabilities (such as the improved capabilities in Enterprise E5 compared to the much more commonly deployed Enterprise E3). In reality, organizations may better satisfy their needs by using a less expensive Office 365 plan and supplementing its capabilities with best-in-class, third-party offerings instead.

- Microsoft is rapidly evolving the capabilities of Office 365, and it is challenging to know when Office 365 – and the wider complementary portfolio of Microsoft cloud services – are adequate to satisfy a particular set of requirements. Microsoft produces useful capabilities, but they are often replaced in short order. Corporate decision makers face the challenge of understanding which Office 365 capabilities are still current, which have been improved, which will be deprecated, and how third party solutions can better satisfy their needs.

- The ground has shifted – collaboration and next-generation productivity tools are now widely available, offering modern tools and approaches for business challenges. But the new challenge is keeping employees from falling for increasingly advanced social engineering scams and malicious attacks, while ensuring data protection for personal and corporate data. Office 365 is a broad-based service that offers collaboration and productivity; are its security capabilities good enough to offer the protection that is necessary? And will "good enough" today be good enough tomorrow?

- While Office 365 is a robust service offering – particular in the basic Exchange Online and Office ProPlus – like all cloud services it is not perfect, as evidenced by Microsoft's own breakneck pace of upgrading the service. Highlighting current issues of concern assists organizations in making effective plans with clear insight about the best path forward for Office 365.

- Office 365 is designed at scale for a set of general use cases, and Microsoft's design parameters for Office 365 may not align with the needs of a particular organization. As with any cloud service, the profile of a particular customer may differ from what is offered by Microsoft. Therefore, the role of this white paper is to explore what does and doesn't work, highlighting potential red flags for an organization's deployment.

Office 365 is designed at scale for a set of general use cases, and Microsoft's design parameters for Office 365 may not align with the needs of a particular organization.

ARCHIVING AND CONTENT MANAGEMENT ISSUES

Few organizations are all-in on Office 365 to the exclusion of everything else. The vast majority have many other content repositories, on-premises data stores, and other Microsoft and non-Microsoft cloud services. The addition of Office 365 to an organization's information management architecture means the addition of new content sources and content types that need to be secured, controlled, and governed. While potentially unlimited storage is available in Office 365, keeping all data and content in perpetuity is a bad approach from business, legal, and information management perspectives. Free data storage doesn't negate the other expenses of information, including:

- **Confusion caused by out-of-date information**
  The wrong information in the hands of the right people will spread misinformation and lead to decision-making on out-of-date, irrelevant and poor intelligence.

- **Time wasted wading through wrong information**
  Information and knowledge workers already spend too much time searching for the right information; keeping unnecessary content around longer than necessary only gets in the way and slows the ability to find, retrieve, and make use of the right information.

- **Legal exposure and risk**
  When information that is responsive or potentially responsive to a legal case has been retained beyond what was necessary, risk increases. Having too much information available increases the legal and discovery costs for searching, identifying, culling, reviewing, and producing responsive content.

- **Supervision**
  The ability to supervise content is also an issue for highly regulated sectors, such as financial services. Office 365 tools will not adequately address this requirement in many cases.

Respondents to the survey that was conducted for this white paper were asked about the content management capabilities that are most important to them, as shown in Figure 3.

Figure 3
Importance of Various Content Management Capabilities
Percentage Responding "Important" or "Extremely Important"

Capability	%
The ability to have in-place search and review eDiscovery capabilities within the Office 365 stack	66%
The ability to have in-place eDiscovery capabilities within the Office 365 stack	63%
The ability to have in-place search and review eDiscovery capabilities across multiple vendors' solutions	53%
The ability to have in-place eDiscovery capabilities across multiple vendors' solutions	48%

Source: Osterman Research, Inc.

In this section we look at the capabilities of Office 365 in the areas of archiving, encryption, litigation hold and eDiscovery. Decision makers will need to examine how best to balance business, compliance, records, legal and IT goals when embracing Office 365, and should be aware of the limitations in Office 365 in these areas.

LACK OF ARCHIVING FOR SOME CONTENT TYPES

Archiving – moving business data out of one business system into a separate, secured location for optimized storage, immutability, and better data governance – is not offered for some important content types in Office 365. These include SharePoint, Skype for Business, additional message types, and third-party content. Organizations that require archiving capabilities should be aware of the following issues:

- SharePoint content, such as documents and list items, can be retained in place through retention policies, or moved to another location in SharePoint when it has expired or become irrelevant. These retention or move actions can be triggered based on specific date-based event triggers only, and for organizations staying within their assigned storage limits for SharePoint, SharePoint's In-Place Records Management in SharePoint may be sufficient. What is not possible, however, is to archive SharePoint content that is no longer current to alternative and cheaper storage systems. Although it is possible to purchase unlimited SharePoint storage capacity, it attracts premium pricing. Organizations with large quantities of SharePoint data are not well served if they want to keep their SharePoint content trimmed and current without incurring additional long-term SharePoint storage fees, or that want to archive content away from SharePoint Online based on event triggers beyond date-based metadata. Moreover, SharePoint is not write once, read many (WORM) compliant; a serious issue for organizations in regulated industries.

- Skype for Business Online relies on Exchange Online for archiving if specific conditions are met. No native archiving service for Skype for Business Online is available. By default, Skype instant messaging transcripts are retained in the Conversation History folder in each user's Exchange Online mailbox, but unless the mailbox is on legal or litigation hold, a user can delete their instant messaging transcripts at will, which doesn't provide an immutable or reliable archive of past messages. The need for legal hold to force the retention of Skype messages means that all Exchange Online mailboxes must be on hold at all times for this to work, which we consider to be an odd design. If a mailbox is on hold, peer-to-peer and multiparty instant messages are retained, as well as content upload activities during meetings. Other actions within Skype for Business are not retained, such as peer-to-peer file transfers, audio/video

for peer-to-peer instant messages and conferences, application sharing, and conferencing annotations.

- Text messages on BlackBerry devices will be archived into Office 365 if a third-party agreement is in place to capture these messages. Text messages on other devices, including iOS and Android, are not captured. With BlackBerry now having a low and dwindling market share in comparison to iOS and Android, capturing only BlackBerry messages is not as useful as it might otherwise be.

- Content from specific third-party messaging, collaboration, social media and other content sources can be archived into Exchange Online in Office 365 as converted email messages if agreements are in place with a third-party data partner. Messages are stored in the Exchange Online mailbox belonging to the specific user, and for content that cannot be tracked to a named individual, a catch-all mailbox is used. Most of the context of content from Twitter, Facebook, Yahoo! Messenger, DropBox and Salesforce Chatter is lost when these rich media sources are converted to email messages, making it difficult to re-create a historically valid chain of events.

BACKUP AND RECOVERY LIMITATIONS

Office 365 does not offer traditional backup and recovery capabilities in the same way as organizations have deployed in on-premises environments in the past because it is a live production system that offers recovery of messages and documents within a rolling time window. Instead, Microsoft uses alternative approaches for safeguarding current production data. For example:

- In Exchange Online, a user can recover a deleted item for up to 14 days by default (although an administrator can increase the recovery window to a maximum of 30 days).

- Data that is sent to the recycling bin from OneDrive will still be recoverable for 90 days, but only the most recent version of that data.

A different option is to use litigation hold or an indefinite legal hold to prevent any mailbox item from actually being deleted. The content will be hidden from the user's view when deleted, but it still exists in the mailbox. In SharePoint Online, there is also the ability to retrieve a deleted file within 30 days of deletion.

It's important to note that Microsoft does not offer point-in-time backup and recovery for organizations that want more traditional backup capabilities. Moreover, it cannot retrieve items that have been deleted beyond their recovery timeframe (assuming the mailbox is not on litigation or legal hold.) Other disaster-level scenarios are also not covered by Microsoft's service offering.

Office 365 does not offer traditional backup and recovery capabilities in the same way as organ-izations have deployed in on-premises environments in the past.

PRACTICAL STORAGE LIMITATIONS IN SHAREPOINT ONLINE

While SharePoint lists and libraries can hold up to 30 million items, there is a limit of 5,000 list items or documents that can be displayed in any one view. This was enforced in SharePoint Online to ensure that all tenants get good performance on SharePoint queries, but it has the practical implication of forcing unnatural content segregation design decisions by SharePoint developers within organizations to try to get around the 5,000-item threshold. It frequently means that end users are stopped from doing their work because the 5,000-item limit has been reached, or a lookup against a list with more than 5,000 list items has failed. This is a long-term issue for customers, and while Microsoft has been working recently to address this issue, it has suffered several false starts. Some customers are so frustrated by the 5,000-item list threshold that they are considering moving away from SharePoint Online entirely.

MESSAGE ENCRYPTION

Microsoft offers two message encryption services in Office 365, both confusingly called Office 365 Message Encryption. The soon-to-be-legacy Office 365 Message Encryption (OME) service was part of the higher-cost Enterprise plans, or as an add-on to other plans.

Under legacy OME, the message was sent as an encrypted HTML attachment – of up to 25 megabytes in size – that could be viewed only on the Office 365 viewing portal (some thought the viewing portal application for mobile was difficult to use and featured a poor UI experience). Legacy OME was powered by Azure Rights Management (Azure RMS). It was designed to enable an Office 365 user to send encrypted email to any recipient without having to know what email service, email client, or encryption capabilities they supported; the recipient's email address was used as the public key, rather than relying on a certificate infrastructure.

Legacy OME suffered from numerous weaknesses, including:

- The decision to encrypt a message was triggered largely by manual action on the behalf of the sender. He or she needed to include the word "encrypt" in the subject line (or something similar), which would then be captured by an Exchange transport rule configured to look for that key word. More automated options were also possible through Exchange transport rules, including the recipient being outside the organization and the presence of certain words or phrases in the message.

- On receiving a legacy OME message, the recipient had to save the HTML attachment, open it in a supported browser, and login to the Office 365 viewing portal using an Office 365 or Microsoft account, or request a one-time passcode. Access was also possible on iOS and Android mobile devices, using a special viewer app for OME messages; users on other mobile devices needed to use a supported browser. These additional steps were required even for other Office 365 users using Outlook 2016 for Windows, the premier and most advanced email client offered by Microsoft. There was no support for fully transparent and seamless delivery of encrypted messages between Office 365 subscribers in different organizations.

- Legacy OME was not able to track or alter what happened to a message after it was sent, meaning that a message could not be revoked, and the sender had no insight into what happened to the message. Even though special actions involving the Office 365 service were required by the recipient to access the message, no post-delivery status information was available to senders or administrators.

In summary, legacy OME did not offer a transparent, end-to-end encryption service that would automatically encrypt and decrypt messages for both senders and recipients without additional per-message steps and authentication requirements. Legacy OME was offered until September 2017, when it was replaced by "new OME." Both services have the same name, but are quite different in design.

New OME ties encryption to Azure Information Protection and Azure Rights Management in order to provide a singular method of sending encrypted messages inside and outside of the organization. It is designed to address some of the issues in legacy OME, such as working seamlessly in Outlook for Office 365 customers, and easing the sign-in restrictions to now also allow recipients to use a Google account or Yahoo! ID, in addition to the other pre-existing options.

It is still early days for new OME, but based on early experiences, we offer the following cautions:

- New OME will encrypt only attached Word, Excel, PowerPoint, InfoPath and XPS documents. No encryption or rights management capabilities are available for non-Office file formats, including PDF, and the document must actually be attached to the message; it cannot be referenced from OneDrive for Business or SharePoint Online.

- A manual action is still required for new OME to encrypt the message. The sender needs to select the "Do Not Forward" or "Encrypt" permissions policy in Outlook on the web, or another similar custom policy if set up. Administrators can also set Exchange transport rules to automatically apply encryption if an exact match to certain words or phrases are included in the message.

- There are two out-of-the-box policies in new OME. The Encrypt permissions policy applies encryption, but allows the recipient to forward, copy, and print the message. The second option of Do Not Forward explicitly ties together encryption and post-delivery rights management, which may be too restrictive for customer scenarios.

- Applying either the "Do Not Forward" or "Encrypt" policies only works in Outlook on the web. While Microsoft says that support for Outlook for Windows and Outlook for Mac are coming, that is not available yet. For users of the desktop apps, therefore, this will require changing their workflow to use the browser version of Outlook whenever a message needs to be encrypted.

- It is unclear whether the subject line of the message, if it contains sensitive information, will be protected through encryption. Legacy OME did not offer this capability, and the early evidence says that new OME does not either.

- DLP rules in the Security & Compliance Center cannot be used to automatically encrypt messages. Only Exchange transport rules (mail flow rules) in the Exchange Admin Center can be used. In other words, the newer tools in Office 365 for data security and protection cannot support new OME.

- There are still no post-delivery insights or reporting capabilities, nor the ability for the sender to revoke access to the message.

Finally, legacy OME will be deprecated at some point in the future. It is unclear what will happen to the messages sent using legacy OME technology, and for how long the ability to decrypt the message on the legacy Office 365 viewing portal will remain on offer.

BASIC LITIGATION HOLD CAPABILITIES

Legal and litigation hold in Office 365 offers only basic capabilities compared to some third-party offerings. Historically, Microsoft offered workload-specific legal hold capabilities for Exchange Online and SharePoint Online, but has recently created a new unified approach in the Security & Compliance Center. It is no longer possible to create new legal holds on SharePoint content from the previous SharePoint eDiscovery Center, and while Microsoft intends to similarly deprecate the ability to create new legal holds on Exchange content within the Exchange Admin Center, customer push-back has delayed its removal. The current In-Place Hold in Exchange Online enables the creation of multiple separate legal holds that are transparent to the user, and that can be based on different parameters such as time-based, search query-based, and indefinite (until further notice).

Legal and litigation hold in Office 365 offers only basic capabilities compared to some third-party offerings.

The litigation hold capabilities in Office 365 suffer from the following issues:

- Current legal holds created in Exchange or SharePoint cannot be migrated into the new experience in the Security & Compliance Center. They are separate objects that must run their course and then expire, rather than being something that can be pulled across for a unified view of current and outstanding legal holds.

- The litigation hold capabilities deal only with content in Office 365, but not content stored elsewhere. Organizations with significant data repositories outside of Office 365 – on-premises and in other cloud services – will require multiple, disparate systems for setting and apply legal holds, creating a complex legal compliance minefield.

- No workflow support for coordinating with data custodians across the organization who may have content that is responsive to the legal hold parameters. While these could be manually created and sent, no audit trail reporting would be created for subsequent review.

- Searches for responsive material are point-in-time, and do not automatically keep the result set up-to-date. Human intervention is required to re-run all current legal hold searches, and then apply a hold to new material.

- Office 365 can search and index only a specific list of file types. If non-supported file types are identified during a content search, they will be flagged for human review. Organizations with file types not on the supported list will face high manual analysis costs for document-by-document review to meet legal requirements.

- After searching for content in Exchange Online, the search preview pane will display a maximum of 200 items for an In-Place eDiscovery Search, listing the mailboxes and items found. However, these items cannot be displayed in the search preview pane; they must be exported to a discovery mailbox for review. Better in-line support for previewing messages directly from the search pane is not available.

- The advanced eDiscovery capability in Office 365 is not "in-place". The advanced tools provide eDiscovery capabilities within the suite of Office 365 applications and are not integrated directly into the data sources. Therefore, the effort is a two-step process, requiring a search and export for data using the limited Security & Compliance Center capabilities, selecting the advanced eDiscovery center as a destination before one can actually run the advanced tools. Therefore, there is no way to iterate and search on the source data without multiple manual repetitive blind operations.

- For content searches based on multiple keywords, the search results do not show which keyword triggered the inclusion of a specific item. The only way for an analyst to know which keyword was responsible in Office 365 is to set up multiple single keyword searches.

eDISCOVERY WORKFLOW

Microsoft offers a range of eDiscovery capabilities for searching for responsive material across Office 365, plus a more advanced eDiscovery service called Advanced eDiscovery that adds text analytics, machine learning, and relevance and predictive coding for early case assessment. The latter is available in the premium Enterprise E5 plan, and as an additional cost add-on to the Enterprise E3 plan. With its latest approach to eDiscovery through the Security & Compliance Center, Microsoft has removed some of the limitations from its earlier attempts to provide enterprise-class eDiscovery, such as limited search scopes (where a maximum of 10,000 Exchange mailboxes could be searched at once in an eDiscovery search), as well as separate eDiscovery tools for Exchange Online and SharePoint Online.

However, none of the eDiscovery tools in Office 365 provide a robust eDiscovery workflow process that will satisfy many organizations' requirements. For instance:

- There is no workflow or project tracking of an eDiscovery case, such as the status of the case, who is involved, and which tasks are being worked on and by whom.

- An eDiscovery case administrator has no ability within the Security & Compliance Center to send legal hold notification alerts, nor reminders or escalations. These have to be handled out-of-band. As above, the lack of workflow and project tracking capabilities is not ideal.

- All cases are created and managed in an ad-hoc way, with a compliance officer entering ad-hoc search terms. It is not possible to create a case template for repeatability and auditing, with standard search queries and locations, key actions and requirements to complete, and an audit trail of what was and wasn't done. This is of particular concern to organizations that are not doing eDiscovery all the time; the ad-hoc approach means that prior learnings and approaches are likely to be forgotten and overlooked in a current eDiscovery case, possibly exposing an organization to sanction for insufficient production of evidence.

- Exports from Office 365 are not protected and so are at risk of alteration and spoliation. The output is a raw native export and not in a preservation format, such as forensic image format, which many eDiscovery collection tools offer. Moreover, there are no additional encryption options provided by Microsoft to encrypt the export.

- Due to batch processing, searches using the native Office 365 functionality are fairly slow. It can take several minutes to run a single search and search time increases based on the number of mailboxes in question.

- The eDiscovery capabilities in the Security & Compliance Center take a unified approach to responsive content in three key Office 365 workloads only: Exchange Online, SharePoint Online, and OneDrive. Other workloads – such as Yammer, Microsoft Stream, and Microsoft Teams – are excluded. Further, an eDiscovery case created in the Security & Compliance Center cannot search for responsive content in non-Office 365 content repositories, such as those maintained on-premises or in other cloud services. This limited approach means that any organization with content outside of Office 365 – including SharePoint 2013 and 2016 on-premises – will need multiple eDiscovery tools, in addition to having to instantiate, perform, and coordinate multiple eDiscovery cases in each separate tool. This is an expensive, complex and error-prone situation.

- Customers have recently been given the ability to import non-Office 365 data for analysis into Advanced eDiscovery. This has to be organized in a particular structure, uploaded into Azure, connected through a series of manual steps, and then processed by Advanced eDiscovery. Once processed, additional new content cannot be added to the Azure container. Another separate non-Office 365 data import has to be organized instead.

- Searching Exchange Public folders is an all or nothing proposition. There is no ability to scope the search to a targeted list. This means far too much information will be exposed to eDiscovery managers.

- It is not possible to configure a more limited search scope for eDiscovery managers searching OneDrive and SharePoint Online repositories, and Exchange mailboxes. Any eDiscovery manager can search any OneDrive folder, SharePoint Online site, or Exchange mailbox anywhere in the world; these should be able to be restricted by geographical region or country to safeguard and protect data.

- It is not possible to set the search scope on email messages to exclude the signature block, so if a keyword appears in email signatures, it will generate a high rate of false positives. This is an annoying time waster for eDiscovery personnel, and expensive for the organization.

- Messages encrypted with rights management protections can be automatically decrypted at the time of export, but a separate export must be run to handle these messages as individual entities. The export of encrypted messages cannot take place in line with any other export activities.

- Search results for Exchange Online, SharePoint Online and OneDrive must be exported from Office 365 to facilitate the review process; the Exchange content as one or more PST files, and the SharePoint and OneDrive content as individual files (with an option for all versions). There are multiple problems with the Office 365 approach: it creates a duplicate set of content outside of Office 365 which must be protected, there is no reporting on actions taken on the exported content in the eDiscovery case in Office 365 because Office 365 is blind to post-export actions, if the search is run again in Office 365 then a subsequent export is required along with integration of multiple sets of data, and there is no connection between what was collected and the coding decisions made to that content in order to inform future cases and reduce the volume of potentially responsive content in Office 365. The need to export content to Azure – with the time delays that are introduced from Office 365 to Azure and then Azure to a local computer – creates unhelpful delays in an urgent process for compliance officers. With GDPR coming on stream in late May 2018, the potential existence of personal data in additional locations will raise significant data governance concerns.

Customers have recently been given the ability to import non-Office 365 data for analysis into Advanced eDiscovery in the next draft.

TENANT ARCHITECTURE AND DATA RESIDENCY CHALLENGES

From the beginning of Office 365, the design of the tenant architecture was that each organization used one and only one tenant, homed in one geographical region, and to which all out-of-region traffic would route for access to the organization's data. This design works perfectly for organizations that are solely active in one geographical region, but can cause significant data sovereignty and data residency challenges for multi-national and cross-regional organizations. The sole tenant location for the organization is set when the organization first signs up for Office 365, and even then, some content types in Office 365 have only been served out of the North American region, regardless of the organization's master region, although this is slowly changing over time.

What this means, therefore, is that under the original design, an organization with significant operations in multiple geographies cannot geo-ring fence content into local Office 365 data centers, which has implications for legal cases, government access, and compliance with data protection regulations. Organizations dissatisfied with the original design have until recently had only one other option, and that was to try to make multiple tenants homed in different geographical regions work as one. Setting up multiple, inter-related Office 365 tenants is a non-trivial technical undertaking, and has several negatives for actual usability. Microsoft has, in general, advised organizations not to pursue this route.

Microsoft used its Ignite 2017 conference to introduce a second and more tenable option for organizations for which one tenant was not a workable answer: Multi-Geo. Once out of private preview, Multi-Geo will enable large organizations (it is aimed at tenants with more than 10,000 Office 365 users) to use a single tenant as before, but with data and content segregated across multiple geographical areas. Multi-Geo is not a free service, and early indications are that the added cost is significant. Here's what we know or can ascertain based on early information:

- In the short term, Multi-Geo will apply only to Exchange Online and OneDrive for Business. The Exchange mailbox and user's OneDrive folder will be moved to the preferred data location set for the user. Since these two workloads are easily divided at the user level, Multi-Geo is conceptually easy to apply in each case, and should work almost seamlessly.

- After setting up additional geographies in a tenant, customers will gain the ability to tailor various policies at the geo level. This includes sharing policies in OneDrive and SharePoint, DLP policies in the Security & Compliance Center, and even eDiscovery managers.

- SharePoint Online is targeted as the third workload for Multi-Geo, but unlike Exchange and OneDrive, which are user-focused services, SharePoint is a team- or group-focused service, which makes some flow-on decisions about data access and data rights more complicated. Each geo-enabled location with SharePoint will have a unique URL namespace, which means that SharePoint access will be less seamless than for Exchange and OneDrive. And organizations with cross-geographical collaboration between employees will constantly have to ask which SharePoint location is the correct one for each new site.

- Some critical services, such as Exchange Online Protection, are not currently targeted as being Multi-Geo enabled. The current intent is that EOP processing will always happen in the tenant's default geo location, rather than being distributed out to each individual geo. Having all email route through scanning services in another geo location may not be good enough for large organizations.

- Multi-Geo is a good step in the right direction, but it doesn't yet deal with all of the workloads in Office 365. Multi-Geo customers will still need to figure out their data residency approach for Microsoft Teams, Skype for Business, Yammer, and other Office 365 services.

INDEXING FILE TYPES

As noted previously, Office 365 can index a specific list of 58 file types, which is weighted in favor of the various file formats in Microsoft Office products. When undertaking an eDiscovery search and performing an Early Case Assessment, any file that is not included in the 58 will be flagged as unprocessed. When applying DLP rules, file types not included in the 58 will not trigger the capture rules. The implication is the need for a manual review of these non-supported file types by a compliance or security officer, adding cost and decreasing timeliness of information exchange. Moreover, keyword searches may also miss relevant content due to the use of a "best-effort" index. If an organization makes regular use of non-supported file types, it should look at third-party tools that will index additional file types.

STORAGE OF AUDIT REPORTS

Office 365 offers a unified audit logging service across key workloads, and is accessed through the Security & Compliance Center. Auditing for most workloads is turned off by default (and thus must be turned on to start the process of collecting audit entries); one prominent exception is audit logging of administrator actions in Exchange Online which is turned on by default. Audit entries in the Security & Compliance Center are retained for 90 days, after which they are purged. A recent change to audit logging of Exchange items means that an administrator can set a higher (or lower) default period. Advanced Security Management – an integrated component of the Enterprise E5 license and an optional add-on for other plans – captures audit log data from Office 365 and moves it to Azure, but even then, such audit log entries are stored only for 180 days. Organizations that need long-term access to audit report items – such as seven years' worth of data under some compliance regulations – should be aware of the limitations of the Office 365 Audit Log service, namely:

- Audit log entries are purged after 90 days, except for Exchange Online audit items if an administrator has specified a longer retention duration.
- Querying the audit log system in Office 365 allows a maximum query period of 90 days. This cannot be changed.
- Exporting audit log items from Office 365 is limited to 1,000 entries unless all results are exported, for which the limit is 50,000 items. An organization with auditing turned on will generate at least 10-20 audit items per individual per day for a light user, and potentially a couple of hundred items per day for an active information worker. Some medium-sized organizations, let alone their larger counterparts, will hit the 50,000 item limit every day. In such a scenario, an administrator will need to specify and generate at least one export every day, and hope that the time delay in capturing audit report entries doesn't mean that items that should be collected are missed from the report.
- Exports are delivered as CSV files, the collection of which must be managed. Paradoxically, as an exported file of audit items, there is nothing to prevent an errant administrator from removing evidence of his or her own wrongdoing; the exported file does not guarantee authenticity of the historical information contained inside.

While Microsoft has increased its capabilities for the storage of audit reports over the past year, their handling of these reports is not as robust as that available from some third party vendors.

LICENSE REQUIRED FOR EX-EMPLOYEES' MAILBOXES

When an employee leaves an organization, but their mailbox must be retained, it was historically true that a full user license was still required to keep the mailbox. Microsoft has removed this licensing requirement, and so-called "inactive mailboxes" in Exchange Online can be retained free of charge. This means that an administrator can put a mailbox on legal hold and delete the associated user account; the mailbox is retained for the duration of the legal hold as an inactive mailbox without incurring any charge to the organization. However, Microsoft has signalled its intent to introduce a new license requirement for

Office 365 can index a specific list of 58 file types, which is weighted in favor of the various file formats in Microsoft Office products.

inactive mailboxes, originally scheduled to come into force from October 1, 2017, but for the time being has delayed the introduction of this cost. It is likely that inactive mailboxes will attract new licensing terms during the next 12-24 months.

OFFICE 365 AND GDPR COMPLIANCE

The European Union's (EU's) General Data Protection Regulation (GDPR), the soon-to-be-enforced data protection regulation covering personal data on EU data subjects, will have significant impacts for organizations doing business in the EU and elsewhere. Organizations using Office 365 will need to ensure the protections offered in the service are up to standard, or they may face punitive fines under the regulation. A holistic approach to data protection, both within Office 365 and beyond, will be necessary for GDPR compliance.

While GDPR will be enforced from late May 2018 and Microsoft has been investing heavily to get Office 365 and its other cloud properties ready for GDPR, there is a lot that is unknown about how GDPR will be enforced in practice. In examining the capabilities offered for security, archiving, encryption, compliance and data protection in Office 365, the following strengths and weaknesses are evident in advance of GDPR's enforcement date:

- Office 365 offers various capabilities for identifying sensitive information across Exchange Online, SharePoint Online, and OneDrive for Business, using the more than 80 pre-built sensitive information types in the Security & Compliance Center. Advanced Data Governance, a service included in Enterprise E5, can proactively and automatically apply sensitivity labels to data as it is being created. For organizations using Enterprise E5, these capabilities will help with the data discovery challenge of GDPR.

- While not part of Office 365, Microsoft's Azure Information Protection Scanner will periodically scan on-premises file servers and repositories for sensitive, confidential and protected data. This will highlight to data controllers what personal data is currently being stored in on-premises systems, and therefore where data protections will be needed. These scan results will also help in planning for migrating to Office 365, Azure or other cloud services, highlighting to where sensitive information will be moving.

- When a DLP policy identifies sensitive information in a document in SharePoint Online or OneDrive for Business, it will block access to the data to everyone but the document owner, last modifier, and the site owner. While this will indeed protect personal data, it will not address use cases where people other than those three have valid business reasons for accessing the personal data in a document contained in a secured SharePoint Online or OneDrive for Business site. Likewise, sensitive data in a document cannot be sanitized while leaving the rest of the document available for review, or partially encrypted to prevent unauthorized access. In summary, Office 365 offers broad and basic ways of applying data protection policies within the organization, but it lacks the nuance, panache and elegance that complying with the GDPR will require.

- DLP policies that identify sensitive information will also lock and block documents in SharePoint Online and OneDrive for Business to prevent them from being shared with external users. This will be the appropriate action to take in some use cases, but not all. For example, there doesn't yet appear to be a way to check if a valid sharing agreement is in place between the organization and external firms or specifically named individuals. End users will need to do out-of-band checks to see whether they can transfer data or not.

- Service integrity and resilience to protect against threats to personal data is a matter of interest in GDPR. From a GDPR compliance perspective, the questions above about whether services like Office 365 ATP are good enough to protect end users from malicious links and attachments become much more than an exercise in comparing

feature effectiveness between competitive offerings. If personal data is compromised in Office 365 because ATP is not good enough, that becomes a real problem for organizations.

- Encryption is specifically mentioned in the GDPR as a method of reducing the impact of personal data being breached, stolen, or inadvertently shared with unauthorized recipients. Beyond its role in doing so, it's a good practice for protecting all types of data. Office 365 uses encryption at many levels to protect data in Office 365, offers Office 365 Message Encryption (for user and policy-based encryption, with some provisos as explored above), and customers newly have the choice of bringing their own encryption keys to add a further level of protection. Since the destruction of a customer's encryption key has catastrophic consequences for access to data in Office 365 (which in itself is a problem under GDPR), organizations will need to ensure appropriate controls are in place to ensure the customer's master encryption key is not compromised in a ransomware or credential phishing attack.

- GDPR is a much more expansive issue than just Office 365. Microsoft's own positioning of its offerings for organizations wanting to work towards GDPR compliance is Microsoft 365, which combines Office 365, Windows 10 (including capabilities like Windows Information Protection), device protection and more. Even Microsoft acknowledges that while Office 365 will need to comply with GDPR requirements, it is not the complete story for organizations.

- Complying with GDPR will require organizations to gain and maintain a holistic and real-time view of data protection threats across all cloud services, applications, endpoints and devices. There is no great gain from a data protection perspective if end users can save documents containing sensitive information to thumb drives or alternative cloud storage locations and use those locations to circumvent Office 365's data protection controls. Microsoft offers some capabilities in these areas, including the Office 365 Cloud App Security and the broader Microsoft Cloud App Security service, as do other vendors. Many employees also grant access to unapproved third-party applications and add-ins that integrate with Office 365 and other SaaS applications. Best-in-class solutions can give organizations visibility and control when it comes to third party applications that may be inappropriately accessing and storing data.

- The data protection requirements of GDPR will bring to light poor data protection practices of modern organizations. For example, storing personal data on customers or subscribers in ad hoc and unsecured Excel spreadsheets is a poor practice compared to using a secured database with field-level encryption and pseudonymization. Perhaps Microsoft's approach to locking and blocking all documents in SharePoint Online and OneDrive for Business that contain sensitive information will prove to be an effective way of forcing organizations to improve their own internal data management and data protection practices.

- The right to be forgotten is one of the core rights of data subjects under GDPR, and means that under certain conditions, all applicable personal data on a given individual must be deleted. However, this requirement is highly nuanced, in that applicability is defined by the legal basis under which the data was originally collected. Applying a blanket deletion to all personal data for the individual is not the intent of the regulation; a highly targeted operation is required instead. Technologies for deleting data in Office 365 will provide brute force capability to ensure a data subject is forgotten, but this must take place only within a strong data governance framework where data provenance requires the deletion action. How Microsoft will address this nuance in Office 365 remains to be seen.

- Until 2017, global organizations were advised to choose one master location for their Office 365 tenant, meaning that all access from outside the region would backhaul across Microsoft's global network. The alternative for organizations with regional compliance and data protection requirements was to try to make multiple tenants work somewhat seamlessly together, which was possible, but messy. With the introduction

GDPR is a much more expansive issue than just Office 365.

of Multi-Geo, albeit still in preview, large global organizations have a new possibility for segregating data access, DLP policies, and sharing policies across Office 365. This may prove to be a beneficial change for organizations with significant operations in Europe and other regions of the world, although Multi-Geo is enabled only for some Office 365 workloads, and services like Exchange Online Protection and ATP are not offered in all geographies. Multi-Geo is currently positioned for organizations with more than 10,000 Office 365 users, but even organizations with 250 employees distributed around the world may benefit from data protection policies and data residency on a regional basis.

Organizations that need to comply with GDPR from May 2018 would be well advised to consider alternative data privacy and protection capabilities beyond those offered in Office 365. While Office 365 will eventually offer more robust and nuanced protections, GDPR needs to be addressed now.

OTHER LIMITATIONS

- **Monitoring**
  Traditional monitoring solutions focus almost entirely on the infrastructure supporting a specific service or application. In an on-premises scenario, there is typically a full understanding of what each piece of the infrastructure is responsible for, what the relation is between components, and what "normal" behavior looks like for each component. In practice, this approach proves to be much less effective with cloud-based services.

  In a cloud-based system like Office 365, visibility is limited, the relationships between key system components are largely unknown to the customers, and the normal performance baseline for the infrastructure components are not readily shared by Microsoft. Secondly, the massive scale of a service like Office 365, coupled with the way users are distributed across several datacenters and hundreds of thousands of servers, make it nearly impossible to maintain the same monitoring paradigm. Within a sea of information, administrators cannot correlate what information is relevant to their organization and what isn't, in part because they don't have complete information about all the components.

  Many of the problems faced by Office 365 customers are not caused by anything Microsoft does or doesn't do. We can classify Office 365 problems or outages into one of three categories:

  - A problem with an on-premises system or hybrid component, such as an AD FS server, a directory synchronization server, or an Exchange server. Problems in this category can be diagnosed and fixed solely by the organization that owns the component – Microsoft and other outsiders can't see or fix these issues.

  - A problem with Internet connectivity between a user and Microsoft. Because users can work on a wide variety of public and organizational networks, these problems can be hard to troubleshoot and may not be fixable by the organization if it isn't on their own network.

  - A problem with a server or component that Microsoft owns and maintains such as an Azure network routing outage.

  Rapidly identifying emerging problems and understanding which of the three classifications they fall into is key to driving down mean time to resolution. It is important to keep in mind that quickly resolving the root cause requires the effort of cross functional teams within the organization, such as the cloud services, networking, security, and messaging departments.

  Because of the fundamental differences between how an on-premises application or a cloud-based system are managed, an *entirely new monitoring approach is required*. Cloud-based systems, such as Office 365, enable organizations and users to work from

virtually anywhere. Because of this, monitoring a service from a specific location, typically the organization's datacenter, no longer represents how applications are used in the real world.

The customer-centered approach to monitoring cloud services maintains a laser focus on measuring and reporting on the end-user experience. The modern, customer-centered approach injects probes into the locations that the customer specifies to carry out typical end-user tasks and reports back on performance. These end-user experience probes provide the necessary data and resulting analytics to ensure complete visibility into performance and service quality at each individual location. Monitoring the experience that end-users have through synthetic tests when using the Office 365 service is critical to identifying and localizing problems. After all, the ultimate measure of any cloud-based service is whether or not the service is available for end-user consumption.

To its credit, Microsoft continues to enhance the monitoring capabilities within Office 365, but customers quickly realize that the out-of-the-box features like the Service Health Dashboard do not provide a sustainable monitoring solution. There is still a substantial amount of time and expense to configure additional services, such as OMS and custom Power BI dashboards to integrate with the Service Health dashboard to obtain a reasonable overview of the Office 365 platform. Even with the additional time and cost to stitch together three different datasets, a complete end-to-end monitoring of hybrid scenarios is not obtained. Additionally, this approach does not provide insight into service quality between the user's location and the Office 365 platform.

In short, administrators must be able to determine what caused an outage or service slowdown so that they can respond appropriately to issues that come up, and so that they can minimize the time required to resolve an issue. Customer-centered monitoring that leverages end-user experience probes, along with real-time synthetic tests, are critical in determining where the problem lies. In the absence of modern monitoring capabilities, quickly understanding where problems are occurring and who is affected may not be obtainable.

- **Supervisory review for FINRA**
  Certain industry regulations, such as those enforced by the Financial Industry Regulatory Authority (FINRA), require the capture and review of communications between particular people, or people in a specific group, to ensure no nefarious or unauthorized topics are being disclosed or discussed. Office 365 previously offered a Supervisory Review capability that could work with Exchange Online messages, which had a range of issues. Microsoft has recently replaced this legacy Supervisory Review capability with a new Supervision tool that requires the Enterprise E5 plan or the Advanced Compliance add-on. We note the following concerns with the new Supervision offering:

  - Every person who is to be covered by a Supervision policy requires an Enterprise E5 license, or the Advanced Compliance add-on. This is a per-user licensing requirement, not an organizational-level option.

  - Supervision works only with Exchange Online in Office 365, but does not address Microsoft's other communication tools, such as Microsoft Teams, Yammer and Skype for Business. This scope of coverage is too narrow in our opinion.

  - Once a supervision policy has been set up, a private shared mailbox is provisioned for receiving captured messages. Supervisory reviewers must connect to the shared mailbox to review and assess each message.

  - It is not possible to use Microsoft's sensitive information types in Supervision policies.

  - When searching for words or phrases, these must match exactly. A misspelt variant will not trigger the supervisory rule. It would be useful if Office 365

Supervision works only with Exchange Online in Office 365, but does not address Microsoft's other commun-ication tools, such as Microsoft Teams, Yammer and Skype for Business.

offered the ability to use fuzzy matching to give a broader impression of what else what happening through Exchange Online.

- Supervisory review works only in Outlook on the web. Although an Outlook client add-in has been promised (and one is available that can be installed, albeit with PowerShell commands), it is non-functional and doesn't work.

- There is no migration support between the old Supervisory Review feature and the new Supervision feature. Policies from the previous approach have to be deleted; they cannot be migrated and updated, and they are not automatically updated by Microsoft.

- While messages are captured for post-delivery or after-the-fact review, there is no ability to quarantine an offending message and have it routed for approval before release. The damage could already be done, since the message has actually been sent and delivered.

While Supervision is positioned as a significant upgrade to the previous Supervisory Review capability in Office 365, the above analysis suggests its capabilities will not be adequate for many organizations.

- **Directory sync issues**
  Azure AD Connect replaced Windows Azure Active Directory Sync (DirSync) and Azure Active Directory Sync (Azure AD Sync), both of which reached their end of support by Microsoft in April 2017. While Azure AD Connect offers useful capabilities, it does have limitations that some third party tools do not. For example, Azure AD Connect does not support failover clustering or automatic failover, it may not offer adequate information for some admins in its event logs, it relies on less secure SSL/TLS encryption for communications with Azure AD, and it requires an Enterprise Admin account in multi-domain and multi-forest environments. Some third-party directory sync tools may be more adequate.

- **Continuity issues**
  Since Office 365 customers may experience periodic service outages, as is the case with any cloud-based platform, a robust business resilience plan, including an email continuity solution, should be implemented. Additionally, outages can introduce a security risk as employees turn to personal email to conduct business during downtime.

SUMMARY

Office 365 provides core and widely-used services for productivity and collaboration to the modern organization, along with capabilities for content archiving, data security, encryption, and eDiscovery, among others. Microsoft has been successful in bringing to market and improving Office 365's capabilities over the past several years. However, even in light of recent updates to Office 365 at the end of 2017, organizations assessing the capability of the platform to meet their requirements in 2018 must be cognizant of areas where third-party solutions will offer better functionality. We have reviewed and explored the impact of these issues in this white paper.

In conclusion, we offer three closing statements:

1. While Office 365 offers core services for productivity and collaboration, it is not a complete offering for archiving, data compliance, and eDiscovery. Microsoft is motivating its customers to adopt add-on services across its cloud services portfolio to gain more advanced capabilities, many of which are not as advanced as those offered by third parties.

2. Microsoft's advanced services for data archiving, compliance, and eDiscovery, among others, will not fully satisfy every organization's requirements. Organizations must

consider supplementing Office 365's basic capabilities in these areas with best-in-class, third-party offerings. In particular, many third party offerings will offer more robust protection against targeted and highly sophisticated attacks than will Microsoft's offerings.

3. It takes only one malicious message that gets through Microsoft's basic and advanced capabilities to wreak havoc on an organization, or one malware-less attack that results in a large financial payment to a malicious actor to completely out-spend in remediation what could have been spent at a lower cost for prevention. In the current environment, this is the core question: do you spend now to create defense and protection, or do you spend later to clean up damage (and try to save your organization's reputation and brand value)?

SPONSOR OF THIS WHITE PAPER

ZL Technologies enables large organizations to govern enterprise content and satisfy corporate needs for regulatory compliance, eDiscovery, records management, analytics, and file analysis. ZL's singular and unified information governance architecture consolidates all applications and billions of documents under one platform, thus eliminating today's fractured data silos and allowing organizations to more effectively protect and manage data privacy. Demonstrating a proven track record with Global 500 customers and strategic partnerships with major players, ZL has emerged as the technology leader in information governance and harnessing big data for strategic advantage. For more information, please visit **www.zlti.com**.

www.zlti.com
@zltechnologies
+1 408 240 8989