AvePoint

COLLABORATE WITH CONFIDENCE

# 4 *Strategies* for Cloud Storage Optimization

EXTERNAL-CORP V1 | 7.2022

AvePoint®

# 4 Strategies for Cloud Storage Optimization

# Introduction

With a pandemic bearing down on the global economy, organizations had no choice but to accelerate their digital transformation journeys. As corporate offices fell silent and employees set up shop at their dining room tables, IT admins rushed to implement cloud-based solutions that would enable flexible, remote work. Tools that may have previously been used from time to time, like chats and video conferences, suddenly became a lifeline to communicate with colleagues from afar. In the first three months of the COVID shutdown alone, Microsoft Teams saw more than 50 percent growth in daily active users. While employees adapted to this fundamental change, a seismic shift was happening behind the scenes as well.

Microsoft Teams saw more than **50%** growth

*in daily active users*

Solutions like Microsoft Teams, OneDrive, and SharePoint made it easier than ever for employees to communicate and collaborate, no matter where they were. But it quickly became apparent that these tools also produced an unprecedented amount of data. Think about your own personal data footprint, for example. Documents and emails probably come to mind. But what about the dozens of Teams chats that you have going at any given time? Or the meetings that you record for posterity each week? These files consume a significant amount of storage, exponentially more than what organizations previously experienced.

**One-third of companies** *have cloud service budget overruns of up to 40%.*

Businesses now find themselves facing significantly higher data storage costs—and even costly overage fees. In fact, a 2021 report found that more than one-third of companies have cloud service budget overruns of up to 40 percent. While storage budgets will likely never return to pre-pandemic levels, you can mitigate your data volume by proactively applying robust data governance and information lifecycle principles—and we're here to show you how.

## Why is Storage Optimization Important?

When you think of storage optimization, you're most likely thinking about the cost reduction you could achieve. While this is certainly a common goal, it's not the only one. Proactively managing your organization's data footprint can actually produce several benefits, including:

**Cost savings.** As the data associated with workspaces, files, chats, meeting recordings, and more continues to skyrocket, holding on to extraneous content can have a significant impact on your organization's bottom line. Did you know that:

- In Microsoft 365, depending on your licensing arrangement, each user could have up to 5TB of storage across all workloads, including email, OneDrive, videos, etc.?
- Each licensed user can create a maximum of 250 Teams?
- A maximum of 500,000 Teams can be created per tenant?

At first glance, you might think these limits are untouchable. However, if you belong to a large enterprise with thousands of employees, you'll race past those numbers in no time if proper controls are not implemented. Another thing to keep in mind? Data storage costs aren't as simple as "X dollars per terabyte." To understand your all-in storage costs, you should also factor in the following expenditures:

- ***Replication.*** Is your organization replicating critical datasets across different regions? If so, you're effectively multiplying the amount of data stored—and your costs, as well.
- ***Transfer.*** Are you regularly transferring data from one storage provider to another? Are you planning for a future data migration? You may be incurring transfer or egress charges.
- ***Management.*** Does your storage provider charge you fees for management activities? These may include moving data across different tiers of storage, monitoring for data integrity, and performing security testing. Typically charged per terabyte, these fees can add up over time.

**Risk management.** One of the best ways to protect your organization from risk is to ensure that you're not holding on to too much—or too little—information. In addition to maintaining compliance with regulations such as GDPR that mandate how long sensitive data should be held for, applying appropriate information lifecycle processes enables you to reduce your risk exposure, including:

- ***Regulatory risks.*** Holding onto personal information for too long can have significant financial and reputational implications. Privacy regulations typically dictate that personal information is only to be kept for as long as is necessary for the purpose it was obtained. Another example: In the United States, Sarbanes-Oxley data retention policies mandate that financially regulated data is only kept for seven years.
- ***Security risks.*** If your environment contains outdated file formats that are no longer maintained or patched, you may be opening the door to additional security risks, including malicious actors or system failures.
- ***Process risks.*** The more cluttered your environment, the harder time your employees may have finding accurate, up-to-date files. Any activities or decision-making based on this ill-informed data could lead to legal, financial, or reputational damage.
- ***Improved performance and efficiency.*** It's very easy for an environment to get bloated by redundant, obsolete, or trivial ("ROT") content, which can cause intense frustration for users struggling to find up-to-date information. A recent survey sponsored by Citrix found that employees spend 25% of their time searching for information they need to do their jobs. By getting rid of this stale content, organizations can drive productivity, reduce risk, and improve the speed and accuracy of search, analytics, and e-discovery solutions.

Employees spend **25%** of their time *searching for information they need to do their jobs.*

# I. Say Good Riddance to ROT

The collaboration workspaces associated with short-lived projects, engagements, and initiatives—and the content contained within—often linger long after a project is over. The first step toward shrinking your data footprint is to identify any redundant, obsolete, and trivial information ("ROT") that may be cluttering up your environment. But what is ROT exactly?

**Redundant:** Due to the lack of versioning available in file systems, there are usually many different copies or self-made versions of content floating around. Possibly the most difficult part of the cleanup process, identifying the locations of this content and discerning what stays and what goes can help play a part in reducing content.

**Outdated:** End users tend to hold on to content that is no longer relevant solely because there may not be a process in place that forces them to dispose of it. Additionally, when people leave departments or organizations, their files often aren't evaluated and/or purged. These conditions lead to a mass of stale content that can likely be eliminated, assuming it doesn't go against your organization's record retention policy.

**Trivial:** The identification of trivial content is a subjective process. Reducing this type of data requires coaching end users on how to determine whether certain files should be removed.

## Weighing Risk vs. Value

Information managers may want to start by categorizing information into broad buckets of active items that are essential to day-to-day business, archived or stored items that are relevant but not accessed regularly, and ROT that can be slated for disposal.

A useful tool for making these decisions is a risk/value matrix or framework that evaluates content based on value- and risk-level. For example, content identified as low-risk and low-value that has not been accessed or modified in a certain period of time (e.g., 18 months) might be moved to archival storage prior to eventual disposal, whereas high-risk, high-value content would be retained.

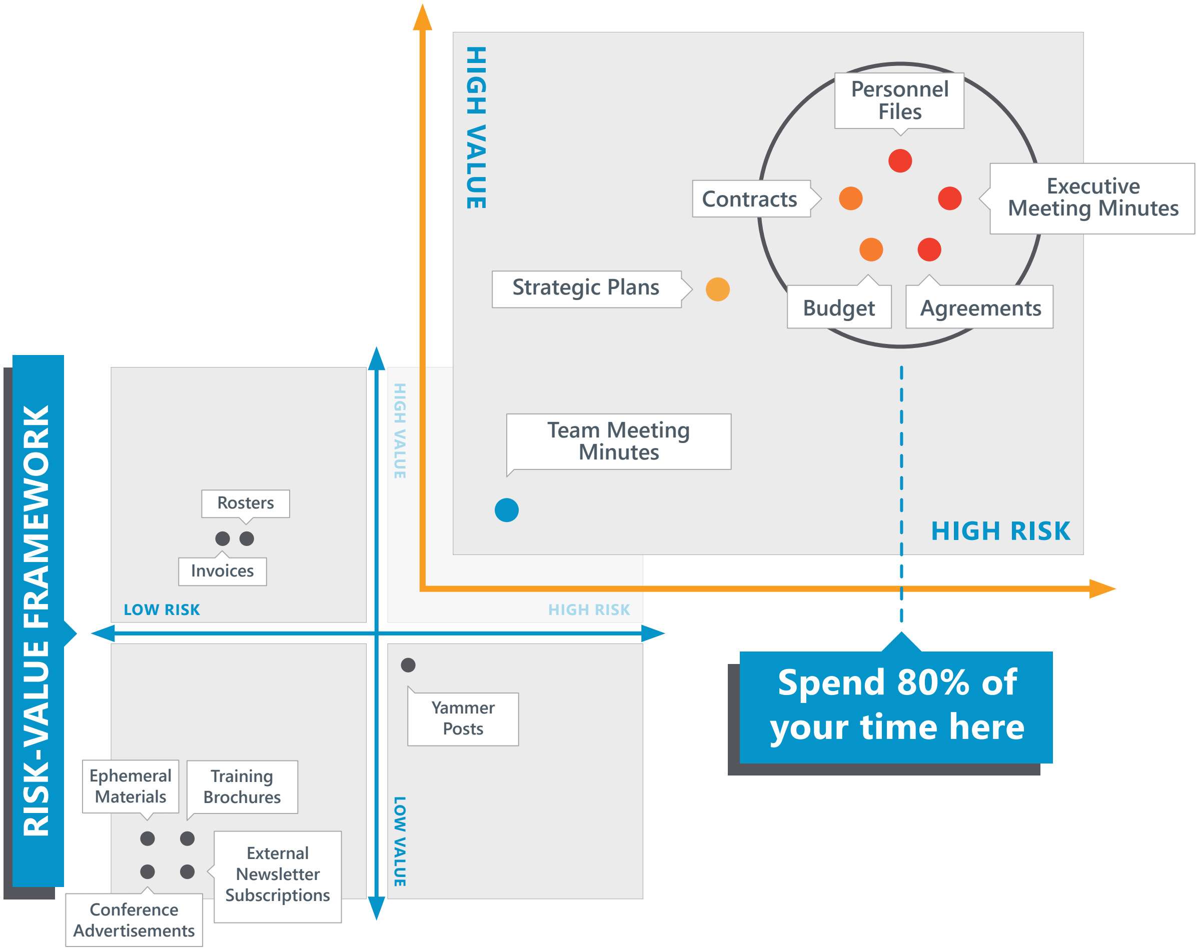

## Work Smarter, Not Harder

Now, it would be quite cumbersome to manually review all your data and identify what to keep and what to retire, especially since first you need to know where these files live. In addition to your Microsoft 365 tenant, you will almost certainly find ROT files in your local folders and other locations.

> **By discovering, mapping, and classifying** *unstructured data, organizations can make more informed decisions regarding which data to keep and which to remove.*

You can utilize available tools—some are even free—to help you identify this data in a very systematic way. A few suggestions:

- **Microsoft SharePoint Migration Assessment** is a tool you can use to scan SharePoint, show all your data, and help you with ROT identification.
- **AvePoint's free discovery tool** helps organizations identify ROT within their on-premises and cloud environments.

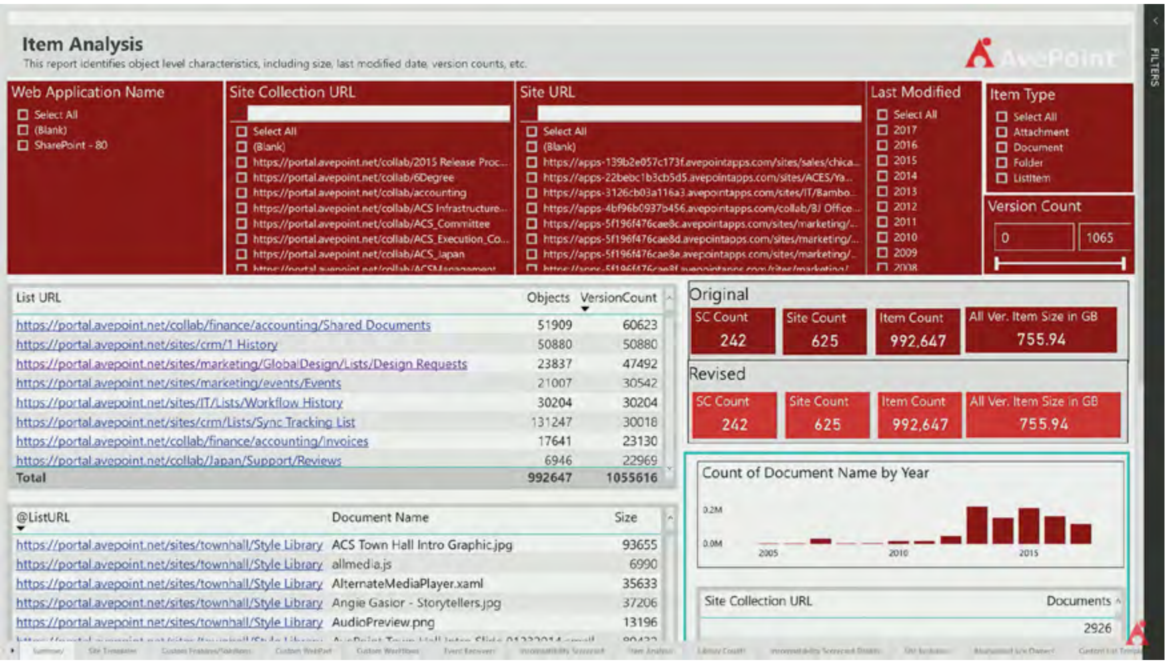

**Item Analysis**

This report identifies object level characteristics, including size, last modified date, version counts, etc.

AvePoint
FILTERS
Web Application Name
□ Select All
□ (Blank)
□ SharePoint - 80
Site Collection URL
□ Select All
□ (Blank)
□ https://portal.avepoint.net/collab/2015 Release Proc...
□ https://portal.avepoint.net/collab/6Degree
□ https://portal.avepoint.net/collab/accounting
□ https://portal.avepoint.net/collab/ACS Infrastructure...
□ https://portal.avepoint.net/collab/ACS_Committee
□ https://portal.avepoint.net/collab/ACS_Execution_Co...
□ https://portal.avepoint.net/collab/ACS_Japan
Site URL
□ Select All
□ (Blank)
□ https://apps-139b2e057c173f.avepointapps.com/sites/sales/chica...
□ https://apps-22bebc1b3cb5d5.avepointapps.com/sites/ACES/Ya...
□ https://apps-3126cb03a116a3.avepointapps.com/sites/IT/Bambo...
□ https://apps-4bf96b0937b456.avepointapps.com/collab/BJ Office...
□ https://apps-5f196f476cae8c.avepointapps.com/sites/marketing/...
□ https://apps-5f196f476cae8d.avepointapps.com/sites/marketing/...
□ https://apps-5f196f476cae8e.avepointapps.com/sites/marketing/...
Last Modified
□ Select All
□ 2017
□ 2016
□ 2015
□ 2014
□ 2013
□ 2012
□ 2011
□ 2010
□ 2009
□ 2008
Item Type
□ Select All
□ Attachment
□ Document
□ Folder
□ ListItem
Version Count
0
1065

| List URL | Objects | VersionCount |
|---|---|---|
| https://portal.avepoint.net/collab/finance/accounting/Shared Documents | 51909 | 60623 |
| https://portal.avepoint.net/sites/crm/1 History | 50880 | 50880 |
| https://portal.avepoint.net/sites/marketing/GlobalDesign/Lists/Design Requests | 23837 | 47492 |
| https://portal.avepoint.net/sites/marketing/events/Events | 21007 | 30542 |
| https://portal.avepoint.net/sites/IT/Lists/Workflow History | 30204 | 30204 |
| https://portal.avepoint.net/sites/crm/Lists/Sync Tracking List | 131247 | 30018 |
| https://portal.avepoint.net/collab/finance/accounting/Invoices | 17641 | 23130 |
| https://portal.avepoint.net/collab/Japan/Support/Reviews | 6946 | 22969 |
| **Total** | **992647** | **1055616** |

Original

| SC Count | Site Count | Item Count | All Ver. Item Size in GB |
|---|---|---|---|
| 242 | 625 | 992,647 | 755.94 |

Revised

| SC Count | Site Count | Item Count | All Ver. Item Size in GB |
|---|---|---|---|
| 242 | 625 | 992,647 | 755.94 |

Count of Document Name by Year

| @ListURL | Document Name | Size |
|---|---|---|
| https://portal.avepoint.net/sites/townhall/Style Library | ACS Town Hall Intro Graphic.jpg | 93655 |
| https://portal.avepoint.net/sites/townhall/Style Library | allmedia.js | 6990 |
| https://portal.avepoint.net/sites/townhall/Style Library | AlternateMediaPlayer.xaml | 35633 |
| https://portal.avepoint.net/sites/townhall/Style Library | Angie Gasior - Storytellers.jpg | 37206 |
| https://portal.avepoint.net/sites/townhall/Style Library | AudioPreview.png | 13196 |

| Site Collection URL | Documents |
|---|---|
| | 2926 |

## II. Implement Lifecycle Management

Once you've gutted your ROT, you want to make sure you are keeping your newly cleaned environment tidy. Information lifecycle management is the practice of assigning rules to different information types that dictate what will happen to it. This might include moving something to a new location, making it immutable through a record declaration process, or running disposition and destroying the content altogether.

> **The most effective** *lifecycle management strategy is twofold, addressing both the workspace itself, as well as the content it contains.*

First, apply an application lifecycle policy to the Microsoft 365 container (Team, Group, Site) itself. Second, manage Microsoft 365 content by applying business rules across workspaces, to manage the information lifecycle upon creation/upload within that workspace.

## Control Your Workspaces

As Teams adoption continues to grow, organizations that permit self-service provisioning will see more and more Teams being created. Without proper employee training and education, these workspaces can multiply exponentially, leading to data sprawl. Unfortunately, once these collaboration workspaces have served their purpose, they often fall by the wayside. If no one (or nothing) intervenes to clean them up afterward, it results in a large amount of redundant data left within the tenant. This not only drives up storage costs but can cause difficulties for users when searching for appropriate content.

To apply appropriate information lifecycle management policies, organizations must accurately capture, track, and record all Microsoft 365 collaborative workspaces as they are created.

AvePoint
Cloud Governance

AvePoint's **Cloud Governance** empowers users to create Teams, Sites and Groups in real-time, backed by a sustainable, efficient, and secure governance framework. Admins can track and easily report on what collaborative workspaces they have, why it exists, who it belongs to, when that information was last verified, and more. Cloud Governance ensures that data owners and admins always have a current and complete inventory of all Microsoft 365 workspaces.

## Tailoring Control Across Departments

| Department A | Department B | Department C |
|---|---|---|
| No external sharing | External Sharing In:<br>👥 T S | External Sharing In:<br>S |
| Expires After:<br>6 Months | Expires After:<br>12 Months | Expires After:<br>9 Months |
| Team Creation:<br>Central IT | Team Creation:<br>Dept IT | Team Creation:<br>Users |
| Member Recertification:<br>3 Months | Member Recertification:<br>6 Months | Member Recertification:<br>12 Months |

Organizations should right-size their policies to ensure that provisioned assets have correct classification, retention, metadata, and access controls in place from the outset. Going forward, workspace owners can also be automatically asked to review permissions and metadata and revise or confirm as necessary. Admins can oversee configuration settings, membership, and ownership change requests, with the ability to delete, revert, or notify of unauthorized changes. Plus, you can proactively mitigate extraneous content (and unnecessary storage costs) with automated, structured end-of-life processes that trigger alerts for potentially idle or irrelevant sites and content.

While widespread collaboration is a hallmark of today's digital landscape, organizations must take steps to proactively manage their storage footprint with robust data governance policies.

AvePoint

# Governance Automation

**User Request**
Lifecycle Management

## SCENARIO

Due to organizational changes, Bob needs to move some documents from his now-outdated team project site to a new site.

**Step 1:** Bob logs into DocAve Governance Automation and requests a new site. Once the request is approved by his manager, the new site is created based on the associated policy that automatically sets the site quota, lease, and retention duration.

**Step 2:** Through content management services, Bob has folders with documents relevant to the project moved to the new site. He archives the rest.

**Step 3:** To make sure the right people have access to this new site, Bob receives a recertification prompt to review permissions.

**Step 4:** Governance Automation notifies Bob that the lease for the old site has expired. Since the site is no longer used, Bob chooses to move the site and its content to more cost-effective storage.

### Your Users

Easily create new collaboration spaces that's tailored to their business needs.

### Your Organization

Save time and cost by streamlining IT services throughout the lifecycle of each site and site collection.

### Your Administrators

Less IT disruption servicing user requests allows greater focus on more valuable business activities.

www.avepoint.com

## END RESULT

FREE EBOOK

## Microsoft 365 Records Management Roadmap

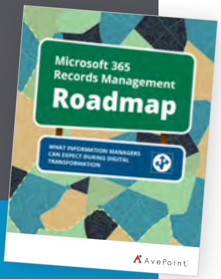

Microsoft 365 Records Management

Roadmap

WHAT INFORMATION MANAGERS CAN EXPECT DURING DIGITAL TRANSFORMATION

AvePoint

As the world continues to adapt to cloud-based collaboration, records and information managers are no exception. Information management practices need to evolve with today's digital transformation trends. Featuring real-life case studies, this eBook dives into planning for, executing, and managing a migration to a digitally centered Microsoft 365 records management system.

Learn how to navigate common records compliance requirements, create retention and disposal schedules, and ensure secure information management.

## Command Your Content

Just as workspaces go through an information lifecycle process, the content within them should be subject to the same. From the very moment a file is created in (or uploaded to) a Microsoft 365 workspace, it enters the information lifecycle. From creation to classification and retention, automated business rules ensure that all data follows a consistent, previously agreed upon process.

With **Cloud Records**, you can easily automate retention and disposal rules to manage your entire content lifecycle without end user intervention. Content from electronic sources including Microsoft 365, SharePoint on-premises, Exchange online, social media, and file systems can be managed directly alongside physical records in a centralized, intuitive platform. With comprehensive oversight and extensive out-of-the-box reporting, organizations can not only understand the current state of lifecycle processes, but easily adjust approvals, terms, policies, and more.

Remember, managing information across an organization is more than just managing the documents themselves. Without knowing the context (i.e., the containers or workspaces where those files live), we wouldn't be able to group related information together.

Consider this example: A contract that was created in a SharePoint library lives out its lifecycle there. However, the approval to move ahead with that contract was likely received via email. Therefore, the approval will sit in the Exchange mailbox, separate from the contract itself, unless the user moves that email to the SharePoint library (which rarely happens). If the users involved in the creation and approval of the contract leave the organization while these records are sitting in separate repositories, anyone else coming in who needs access to that information is not going to have a complete picture. They may have access to the contract itself, but unless they have access to the individual Exchange mailbox, they won't have any visibility into the approval process.

▶ So you see, **content and container are fundamentally linked**. In a collaboration platform like Microsoft 365, the workspace (Team, Library, Group etc.) provides critical context about the topic or subject that a piece of content is related to. Not only does this context empower users to make informed decisions, it also helps them store information in the correct location—all of which contributes to a clean, organized environment.

# III. Ensure Defensible Destruction

For many organizations, archiving is the most common solution for inactive workspaces or old data. While using different storage tiers is certainly an integral piece of the data storage solution, keep in mind that archiving just moves your problem to another place. It may produce short-term savings, but it's not sustainable as a long-term strategy. Holding on to information that is no longer useful to the business can actually create more risk than deletion.

The destruction of records is a very necessary, incredibly important activity that must be conducted in a way that can be defended in the future, if required. If records are destroyed improperly or if there's any appearance of "spoliation" (the deliberate spoiling of potential evidence in litigation), the organization could face allegations of attempted obstruction.

Operational issues can also arise when critical records are inadvertently destroyed. Imagine you go to pull up a report from two years ago, only to find it's not there because another user deleted it, thinking it was "old" and no longer required. At best, it would disrupt your workflow and productivity. At worst, you're facing compliance or public relations issues.

So, how can you destroy records "defensibly"? First, make sure you have information lifecycle processes and policies in place. The decision to destroy should not be subjective, ad hoc, or left up to an individual employee—it should be guided by already-established policies and records retention schedules.

A defensible destruction program might consist of one, or many, of the following processes:

- **Disposal Class Capture.** The ability to record a disposal class or disposition authority number as part of each disposal rule. The class or authority number is the official authorization to destroy and proves under what authority the action was taken.
- **Metadata Stubs.** The ability to retain a metadata stub of the content following its disposal. This shows that the organization did at one stage have the information, but that it was ultimately destroyed as part of an approval process with the metadata of the content retained.
- **Manual Review.** The ability to involve business decision-makers in disposal outcomes. This means asking the various business areas to determine if something is no longer required and take ownership in the disposal process.
- **Destruction Certificates.** The ability to produce a disposal certificate following the process. This shows everything that was destroyed as part of a particular batch and can be saved back into the system as a record.
- **Audit Trail Retention.** Retaining the audit trail of what you have done is almost as important as the information itself. This is especially true when it comes to defensible disposal process and the ability to show exactly what happened, who did it and when. The audit trails of these actions should be retained as part of a defensible process.

With these measures in place, the destruction of records can be justified and defended should the need arise. If, for some reason, the organization wants to destroy records outside of the standardized process, they must do their due diligence to ensure that the records can be safely purged without operational or legal ramifications *(Translation: Make sure there are no active or pending lawsuits, investigations, or audits).* Even then, destruction should always follow established, written procedures.

Finally, it's important to realize that the decision to destroy is separate from the act of destruction itself. Depending on the type of record, or the content contained within, "delete" may not be enough. To ensure that the content can't be recovered later, you may need to use special software or processes to destroy the records.

# How Records Management Works in M365

*(An excerpt from our Microsoft 365 Records Management Roadmap e-book)*

It’s important to identify your organization's information management requirements and determine whether they can be fulfilled by native capabilities. To get you started, we’ve outlined some common compliance requirements that highly regulated companies might face, along with the out-of-the-box features available in Microsoft 365.

| Common Records Compliance Requirements | Microsoft 365 Records Management |
|---|---|
| **AUDIT TRAIL RETENTION** | |
| To provide record integrity, audit trails must be retained for as long as the record is held for. | Audit trails are retained in Microsoft 365 for either 90 days (E3) or 1 year (E5) depending on licensing |
| **DEFENSIBLE DESTRUCTION** | |
| Records should be destroyed with processes in place that ensures their destruction can be defended in the future. This might include:<br>• Retaining metadata stubs<br>• Business owner approval prior to destruction<br>• Audit trail retention<br>• Retention authority class capture | Microsoft 365 offers some defensible disposal with E5 licenses. This is available in the Disposition tab – Disposed items list, however details of the items disposed will only be kept for 7 years from the time the item was disposed |

| Common Records Compliance Requirements | Microsoft 365 Records Management |
|---|---|
| **DISPOSAL AUTHORITY TRIGGERS** | |
| Organizations need to be able to configure retention and disposal rules in accordance with disposal authorities. This may include a variety of disposal triggers or criteria. | Microsoft 365 disposal triggers are limited to:<br>• Date Created<br>• Date Modified<br>• Date Label Applied<br>Event-based retention can be configured, however it requires significant configuration and oversight. |
| **DISPOSAL PROCESSES** | |
| • Disposal process kick off<br>• Multi-stage business owner approval prior to disposal<br>• Access to approval to dispose of records | • In Microsoft 365, disposal processes are run automatically each week. They cannot be established on demand.<br>• Business owner approval is limited to a single user.<br>• All reviewers will be able to see all records to be approved for disposal. |
| **UNAUTHORIZED DISPOSAL** | |
| Inadvertent or deliberate disposal of records | Adjusting user permissions to remove disposal privileges causes additional administrative burden. M365 retention labels can be used to enforce retention, however because only one outcome can be assigned to a label, if the labels are used for this, they cannot be used for any other record lifecycle outcome. |

▸ Find out how **AvePoint Cloud Records** enhances Microsoft 365's native functionality.

# IV. Backup Content to Remove Inactive Accounts

A significant portion of your storage may be occupied by prior employees' data. When people leave the organization, you're faced with a question of what to do with their content, including emails, files, recordings, and more. There are a few options to consider:

- **Keep the account alive.** When storage was cheap and plentiful, it was common to keep the accounts of departed employees alive for fear of deleting something important. This option makes much less sense now.
- **Share the departed employee's account.** A common practice is to share the departed employee's inbox or OneDrive with their manager or colleagues. While perhaps a practical way of sharing data, it raises significant privacy concerns.
- **Turn on litigation holds or retention in Microsoft 365.** Yes, you'll have exerted some control over the employee's content, but their data footprint will still count toward your storage capacity.
- **Consider any compliance requirements.** For departing employees, these compliance requirements might include assessing whether Mailboxes or OneDrives still hold corporate information that needs to be moved to a more appropriate location. Some organizations (such as the U.S. Public Sector) may have specific requirements around email capture like the Capstone requirements that dictate all emails to and from employees of a certain level must be retained as permanent records.

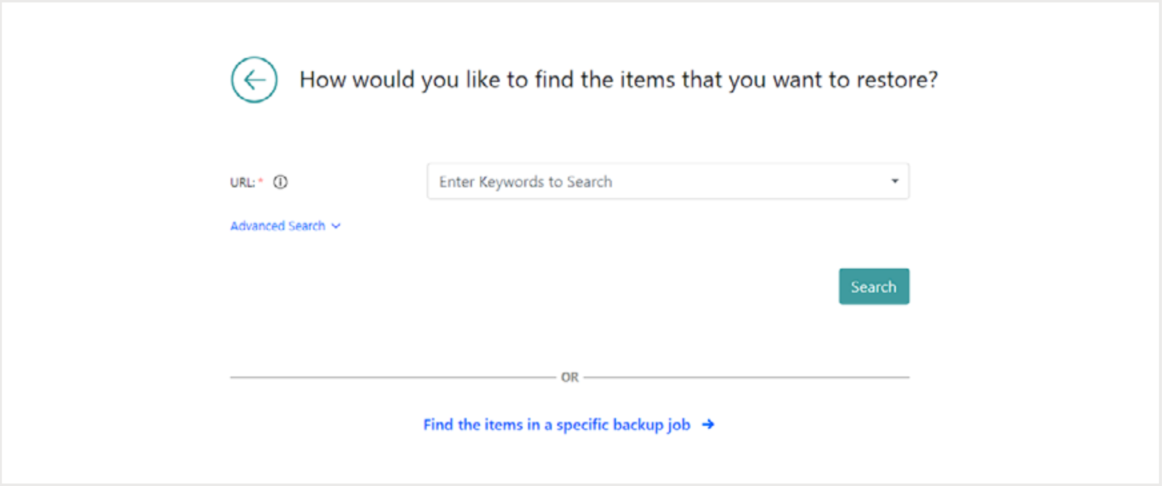

How would you like to find the items that you want to restore?

So, what's the solution? How can you respond to an employee's departure in a way that reduces your storage, while also holding onto their data—just in case? ***Back it up so you can restore it if you need it.***

AvePoint Cloud Backup

A cloud backup solution with granular restore (such as AvePoint's Cloud Backup) would enable you to restore the former employee's mailbox or OneDrive should the need arise. You've likely experienced the frustrating situation where you need access to a file that a former colleague worked on and you can't get your hands on it because it's been deleted or removed. With Cloud Backup, an admin can search for something as specific as a document produced between certain dates or a certain email subject line. Better yet, once they locate it, they can restore it and send it directly to a current employee's OneDrive or email. Problem solved.

# Conclusion

As organizations move past the nascent stage of their cloud journey, collaboration has become king. Tools such as Microsoft Teams, OneDrive, and SharePoint have become a cornerstone of the modern workplace, empowering users to co-author, co-present, and communicate with ease. What users don't see, however, is the heavy digital footprint that these solutions leave behind.

Just as cloud solutions must be continuously evaluated and optimized, so should cloud storage. The explosive growth of collaboration data is going to continue for years to come, which means that while short term solutions may help you avoid next quarter's invoice for data overages, forward-looking organizations must implement long-term tactics to manage costs. We've covered critical strategies including:

- ☑ **Take out the trash that's "ROT"-ting in your environment**
- ☑ **Gain control of your workspaces and content with lifecycle management**
- ☑ **Make sure your deletion and destruction process are defensible**
- ☑ **Free up space by backing up inactive accounts**

By implementing robust information management protocols, organizations can move from defense to offense—isn't that what we all want? To learn more about how to optimize your storage costs, protect against surprise overages, and safeguard your organization against legal and reputational risk, visit www.AvePoint.com.

AvePoint

**AvePoint US Headquarters, R&D Center**

525 Washington Blvd, Suite 1400 | Jersey City, NJ 07310
+1.201.793.1111 | sales@avepoint.com