Integrity360

your security in mind

Q&A follow-up

Zero to Purview hero: unlocking the full potential of Microsoft's Data Governance & Protection Suite

July 9th 2025

Topics covered

1 Compliance manager & standards

Q: Which compliance standards does Compliance Manager support?
A: Microsoft Compliance Manager comes with built-in assessment frameworks for standards including GDPR, HIPAA/HITRUST, NIST SP 800-53/CSF, PCI DSS, SOC 2 and more. You can customise the dashboard, track control maturity, assign remediation tasks, and export evidence for audits.

2 Selective blocking policies

Q: Can you implement blocking policies selectively on users/groups?
A: Yes—Microsoft 365 allows for highly targeted enforcement using security groups, Microsoft 365 groups, or dynamic EntraID (formely Azure AD) groups. This enables you to apply blocking policies, DLP rules, or compliance configurations only to relevant departments, roles, or user types—without affecting the entire organisation.

To improve accuracy and scalability, many organisations integrate their HR system with EntraID. This ensures job titles, departments, and employment status are consistently reflected across your identity infrastructure. You can then build dynamic groups based on these enriched attributes (e.g., "All employees in Sales UK", "Contractors", "Finance Team") and assign policies accordingly.

Example use cases include:

- Blocking file uploads to third-party apps for contractors
- Applying stricter DLP rules to HR and finance departments
- Tailoring Communication Compliance to executive groups

Ensure your HR data is clean and synced to EntraID via SCIM, Azure AD Connect, or a custom connector. This reduces manual group management and makes policy enforcement more precise and dynamic.

3 EntraID clean-up for peer groups

Q: If using EntraID for peer groups, should you sanitise job titles first?
A: It's best practice to maintain clean and accurate EntraID data (job titles, department, location) before deriving dynamic groups. Poor metadata can lead to mis-grouping/ inaccurate scopes being applied.

4 Limiting AI prompts

Q: Can you apply policies to limit AI prompts?
A: Yes with a layered, well-configured approach. The first step is enabling Microsoft Audit and ensuring Purview is properly integrated. Once this is in place, Purview can log user and admin interactions with AI tools like Copilot, capturing the actual prompts being sent. This visibility enables you to detect potentially malicious or risky behaviour, such as the unauthorised sharing of sensitive data.

From there, you can build DLP policies to block or monitor AI-related activities. These policies can be applied not just to Copilot, but potentially extended to third-party tools like ChatGPT, depending on your setup. While full GPT-style prompt control is emerging, you can prevent PII from being shared with external AI tools using Comms Compliance and DLP. To make this effective, Sam recommended:

- Enabling **DLP for endpoints**
- Using the **Purview browser extension**
- Reducing the number of **non-sanctioned AI apps** in the environment to minimise the attack surface

5 AI apps on mobile devices

Q: What about users typing into AI apps on mobile devices—can we extend those same controls?

A: Full enforcement on mobile requires the device to be managed. For personal or unmanaged devices, Microsoft offers two key options:

Microsoft Defender for Cloud Apps (formerly MCAS) – This acts as a reverse proxy for unsanctioned apps like ChatGPT. When integrated with Single Sign-On, it can intercept traffic, enforce session controls, and block high-risk actions—even on personal devices.

Mobile Application Management (MAM) policies – These provide granular app-level controls (e.g. disabling clipboard, restricting file sharing) without requiring

full device management. While these don't provide traditional DLP enforcement, they help contain sensitive data in mobile scenarios.

"*There isn't a single silver bullet—it's about defence in depth, using multiple Microsoft tools together to protect your data wherever users interact with AI.*" Richard Ford

6 Communications compliance for profanity/sexual harassment

Q: Have you implemented communications compliance for profanity/sexual harassment for customers?
A: Absolutely - particularly within the education sector, where preventing inappropriate language or harassment is a high priority. Sam Lambert explained that Communication Compliance can be an effective control, but it's essential to enable it proactively. Microsoft's privacy-by-design approach means that unless Communication Compliance is turned on in advance, it won't retain the necessary audit data for investigations.

One of our enterprise clients required automated monitoring of Teams and email for harassment, profanity, and sexual content. We configured custom sensitive term dictionaries and policies in Microsoft 365 Comms Compliance. The implementation successfully flagged and blocked over 350 incidents in a month, with real-time alerts to HR teams.

7 Excluding automated emails from mandatory classification

Q: Can you exclude automated emails from mandatory classification?
A: Yes. You can create exclusion rules based on sender patterns (e.g., service@myapp.com, no-reply@domain.com). Use Exchange transport rules or DLP rule exceptions to bypass mandatory classification for specific senders or templates, ensuring automation systems aren't disrupted.

8 Labelling Best Practices for email & sharepoint

Q: Best practice to follow when implementing labels. We have recently deployed Info Protection for PII data in email, but we are expanding, and SharePoint is next on the DLP migration path, but we would like to make use of labels as demonstrated in your presentation.
A: Yes, there are clear best practices to follow—especially when expanding labelling from email (where you've already applied Information Protection for PII) to SharePoint and other Microsoft 365 workloads.

Start with "Know Your Data"

Begin by identifying where sensitive information resides across your Microsoft 365 estate—Exchange, Teams, SharePoint, OneDrive, etc. Understanding your data landscape is crucial before assigning any labels.

Use both sensitivity and retention labels

Don't just focus on sensitivity. Retention labels help identify stale or unmodified data (e.g. files untouched for 5+ years) and can trigger workflows for archival or deletion—reducing risk and cleaning up your environment.

Customise information types

Go beyond the defaults. Create custom sensitive information types that reflect your organisation's specific data classification needs.

Test in simulation mode first

Before full deployment, run your label policies in simulation mode to observe impacts without disrupting users. This helps reduce friction and uncover any conflicts early.

Consider policy interactions across platforms

When expanding to SharePoint, think about how your labels will interact with those already applied to email. For instance, what happens when a SharePoint document with one label is sent via Outlook, where another policy applies? Ensure these scenarios are reconciled in advance to avoid conflicts.

In short: build gradually, test thoughtfully, and plan holistically—especially when moving beyond email into broader data governance in SharePoint and beyond.