Balancing AI innovation and data security risk within public sector

As businesses increasingly adopt AI tools, the benefits in terms of productivity and efficiency are hard to ignore. But with great power comes great responsibility.

Balancing AI innovation and data security risk within public sector

Last month, the Office of the Victorian Information Commissioner (OVIC) **released a report** into the investigation into the use of ChatGPT by Child Protection workers. The report has sparked a flurry of conversation and healthy debate about the use of generative AI (GenAI) within public sector, including conversation on how to protect sensitive data (such as personal information about individuals), mitigate privacy risk, ethically govern AI use, and manage bias (to name a few).

Some departments within the Victorian government have chosen to place a blanket ban on all GenAI – including Microsoft's Copilot products - while it addresses the data security risk and seeks further guidance on how to mitigate privacy risks.

The benefits of GenAI within public sector are undeniable. From enhancing the way services are delivered to the public, to automating processes, improving collaboration, improving individual productivity, and informing policy. The use cases for GenAI to transform the way the public sector support our community are endless. A **Copilot trial** completed by the Australian government found that:

"There are clear benefits to the adoption of generative AI but also challenges with adoption and concerns that need to be monitored."

- *77% were optimistic about Microsoft 365 Copilot at the end of the trial.*
- *1 in 3 employees used Copilot daily.*
- *86% of trial participants wished to continue to use Copilot.*
- *Senior Executive Service (SES) staff (93%) and Corporate (81%) roles had the highest positive sentiment towards Copilot.*

Respondents estimated that Copilot saved up to an hour a day…Post-use survey respondents perceived Copilot contributed the highest time savings in activities related to information summarisation preparing first drafts and information searches."

Summary report – evaluation findings | digital.gov.au

But – it's still early days. The pace at which AI-powered products are being developed is almost unprecedented. It's critical that all organisations – across both the public and private sector – understand not only the risks but also the value AI could deliver. As OVIC states in their AI privacy guidance:

"While the use of an AI system to process personal information can deliver significant benefits, VPS organisations should consider whether the deployment of such a system is necessary to address an identified problem, and whether it is the best solution to that problem – AI systems should not necessarily be deployed simply because they are available."

Artificial Intelligence – Understanding Privacy Obligations – Office of the Victorian Information Commissioner

In this document, we share our thoughts on how public sector organisations can balance assessing the risk while also pushing forward innovative ways to better serve the community using the power of GenAI.

While there are many AI tools 'on the market', we'll focus our thoughts on Microsoft's suite of AI products (Microsoft 365 Copilot, Copilot Studio, Azure OpenAI etc.) and talk through a real scenario for how these tools could be used to support the process around policy creation and approval.

Before we jump in, let's quickly get you up-to-speed on what concerns were raised in the OVIC report:

OVIC's report brings several important issues to light, and it's essential for businesses to be aware of these as they integrate AI into their operations:

Data Privacy and Compliance

AI tools can handle a lot of sensitive personal information, which raises questions about how companies manage that data.

Bias and fairness

AI systems can unintentionally reflect biases in their training data. This means businesses need to be on guard to ensure fairness and prevent discriminatory outcomes.

Security vulnerabilities

The use of shadow IT such as ChatGPT can expose new data security risks such as inputting of sensitive data into external systems. There's also concern that implementing AI tools within organisations will expose existing security risks, particularly around unauthorised access to sensitive data. Organisations need to take steps to protect their information.

Ethical governance

Companies must develop ethical guidelines to govern AI usage. This includes being transparent about how AI makes decisions and ensuring accountability.

Education

It's not just about using the tools; employees need to understand how to use them responsibly, too. It's critical employees understand what can AI do and what can it not do and what are the strengths and limitations of each AI solution.

The report raises an important question: How can organisations embrace AI-powered productivity while governing data appropriately?

The challenge for government organisations is to innovate responsibly—testing AI's potential while managing risks effectively. To achieve this, organisations need to adopt a structured approach that puts the right guardrails in place to foster innovation and ensure GenAI, like Microsoft's Copilot products, is implemented with precision and purpose. This includes going beyond technical controls; providing the right education and support to employees is of equal importance. Generative AI isn't confined to the workplace. It is fast becoming a staple in our virtual world. Outside of the workplace, we can interact with AI on everything from social media to more sophisticated advertisements, and search engines; the use cases of generative AI are growing larger and larger by the day. With this backdrop, it's fair to say that AI is here to stay and that testing its capabilities is a must-have on any technology strategy.

A methodology for assessing risk, testing and implementing AI

AI implementation should never be rushed, especially within highly regulated sectors such as government. The value of AI must be tested in a systematic way, ensuring that both its potential and its risks are understood.

While a whole-of-vic-gov AI strategy and mandatory risk assessment is yet to be determined, there are minimum expectations that OVIC has released to assist government departments, this includes assessing the maturity of their existing information security program prior to signing up for the integration of Copilot.

This includes:

- identifying existing information holdings and systems that may be impacted by the introduction of Copilot, including due consideration of the security value of these;
- considering how any newly generated information by Copilot will be assessed, valued and securely managed, including applying appropriate protective marking for newly generated information assets;
- conducting an updated information security risk assessment considering the integration of Copilot, including due consideration of adjusted risk profiles and development of treatment plans to address relevant Copilot features;
- implementing any updated treatment plans by rolling out any new or changed controls. Specific consideration should be given to a review of the organisation's existing logical access controls, informed by an audit of all current access controls to protect, detect, and handle against unauthorised access to information holdings and systems;
- implementing a formal process for the ongoing monitoring and review of risks and controls, especially critical given the dynamic development and release of enhancements and new features of Copilot.
- explicitly disable data sharing with Microsoft in platforms such as Power Platform and Dynamics 365;
- undertake a privacy impact assessment to understand the ways in which Copilot will be utilised, the risks it presents to the privacy of individuals, and how to mitigate those risks;

The full list of minimum requirements can be found here: Public Statement: Use of Microsoft 365 Copilot in the Victorian Public Sector

Get help with your Copilot journey

Our Copilot implementation framework helps you meet the minimum expectations while putting in place a plan to rapidly develop, test, and safely implement GenAI across your organisation.

Technical readiness
Ensuring the Technical readiness & deployment of Copilot to end users incl. of licensing assignment.

Art of the possible
Build an understanding of core Copilot functionality and how it applies to your organisation.

Copilot Pilot
Implement 50 - 300 Copilot licenses, focus on direct enablement and build knowledge and capability in line with business scenarios relevant for your organisation.

Copilot Business Case
Establish a business case to gain executive buy in for recommended next steps of Copilot, through an understanding of value opportunity at your organisation.

Scaled implementation
Implementation of > 300 Copilot licenses, with a focus on driving knowledge and capability in line with the high value scenarios identified in EAP and expanding scenarios relevant for your organisation.

InfoSec Remediation

Information Security Maturity

Copilot extensibility and integration

ENGAGE SQUARED

Let’s look at our methodology in a little more detail

1. Start with technical readiness

Next, it’s crucial to assess the organisation’s technical readiness for AI deployment. This involves evaluating the existing IT infrastructure, ensuring it’s capable of supporting new AI tools. Microsoft’s Admin Center offers several built-in capabilities that help in this process.

A technical readiness assessment should cover:

- Compatibility checks: Ensure that current systems are compatible with Microsoft Copilot and other AI tools.
- Resource planning: Assess the need for additional storage, processing power, or cloud infrastructure.
- Compliance alignment: Confirm that the infrastructure meets the regulatory requirements relevant to public sector operations.

This phase also includes identifying potential gaps in the technical environment that could hinder AI performance and addressing these before proceeding to the next steps.

2. Value discovery

Before any AI deployment, it’s essential to begin with a value and value workshop. This workshop brings together key stakeholders—IT, security, and departmental leaders—to identify where AI can drive the most impact and start a discussion around security, governance and data access. A short workshop like this can also be a perfect opportunity to showcase how AI could transform business operations, which roles would benefit from having access to tools like Microsoft 365 Copilot, and what the return on investment could be.

Workshop Framework

Inspire	Assess	Design
• Overview of Microsoft 365 Copilot’s value to inspire and engage stakeholders. • Address any security, governance, and data access questions • Explore strategies to accelerate the adoption and change management process using Microsoft Viva.	• Showcase the intelligence added to employee experiences with Microsoft 365 • Explore how you can boost creativity, unlock productivity, and spend more time on work.	• Develop a plan to implement recommendations based on prioritised scenarios • Develop a business case, including ROI considerations, to justify the investment in Copilot. • Define next steps and timeline to develop and implement Microsoft 365 Copilot • Deliver executive summary of key findings, recommendations, and next steps for implementing Copilot.

Expected Outcomes

Prioritised business scenarios that can be addressed by deploying Microsoft 365 Copilot

Recommended activities to help prepare your people and technical environment

A roadmap outlining potential work streams and dependencies with **clear next steps**

3. Information security remediation

We recommend that this step happens in tandem with technical readiness and value discovery activities. It is critical in ensuring the seamless integration of AI tools like Copilot into existing workflows.

Key steps include:

- Upgrading infrastructure: If necessary, upgrade servers, cloud capacity, and networks to handle increased data processing.
- Configuring data governance policies: Leverage **SharePoint Advanced Management** and **Microsoft Purview** to establish and enforce data governance across the organisation. This includes reducing accidental oversharing, cleaning up inactive sites, classifying data, applying data loss prevention policies, and setting access controls.
- Enhancing cybersecurity: Strengthen the organisation's cybersecurity posture by updating firewall settings, enabling multi-factor authentication (MFA), and applying role-based access controls (RBAC) within Microsoft's Admin Center.

By optimising the technical environment, government organisations can ensure AI tools perform efficiently and securely without creating bottlenecks in existing systems.

4. Copilot Pilot

As AI tools like Microsoft 365 Copilot or Copilot Studio are deployed in a controlled environment, it's important to observe and capture functional use cases. Monitoring the effectiveness of AI in the context of specific workflows within your organisation will provide insights into its value and help determine a business case for broader implementation.

During this phase, government organisations should focus on:

- User feedback: Gather input from employees who are using AI tools to understand their experience and identify areas for improvement.
- Security incidents: Monitor for any potential security breaches or data leaks, ensuring that the AI deployment remains compliant with security policies.
- Capturing AI use cases: Establishing a mechanism to capture, triage, assess, and prioritise use cases
- Performance metrics: Track key performance indicators (KPIs) tied to each AI use case, such as time savings, accuracy improvements, or cost reductions.

This phase allows the organisation to iterate on its priority use cases, refining where AI is most beneficial and ensuring that risks are mitigated before expanding its use.

5. Copilot business case

Having a clear business case that articulates use cases (specific to your organisation) and provides data on AI usage (e.g. how, when, and where it's used) mixed with qualitative feedback is critical to understand the direct and indirect benefits and costs. A clear business case should also include the proposed roadmap for scaling AI use within an organisation and recommendations for how to govern the technology (to ensure sustained ROI).

Data Review

Provides insights into Copilot usage analytics and estimated cost savings.

Desktop research

Empirical research to contextualise and compare impacts of Copilot.

User interviews

1:1 interviews with Copilot sponsors to generate insight for future expansion.

Focus groups

Group discussions with departmental representatives to gather Copilot use-cases and feedback

Copilot Business Case

Synthesising insights gathered through our research to develop a business case to invest in Microsoft Copilot

Key Opportunities

OPPORTUNITY 1:

Employee productivity increase and time savings

Based on the current usage levels of Copilot, analysis indicates users are saving

3.6 hrs per month

Outcomes:

- Less time doing low-value, mundane tasks e.g. finding files or making notes.
- More time applying their skills and focus on higher value activities leading to high quality outputs delivered with more pace.
- Overall productivity and efficiency uplift.

Challenges it will help address:

OPPORTUNITY 2:

Improved wellbeing and employee experience

90% had more time to focus on important work

8.94/10 satisfaction score among the Pilot users

Outcomes:

- Human hours back boosting wellbeing.
- Improved working experience with less time spent on mundane tasks.
- Improved access to relevant information and people - less duplication of effort and more knowledge sharing.

Challenges it will help address:

OPPORTUNITY 3:

Achieve Organisational goal of Cultural Simplification

88% can complete tasks faster and more easily

94% spent less time searching for information

Outcomes:

- Ability to leverage Copilot's functionality to simplify individual work tasks.
- Copilot powered search helps users to overcome complicated and scattered storage to find items quickly.

Challenges it will help address:

Top challenges

High value work being blocked by mundane tasks.

Employees are time poor.

Employees are finding it difficult to locate what they need.

Example output of opportunities from AI business case development, completed by Engage Squared

6. Scaled implementation

Once AI tools have been thoroughly tested and validated in discrete groups, the final step is to scale the implementation. At this stage, your organisation should have a solid plan that includes key audiences, prioritised use cases, and a comprehensive change program to support employees through the change.

Scaling should involve:

Change strategy: A comprehensive approach that prioritises user-centric initiatives, ensuring successful adoption and utilisation of Copilot through tailored change management, stakeholder engagement, and robust training programs.

Communications: Development of communications to drive awareness & knowledge among new Copilot users using communications channels including Enterprise Social Network

Training & toolkits: Development and delivery of toolkits for leaders, as well as use case and persona-based training for users to support maximum and sustained adoption

Success measurement: Adoption measurement and feedback, incl. adoption dashboard reviews, user feedback and recommendations for actions

Continuous data security monitoring: Use Microsoft Purview to maintain ongoing oversight of data governance, ensuring that all AI interactions remain compliant with regulations. Feedback loops

Adjusting controls as needed: As AI tools are scaled, it may be necessary to adjust access controls, DLP policies, and data classification rules to accommodate the growing user base.

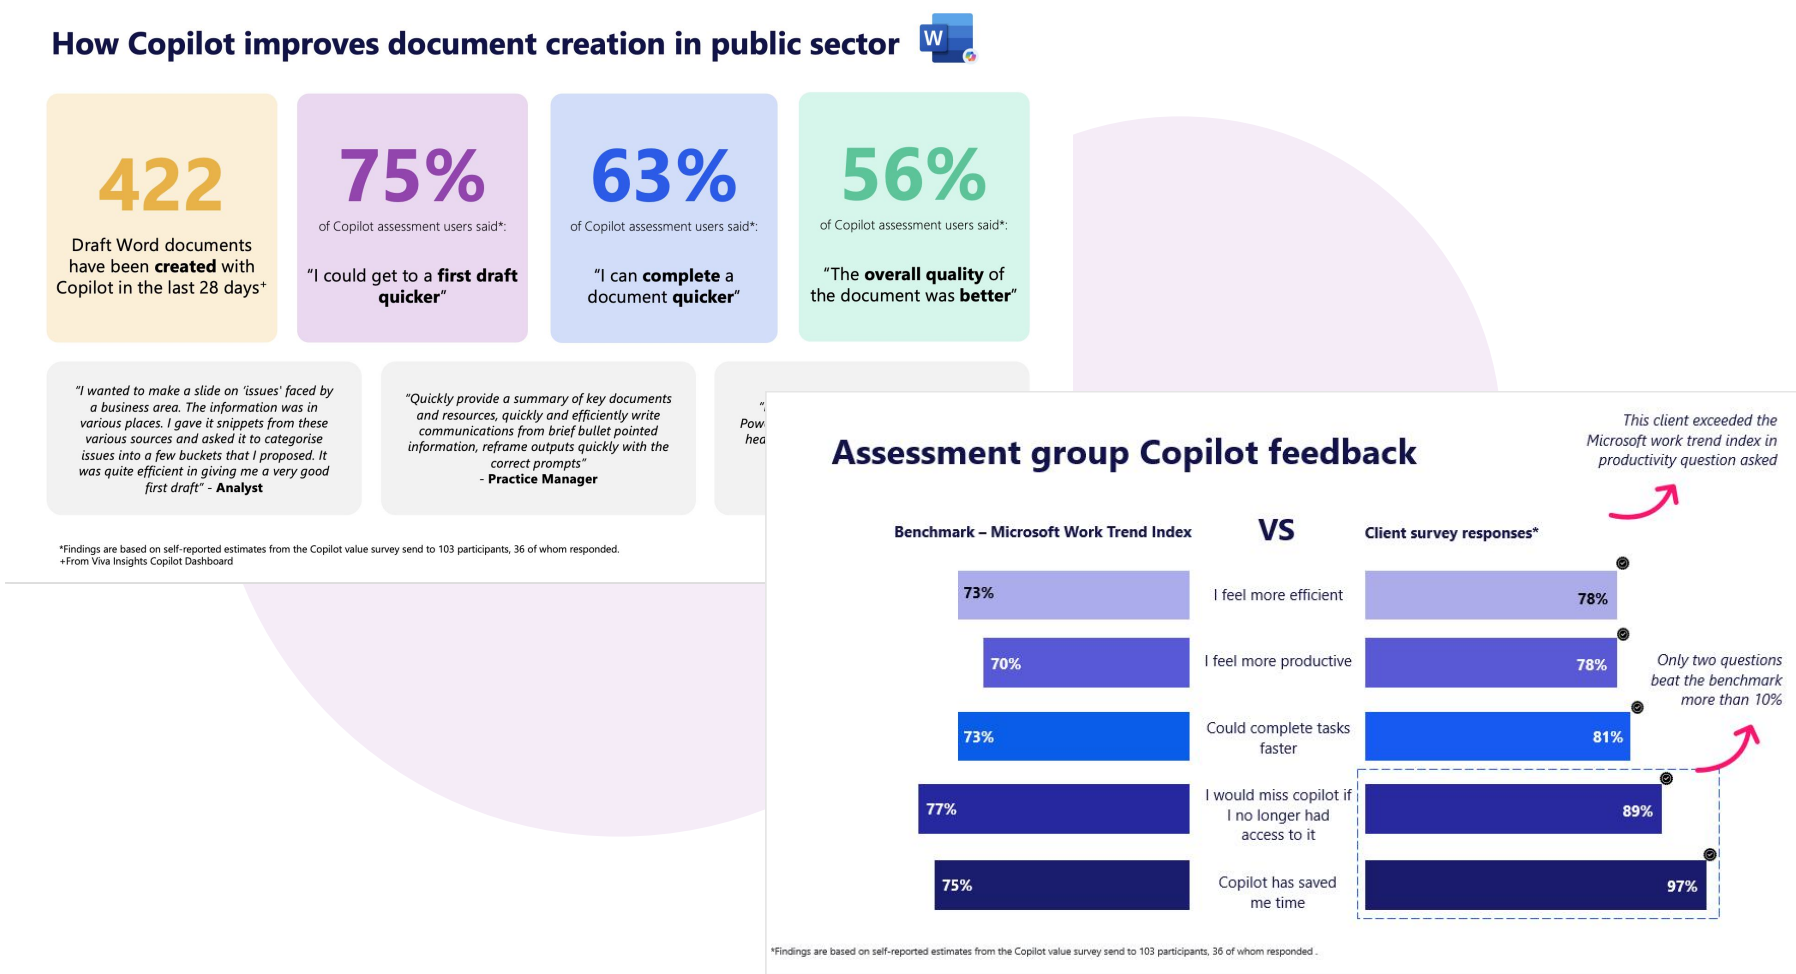

Engage Squared work example from Copilot Pilot

By following a structured, methodical approach—beginning with technical readiness and value discovery—government organisations can harness the power of AI while safeguarding the sensitive information they are entrusted to protect.

Let’s look at a real-life use case for AI within Public Sector

Picture this – you are a Senior Policy Advisor within a state government department. In your role, you’re responsible for tasks like:

- Developing and tracking policy briefs: Create and manage policy briefs from initiation to approval, with real-time tracking and status updates.
- Coordinating and consulting with key stakeholders – within and across agencies: Share and receive feedback on policy briefs efficiently, ensuring all stakeholders are kept informed.
- Analysing policy impact: Use the platform to aggregate data and generate reports on the potential impacts of proposed policies.

You manage the creation of your briefs through **Brief Connect** which has built-in information protection policies for cabinet-in-confidence documents. While Brief Connect is a big step up from how your Department used to work, you’re eager to explore how GenAI could speed up tasks even more and have heard that it could even **pre-populate briefing templates with reliable source data**, cutting down your creation of the initial brief from hours to just minutes.

Testing Gen AI within discrete solutions that already meet information security, compliance, and records requirements is a powerful way to measure ROI and value within a controlled environment.

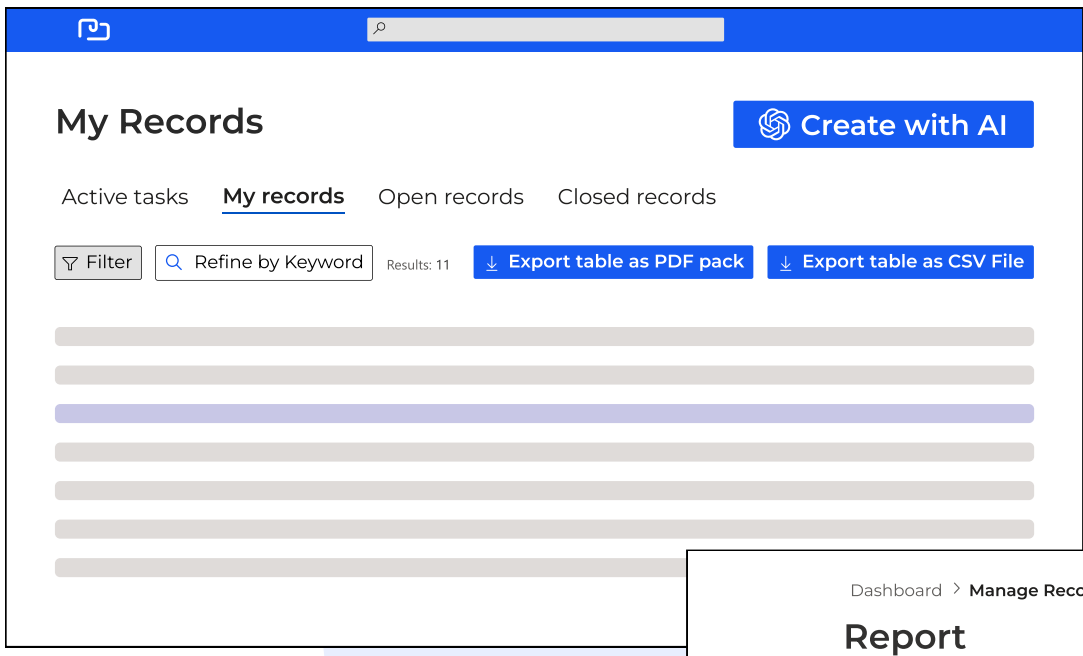

OpenAI

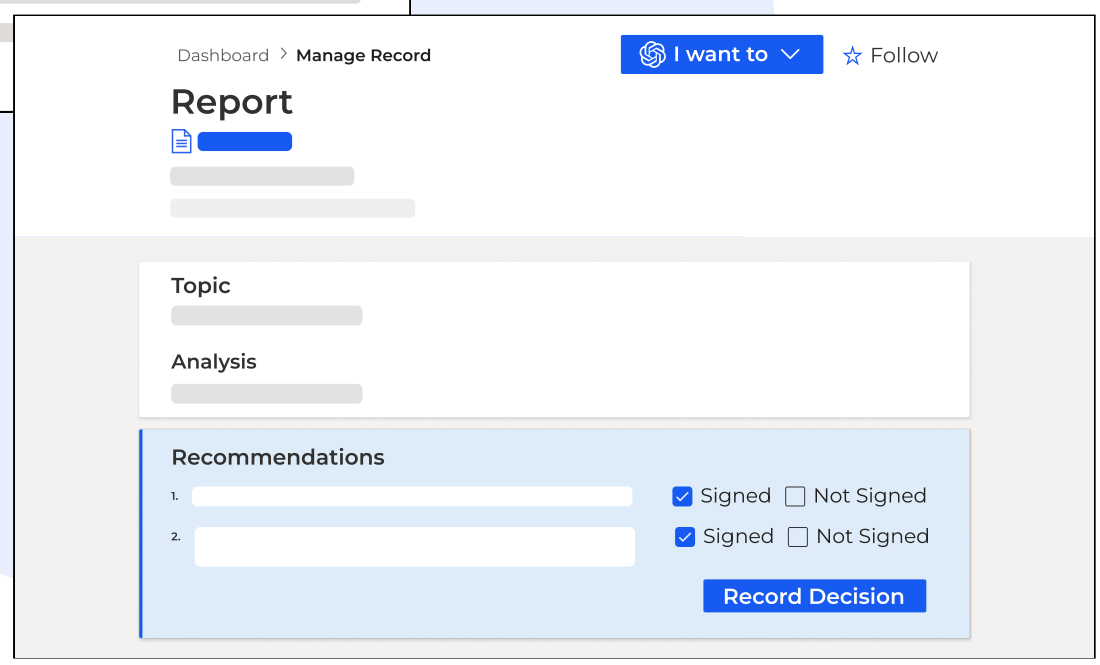

ENGAGE SQUARED

Day in the life of a Senior Advisor

Name: Alex

Role: Senior Advisor at a State Government Department

Context: Alex is a Senior Advisor responsible for reviewing all briefs before they reach the Deputy Secretary.

Today, Alex is managing a brief regarding a proposed environmental regulation before an upcoming meeting.

8:00am — Stay one step ahead

Alex begins the day on her phone, opening Microsoft Teams. To quickly prioritise tasks, Alex uses Microsoft 365 Copilot to summarise any urgent emails or Teams messages related to the environmental regulation brief received overnight.

Copilot in Teams
Prompt*: Summarise any urgent emails or Teams messages related to the environmental regulation brief received overnight.*

9:30am — Review dashboard of upcoming briefs

Now at her desk, Alex logs onto Brief Connect within Teams. Alex uses her personalised dashboard displaying status of open briefs, briefs in progress, pending tasks assigned, and highlighted deadlines for submission.

Brief Connect dashboard
Immediate visibility into the brief's status allows Alex to schedule dedicated time for thorough review and collaboration.

10:30am — Collaborating on draft briefs

The lead author has assigned Alex as a collaborator in Brief Connect with the task to peer review the draft brief and provide feedback by 10:00 AM. Alex accesses the draft directly from a notification, reviewing the new data incorporated overnight with a focus on policy risks.

Copilot in Word
To ensure the brief adheres to departmental writing standards, Alex brings up Copilot in Word while working on the document. By /referencing The Department Written Style Guide, Alex asks Copilot to suggest changes to the document content based on the guidelines. Copilot analyses the brief and provides recommendations for improvements in language, tone, and formatting to align with the department's standards.

11:00am — Managing workload efficiently

Automated alerts from Brief Connect notify Alex of tasks requiring attention, including other briefs nearing deadlines, updates on briefs revised based on feedback, and reminders for upcoming meetings related to the brief.

Lunch

2:00pm — Generating a status report for the Deputy Secretary

The Deputy Secretary requests a status update on the brief. Alex accesses Power BI dashboards to view real-time data on the brief's progress and exports relevant visuals to PowerPoint, creating a concise and informative report.

Microsoft Copilot in PowerPoint
Prompt: Create a summary slide highlighting the current status of the environmental regulation brief, including key milestones.

3:00pm — Final preparations for approval

Advancing the brief to the approval stage in Brief Connect, Alex confirms all sections are complete and compliant with departmental standards. Ensuring that version control reflects the most recent edits avoids discrepancies. Alex sets automated reminders in Brief Connect for the Deputy Secretary to review the brief.

4:00pm — Planning for potential outcomes

Microsoft 365 Copilot
Prompt: Generate potential questions that might be asked about the proposed environmental regulation and draft suggested responses

5:00pm — End of day

Alex uses Microsoft 365 Copilot to generate potential questions that might be asked about the environmental brief to prepare for an upcoming meeting.

Alex prepares to log off, confident that the day's objectives have been met and that the environmental regulation brief is on track. The combination of Brief Connect and Microsoft 365 tools has streamlined the workflow, making the complex task of managing briefs more manageable.

Copilot in Teams
Prompt: Based on today's emails and Teams messages, create a to-do list of actions I need to take tomorrow.

Need help assessing your AI data security risk?

Engage Squared offer Copilot health checks (aligned to OVIC guidance) to assist public sector organisations to manage and secure the data in their M365 environment.

Copilot inherits your existing Microsoft 365 security, privacy, identity, and compliance policies. Your data is logically isolated and protected within your M365 tenant, and always within your control.

In this engagement, we will work with your technical and security teams to **Interim guidance on government use of public generative AI tools - November 2023 | aga** undertake an assessment of internal security controls, focused on identifying gaps and providing recommendations for an uplift to your information security before it's exposed to AI.

As an outcome of this engagement, we'll produce a report to help you:

- Understand how compliant your information landscape is (aligned to your state's legislation),
- Understand how secure your information is,
- Understand the level of risk of oversharing, and
- What specific next steps your Department needs to take to remediate.

The report will detail our evaluation of the risks of oversharing in your modern group-enabled and non-group-enabled SharePoint sites, focusing on the potential for compromised data security, regulatory non-compliance, and operational disruption in a collaborative environment.

The report will highlight areas of concern and suggests direction on measures to enhance data governance, control access, and increase awareness to mitigate these risks.

Finally, the report will help guide which recommendations are implemented by our team as part of this engagement.

Want to learn more about how we can help you?
Contact us today to book in a call with one of our AI experts.

ENGAGE SQUARED

Visit our website for more info