Data#3 | Microsoft

Data integrity: the foundation for secure analytics, cloud and AI.

Transform data into a strategic driver of organisational goals.

Contents

Meet our experts

Peter Heydon

Data and Analytics Solutions Lead

Bruce Hiddle

National Practice Manager, Information and Analytics, Business Aspect

Paul Green

Principal Consultant, Business Aspect

Mark Hindson

Microsoft Security Solutions Architect

Scott Gosling

National Practice Manager, Microsoft

Richard Dornhart

National Practice Manager, Security

Executive summary

The organisations winning with AI aren't the ones that moved fastest. They're the ones that turned data into a strategic advantage.

The problem

Most Australian organisations don't have a data shortage problem; they have a fragmentation problem.

Data is trapped in silos across cloud, SaaS, on-premises and collaboration tools, eroding trust, slowing decisions and quietly increasing breach exposure.

AI amplifies the impact of bad data.

Where poor-quality data once affected the analyst who found it, AI distributes bad outputs across the whole organisation at speed traditional controls weren't designed for.

Accountability sits with the organisation, not the model.

When AI produces a wrong or inappropriate answer, such as surfacing sensitive information to the wrong person, or informing a strategic decision based on flawed inputs, that is an organisational risk, with legal and commercial consequences.

The lever

Governance is the lever that turns fragmented data into a trusted asset.

Done well, it turns data into a strategic enabler, allowing analytics to move faster, cloud migrations to proceed with confidence and AI to deliver real value on solid foundations.

Data Security Posture Management (DSPM) replaces periodic audits with continuous visibility.

DSPM provides real-time insight into who can access what across the entire data estate by identifying excessive access, role creep, outdated permissions and posture drift before they become incidents.

M365 Copilot and other AI tools amplify both good and bad data equally.

Oversharing remediation, sensitivity labelling and access hygiene are pre-conditions for scaled AI deployment, not tasks to address after rollout.

How Data#3 helps

Data#3 applies a five-step framework to control your data.

These steps give organisations a structured, repeatable path from fragmented estate to measurable, ongoing control.

Step 1: Define

Step 2: Identify

Step 3: Classify

Step 4: Assess

Step 5: Implement

Engagements match your maturity and ambition.

Whether you're establishing baseline controls across the highest-risk workloads or operationalising governance as a continuous, enterprise-wide capability, Data#3 meets you where you are.

Each engagement is delivered in three stages: Strategy, Solution and Deployment, to ensure that any technology solution supports business outcomes.

The Microsoft Data Envisioning Workshop is the recommended starting point.

It surfaces sensitive data exposure and oversharing risk, identifies policy and control gaps, and produces a concrete remediation roadmap often before a single Copilot licence is deployed at scale.

Introduction

With AI, organisations can access data in ways they've never been able to before.

Technological advances, however, come with their own risks. Innovations in data access also cause issues such as data leakage and oversharing which diminish the benefits AI brings to the table. People need to trust that when they ask AI tools a question, sensitive data isn't going to be exposed to people who shouldn't have access to it.

The way an organisation's data is stored, governed, controlled, accessed and secured is more important than ever. When data integrity is compromised, employees revert to instinct, spreadsheets or shadow systems, doubling up effort and undermining both productivity and governance. When it is protected, however, it leads to better decisions, more innovation, faster execution, operational efficiencies and competitive advantage.

In a period where almost every organisation is racing to operationalise AI, questions of integrity have become urgent. Data is the input, training set and source of truth for every analytics dashboard, cloud workload and AI agent. If it isn't trusted, neither is the output.

Build a strong foundation

Maintaining data integrity requires a comprehensive approach that factors in governance, security, architecture and operations.

It's not a "technology decision", but more a co-ordinated effort across people, policy and platforms that recognises data as a living asset moving constantly between SaaS apps, cloud workloads, mailboxes and collaboration tools.

On top of this, your organisation needs to think beyond the here and now: how does it manage data across its lifecycle? Bringing data governance in line with best practice is a complex undertaking, which is why Data#3 has developed a proven methodology to help your business transform the way it uses and accesses data. We establish a healthy data foundation so your organisation can move forward confidently with AI, and your leaders are empowered to make faster, more accurate decisions.

Reposition data integrity as a strategic advantage

Strengthening data integrity isn't just a box that needs to be ticked. It can be the strategic enabler that drives value from analytics, cloud and AI. Integrity inspires adoption, adoption unlocks value, and value translates to return on investment (ROI) from better decision-making, innovation and AI.

This eBook explores how to transform your data into a trusted enabler of business outcomes – through secure analytics, cloud and AI. Across the chapters that follow, we look at how to CTRL your data so you can trust your insights, SHIFT your data with confidence as it moves through the cloud, and DELETE uncertainty from AI by giving it only the data it should be allowed to see.

> ***"AI amplifies bad data and, in extreme cases, can be used blindly to make decisions predicated on incorrect data."***
>
> **Peter Heydon, Data and Analytics Solutions Lead, Data#3**

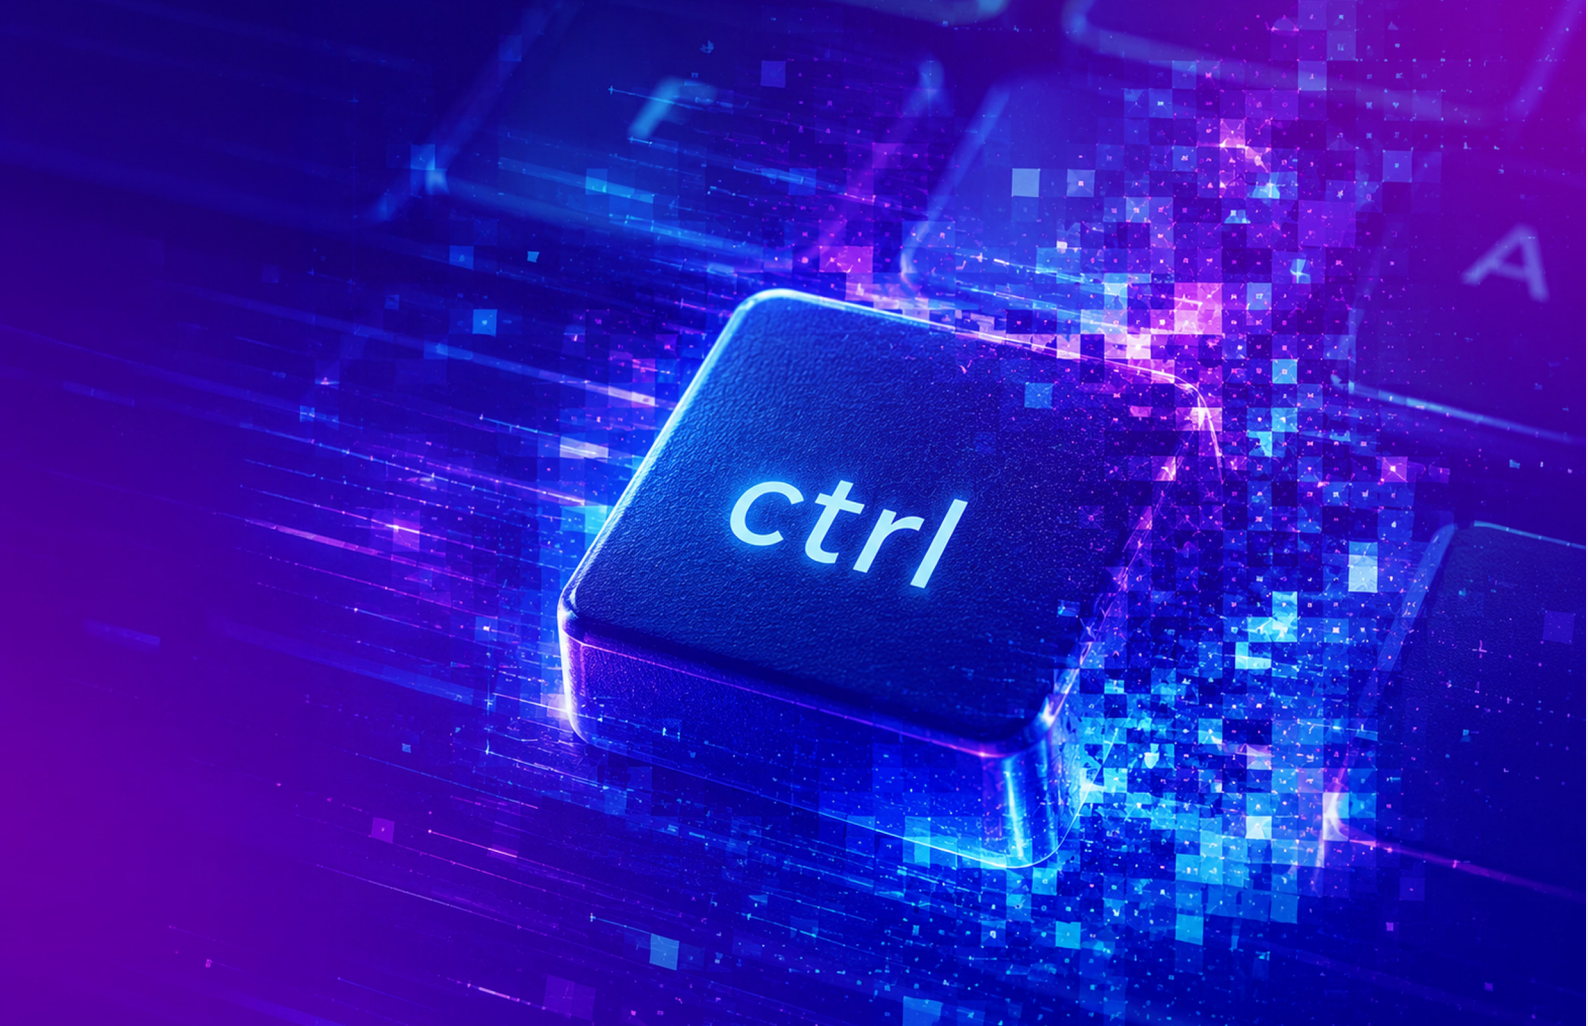

Data fragmentation

CTRL your data to trust your insights.

Data is everywhere, spread across cloud platforms, SaaS applications, on-premises systems and legacy databases that have quietly accumulated over time. Volume keeps growing, but value isn't keeping pace.

This exponential growth has created a situation where data estates are sprawled across locations while also being trapped inside silos.

The result is an opaque, unstructured environment with limited visibility and minimal integration between systems. According to McKinsey, data that should be flowing freely to support decisions is instead trapped inside the systems that produced it, dramatically reducing its usefulness.[1] When data is fragmented, organisations pay a very real price.

Decisions take longer because leaders must compare competing versions of the "truth", reporting becomes inconsistent because teams rely on different datasets, and compliance is harder when no one can confidently say where sensitive information lives.

AI readiness also suffers, because models are only ever as good as the data they are trained on.

There's also a productivity tax. Working with fragmented data is slow and inefficient, with analysts spending more time looking for the right source than producing insights, and innovation stalls while teams debate definitions.

The bigger and more federated an organisation becomes, the more this tax compounds.

1 The data-driven enterprise of 2025

The scale of the challenge

76%

of data leaders say that governance hasn't kept pace with rising AI use.[2]

86%

plan to boost data management investment to support AI growth.[3]

40%

of agentic AI projects will be cancelled by 2028 due to inadequate risk controls, unclear business value or escalating costs (Gartner).[4]

From fragmentation to confidence

Overcoming fragmentation requires a unified view across the data estate, enabling organisations to work across systems without large-scale replacement or enforced standardisation.

That way, organisations get the visibility to know what data exists, where it lives, who can access it and how trustworthy it is. With that foundation in place, analytics can run on safe, secure, properly governed data, and insights become something the business can act on with confidence.

This is the first step towards true data security. Regain control of a sprawling estate so that the dashboards, reports and AI prompts built on top of it are anchored to data your people trust.

> ***"Data resides within diverse places - you usually can't avoid that in enterprise environments. However, creating a unified data platform provides the foundations for rich insights and opportunities for AI."***
>
> **Peter Heydon, Data and Analytics Solutions Lead, Data#3**

2 CDO Insights 2026
3 CDO Insights 2026
4 "Gartner predicts over 40% of agentic AI projects will be cancelled by end of 2027", Gartner 2025

The importance of data governance

Build trust in data across its lifecycle.

If fragmentation is the problem, governance is the solution. It's the antidote to the chaos of most fragmented data estates. Without it, decision-making lacks insight, compliance and security risks increase, access is inconsistent and data quality deteriorates. According to the OECD, poor data governance is a direct contributor to low trust, inefficiency and duplication.[5]

Proper governance ensures data is aligned across business processes, roles and responsibilities. When applied across the data lifecycle, information is accurate, accessible, secure and fit-for-purpose.

In other words, it is the discipline of turning disconnected data sets into one managed, trustworthy asset. As PwC put it: *"Data without governance becomes chaos, and without security it becomes a liability."*[6]

5 Data governance
6 "The future of trusted innovation with data governance and security", PWC 2025

The regulatory imperative

For Australian organisations, data governance has become both a strategic priority and an escalating legal obligation.

The regulatory environment is tightening across multiple fronts, and audit-ready governance programs are increasingly the difference between a managed incident and a material breach event.

Australian regulatory drivers your governance program must address

Privacy Act reforms. Proposed changes to the Privacy Act 1988 will introduce a statutory tort for serious invasions of privacy, strengthen individual rights and expand regulatory powers. As a result, organisations without mature governance frameworks will face significant new exposure.

OAIC Notifiable Data Breaches (NDB) scheme. Entities covered by the Privacy Act must notify the OAIC and affected individuals of eligible data breaches. Persistent oversharing and weak access controls are among the most common root causes.

APRA CPS 234. Financial services entities regulated by APRA must maintain information security capability in line with the size and extent of threats to their information assets. Data classification and access governance are foundational requirements.

SOCI Act. Operators of critical infrastructure assets face mandatory incident reporting and risk management program obligations. Data integrity and security posture are directly in scope.

Each of these frameworks creates specific obligations around data visibility, access control and incident response. A governance program that addresses them systematically rather than each as a separate compliance exercise is more efficient and more defensible to regulators.

ASD Essential Eight. The Australian Signals Directorate's baseline mitigation strategies include application control, patching, restricting admin privileges and multi-factor authentication; all of which depend on a well-governed, classified data estate.

The cost of getting it wrong:

83%

Of organisations had more than one data breach.[7]

$3.9m

is the average cost of a data breach to an Australian organisation.[8]

20%

of data breaches are from insider threats.[9]

Governance: a driver of business outcomes

Future-focused organisations are reframing governance from a defensive to an offensive capability.

When in place, analytics teams can move faster because they don't have to check data quality on every project. Cloud migrations can proceed with confidence because classification and access policies travel with the data.

AI initiatives can launch, with leaders confident that only properly governed information will be surfaced to tools like M365 Copilot.

> ***"When people see data governance as a brake, they apply it like one – and the business slows down. Instead, it should be viewed as a strategic enabler, helping every executive, analyst and AI move faster, because they're finally working with information they can trust."***
>
> Richard Dornhart, National Practice Manager – Security, Data#3

7 "The crucial role of data security posture management in the AI era", Microsoft 2025
8 "New world, new rules: cybersecurity in an era of uncertainty", PwC 2025
9 "The crucial role of data security posture management in the AI era", Microsoft 2025

Bringing order to data estates

Structure that builds integrity.

Modernising a data estate demands a well-structured approach with the goal of creating a platform that can absorb and process data in real-time, incorporate machine learning-driven insights and drive more informed business decisions.[10]

Data#3 takes a risk based, end-to-end approach to data management, focused on ensuring data flows are properly managed and controlled.

This reduces reliance on user behaviour and limits oversharing in the places it most often occurs: SharePoint, OneDrive, Teams and email. That means defining the rules once, embedding them and letting governance work in the background.

10 "Transforming data architecture for intelligent business decisions." Medium 2025

Microsoft Purview: the platform that delivers governance at scale

Data governance capabilities such as classification, labelling, lifecycle management, insider risk detection, legal hold and AI oversight are delivered natively through Microsoft Purview.

Understanding which Purview component does what allows organisations to build a purposeful architecture rather than activating features in isolation.

Microsoft Purview components and what they govern:

Information protection.
Applies sensitivity labels to documents, emails and data across M365 and connected cloud services. Labels drive encryption, access restrictions and visual markings, and travel with the data wherever it moves.

Data lifecycle management.
Governs how long data is retained and when it is deleted, ensuring the organisation meets legal hold obligations without accumulating unnecessary data that expands the breach surface.

eDiscovery.
Enables legally defensible search, hold and export of relevant content for litigation, regulatory inquiry or internal investigation, supported by content that is labelled and catalogued in advance.

Insider risk management.
Detects anomalous data movement and potential exfiltration by correlating user activity signals across M365 workloads, supporting both HR and security response workflows.

DSPM for AI.
Specifically designed for AI-era governance, it surfaces sensitive data that AI tools can access, flags oversharing risks before Copilot deployment, and helps enforce data boundaries within AI environments.

These components work best as a co-ordinated program rather than stand-alone tools. Data#3's five-step framework is designed to activate them in the right sequence: establishing classification before lifecycle, lifecycle before eDiscovery readiness, and all of the above before AI deployment.

The Data#3 five-step data management framework

Our framework provides a repeatable path to achieving measurable control over your organisation's data across its lifecycle.

Step 1

Define a custom Information Classification Policy.

Step 2

Identify a clear picture of your data.

Step 3

Classify and track data.

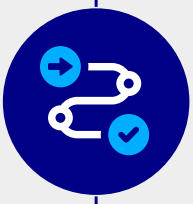

Step 4

Assess controls and develop an improvement roadmap.

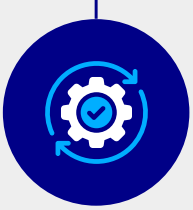

Step 5

Implement and optimise new data security controls.

Data governance delivered in three stages

Once your organisation has its information management under control, Data#3's three-stage delivery model sets and maintains the momentum needed to elevate data governance to industry-leading standards.

Stage 1: Strategy

Our business advisory team works with you to understand organisational needs and define the target state.

Stage 2: Solution

We design the base technology layer needed to deliver against that target state.

Stage 3: Deployment

Our professional services team partners with you to execute the program and embed it into the way the business runs.

The Data#3 approach is holistic, but we're flexible in where we can help.

It could be a modular engagement that targets one high-risk workload, or an end-to-end program that brings the entire estate under structured governance. The right answer depends on your starting point, not on a one-size-fits-all template.

Regardless of where you start, the outcome is the same: data, controls and policy moving together as a single system, rather than three separate workstreams that constantly drift in and out of sync.

> ***"When data governance is built-in from the beginning, and not just bolted onto a finished project, security and compliance benefits stop being trade-offs. They become accelerators, and every new workload starts from a position of strength."***
>
> **Bruce Hiddle, National Practice Manager Information and Analytics, Business Aspect**

The power of visibility and control

SHIFT your data with confidence, no matter where it lives.

As data moves between cloud workloads, SaaS apps and collaboration tools, the question of who has access to what, and whether they should, becomes harder to answer with any confidence. Periodic audits and manual permission reviews simply cannot keep pace with the speed of modern work.

This is where data security posture management (DSPM) comes in. DSPM offers a 360-degree view of data at rest and in-transit, and the risk associated to users.[11]

It is the layer that gives security and data teams visibility and control across sprawling data estates. Instead of providing a point-in-time snapshot, it continuously assesses data security posture, shows gaps and helps remediate risks under a single, comprehensive framework.

11 "The crucial role of data security posture management in the AI era" Microsoft, 2025

Identity: the foundation of access governance

DSPM tells you where your data is and who can see it. But without a robust identity layer, the access controls it enforces are only as strong as the credentials behind them. Identity governance is the precondition for meaningful DSPM and for any AI deployment that handles sensitive data.

Microsoft Entra ID and the identity controls that underpin data security

Microsoft Entra ID is the identity platform that authenticates users and enforces access policies across M365, Azure and connected applications. Clean Entra ID governance is the baseline for every access control above it.

Microsoft Entra strengthens data governance through:

Conditional access.

Enforces access policies based on user, device, location and risk signals. It prevents data from being accessed by unmanaged devices or high-risk contexts, regardless of whether the user has valid credentials.

Privileged identity management (PIM).

Governs just-in-time access for administrative roles, ensuring elevated permissions are temporary, audited and require justification. It also directly reduces the blast radius of an insider threat or compromised admin account.

Data#3's architecture-led approach addresses identity and data governance as a unified program. Microsoft Entra ID hygiene, Conditional Access policy design and PIM configuration are scoped alongside the DSPM and Microsoft Purview workstreams.

What DSPM does for your data estate

A DSPM contributes to building a secure data environment by flagging overexposure and identifying the issues that quietly accumulate over time:

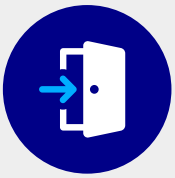

Excessive access.

Users and groups with broader permissions than their roles require.

Outdated permissions.

Access tied to dormant accounts, former projects or decommissioned systems.

Role creep.

Permissions that have been added over time and never revoked, even after a role change.

Posture drift.

Environments that gradually fall out of alignment with policy as configurations change.

Just as importantly, DSPM detects when policy deviations occur, alerting teams to changes that move the estate away from its target state.

The effect is to keep the data estate aligned with governance policies in real-time, rather than relying on quarterly reviews to surface problems.

DSPM and the SOC: from visibility to response

A DSPM that operates in isolation is a reporting tool. However, one that is integrated with security operations becomes a response capability. Data#3 architects DSPM deployments that feed signals into the broader security stack, closing the loop between detection and action.

How DSPM signals flow into SOC operations

Microsoft Sentinel

DSPM findings such as posture drift, excessive access and policy deviations are ingested as signals into Sentinel, where they are correlated with identity, endpoint and network telemetry to surface high-fidelity incidents rather than isolated alerts.

Microsoft Defender XDR

Data exposure findings from Purview DSPM for AI and Information Protection integrate with Defender XDR's unified incident queue, giving SOC analysts data-context alongside threat context when triaging an event.

Microsoft Security Copilot

SOC analysts can query Security Copilot in natural language to investigate data-related incidents, identifying what sensitive data was involved, who accessed it, from where, and what policy was in place at the time. This materially reduces mean time to investigate (MTTI) for data-related events.

This integration converts a governance program into a live operational capability. It is the difference between knowing you have a posture problem and being able to act on it, and between a data security program that satisfies an audit and one that responds to a real event.

How Data#3 puts DSPM to work

Data#3 brings an architecture-led approach to visibility and control.

We use DSPM to show where sensitive data lives, where it is exposed and where it is overshared.

From there, we deploy Microsoft Purview to put the controls, monitoring and remediation workflows in place so that security is integrated across the data estate, not bolted on at the edges.

This is what SHIFT looks like in practice. Data moves freely across your cloud and collaboration platforms but does so under a layer of visibility and control that makes the movement safe. You know where your data is, who can see it and how it's protected at every point across its lifecycle.

> ***"DSPM transforms data security from a periodic audit into a continuous capability. It gives you a live view into who is seeing what, which is exactly what an AI-ready data estate demands."***
>
> **Scott Gosling, National Practice Manager - Microsoft, Data#3**

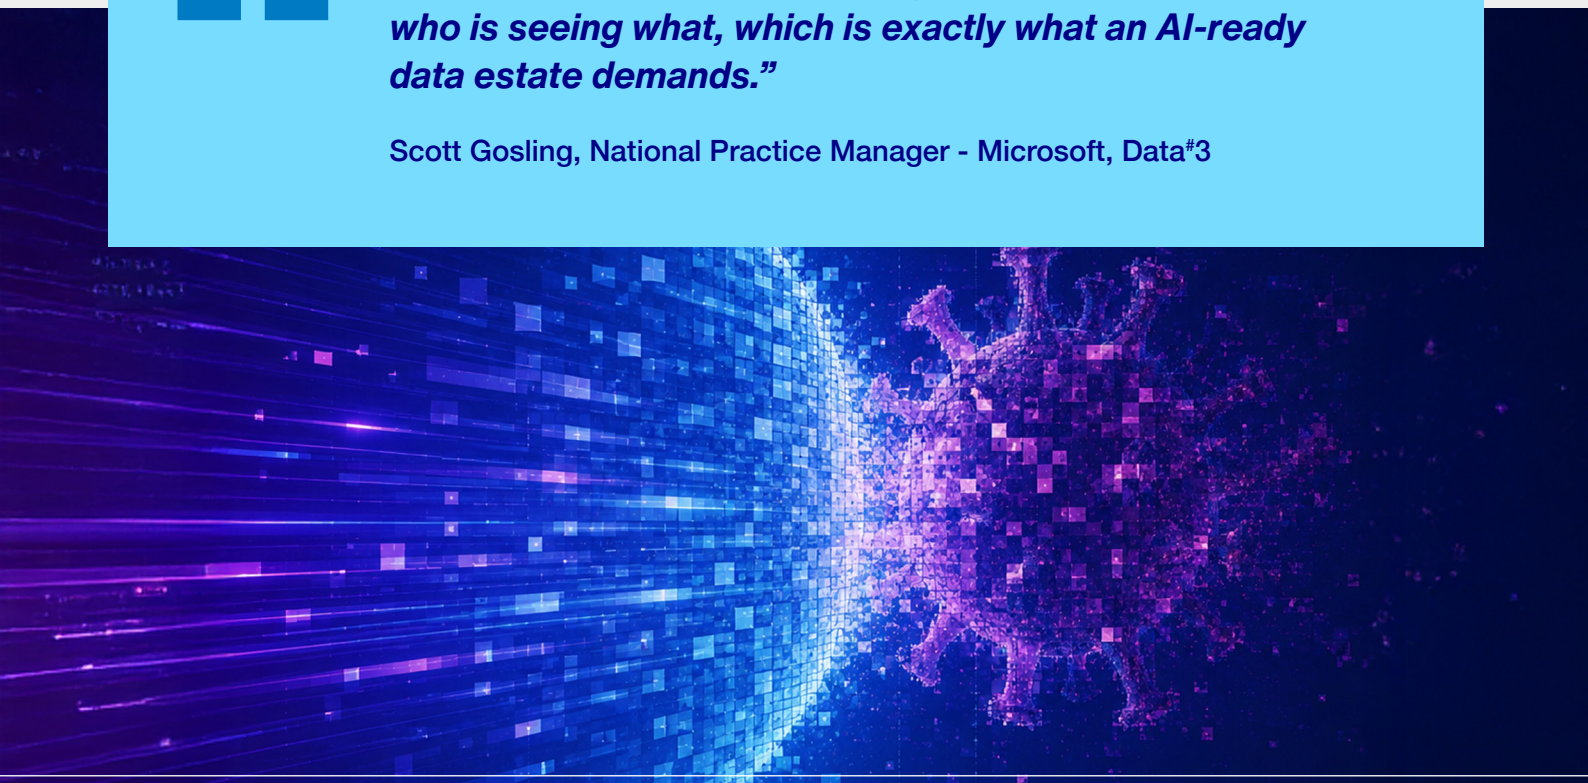

Data and AI

DELETE uncertainty from AI.

As we've already mentioned, AI's potential to transform organisations is only as effective as the data access controls that support it. People are far more likely to back AI when they are confident that critical data can't be viewed by the wrong people, and that outputs are based on information they can trust.

That is the core tension in most AI programs today. Organisations are moving forward with AI without cleaning up data first, increasing the likelihood of reputational fallout from a security or compliance breach.

Why poor data is so dangerous in the AI era

Subpar data governance increases the likelihood of data breaches, but the risk profile is broader than that.

AI tools such as Microsoft 365 (M365) Copilot and agents are only as good as the data they use. When that data is wrong or outdated, the consequences cascade through the business:

Executives

may make strategic decisions on AI outputs built from bad data.

Sensitive information

may be surfaced to employees who shouldn't have access to it.

Customers

may receive responses that contradict the organisation's own policies.

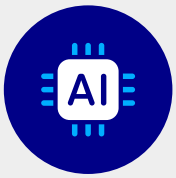

Models trained on poor data

will continue to produce unreliable outputs, eroding trust in AI itself.

According to IBM, poor data quality is one of the most common reasons AI initiatives fail.[12] Moody's analysis goes further, finding that organisations succeeding in GenAI share three traits, the first being a rigorous data governance framework.[13] Without those foundations, AI stalls.

12 "Why AI data quality is the key to AI success"
13 "GenAI and data quality: paving the path to AI success" Moody's, 2025

Building AI-ready data foundations

Your organisation's data readiness can boost ROI from AI investments and kick-start new, more innovative initiatives.

When AI learns from and draws from properly governed data, your organisation unlocks its potential safely and responsibly.

Data#3 offers a Microsoft Data Envisioning Workshop to address AI readiness concerns and develop a remediation roadmap. The workshop boosts visibility into sensitive data exposure, identifies oversharing and surfaces the policy and control gaps that would otherwise undermine an AI rollout.

For many organisations, this is a critical first step on their AI journey, often before a single M365 Copilot license is deployed at scale. Done well, data readiness is what makes AI a strategic capability in your organisation, rather than an unwieldly, chaotic tech blunder.

According to PwC, investing in data and AI together is a way to unlock benefits that can "catapult organisations into new realms of competitiveness", but only if the data underneath is fit-for-purpose.[14]

14 "Crafting a fit-for-purpose data and AI strategy in the era of generative AI" PwC, 2024

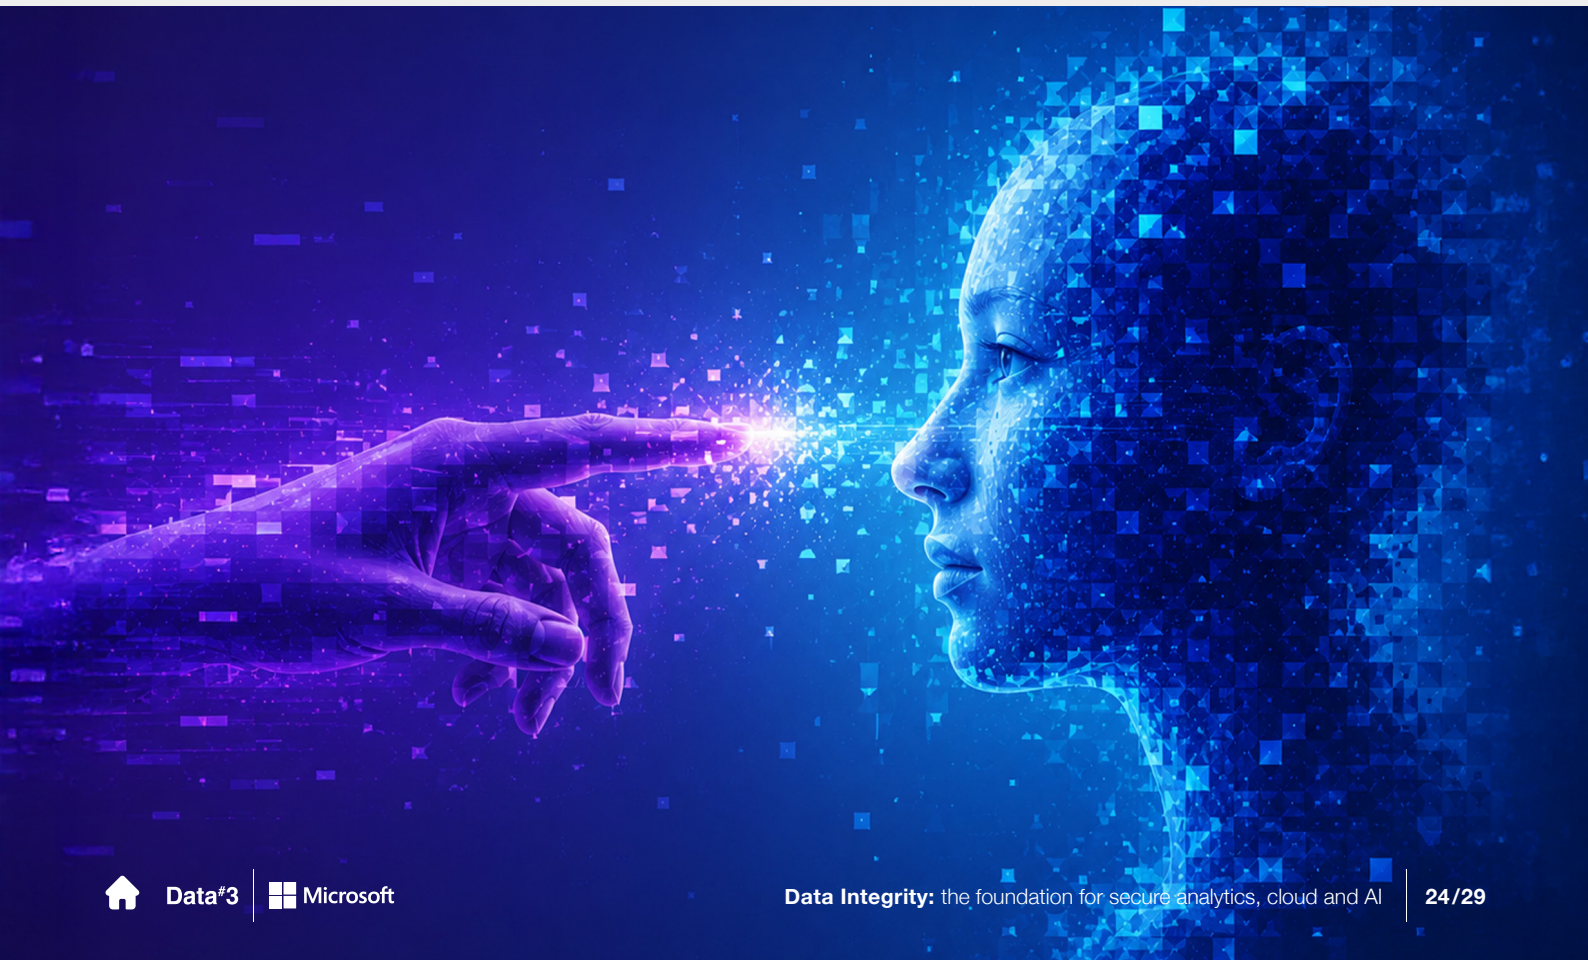

Understanding the shared responsibility model

A common misconception is that Microsoft's platform-native controls are sufficient on their own. In practice, Microsoft operates on a shared responsibility model, securing the infrastructure and platform while the customer is responsible for their data, identities and how access is configured within that platform.

Microsoft's responsibility

Physical infrastructure, platform availability, baseline encryption in transit and at rest, and the security of the Microsoft cloud itself.

The customer's responsibility

How data is classified and labelled, who has access to what and why, how long data is retained, how AI tools are scoped and governed, and how incidents within your tenant are detected and responded to.

Where Data#3 adds value

Data#3 designs the governance architecture, tunes the Purview configuration, establishes the identity hygiene, integrates DSPM signals into your SOC and operationalises the policies that the platform alone cannot enforce.

> ***"AI doesn't whisper, it shouts. A single bad data set can inform strategic decisions before anyone realises the original inputs were flawed. That's the real risk of poor data in the AI era – not just the wrong decisions, but the speed at which these can travel across your organisation."***
>
> Mark Hindson, Microsoft Security Solutions Architect, Data#3

Building the business case

Aligning data security with business outcomes.

Many organisations are accelerating AI investment, but proving ROI is slow and difficult to isolate. According to Deloitte, 91 per cent of organisations plan to increase their investment in AI over the next 12 months, yet most report not achieving a satisfactory ROI until two to four years in.

That is substantially longer than the typical seven to 12 month window leaders are accustomed to from other technology investments. Only six per cent report returns in under a year, and even among the more successful projects, only 13 per cent saw returns within 12 months.[15]

Part of the reason is that AI rarely delivers value in isolation. It is usually introduced alongside efforts to improve data quality, streamline operations or reconfigure teams.

This makes its specific contribution difficult to isolate from broader transformation work.

The business case for AI, in other words, is really a business case for the data and process foundations that make AI useful.

15 "AI ROI: The paradox of rising investment and elusive returns" Deloitte 2025

Securing your investment from day one

Robust data security can play a critical role in helping organisations justify and sustain investment in AI.

When data is secure, monitored and controlled no matter where it is stored, accessed or moved, it builds trust with the employees who use it and enables your business to pursue outcomes with confidence. That trust is the precondition for adoption, and adoption is the precondition for ROI.

Data#3's advisory team brings the expertise to help build the business case for AI.

By aligning data security with business outcomes like reduced breach risk, faster project delivery, lower remediation costs and stronger employee trust, we create a strong argument for increased AI investment.

Meeting you where you are

Whether your business is establishing solid data foundations or is already relatively advanced, Data#3 meets you where you are. Our engagement is tiered to match your maturity and ambition:

Good

Our foundation deployment that establishes baseline visibility, classification and control across the highest-risk workloads.

Better

Our governance-led deployment that embeds governance practices, policies and ownership across the broader estate.

Best

Our sustainable governance approach that operationalises governance and DSPM as a continuous capability that scales with the business

Organisations that invest in robust data foundations are better positioned to realise long-term value. The investment in security and governance is the lever that makes AI ROI defensible to the board, resilient across change and scalable across the business..

"Organisations rarely lose AI budget because the technology created no value. They lose it because they can't articulate, in plain language, why the next dollar will be safer than the last. Secure, well-governed data helps you craft a compelling story for executives, demonstrating that data integrity is how AI systems deliver trusted results, scale safely and support long-term ROI."

Paul Green, Principal Consultant, Business Aspect

Data#3 | Microsoft

About Data#3

Data#3 is one of Australia's leading technology services and solutions providers, partnering with organisations to turn data into a trusted, strategic enabler of business outcomes. As Microsoft's #1 Australian partner, we offer deep expertise across Microsoft security, data and AI; and our teams combine advisory, architecture and managed services to help customers build secure foundations, modernise their data estates and adopt AI with confidence.

Data#3 brings a proven, tested end-to-end approach grounded in the Microsoft platform and shaped by the realities of the Australian market. We help our customers CTRL their data, SHIFT their workloads with confidence, and DELETE uncertainty from every AI initiative they take on.

Learn more at Microsoft Security Solutions | Data#3

1300 23 28 23

www.data3.com.au

facebook.com/data3limited

x.com/data3limited linkedin.com/company/data3

youtube.com/user/data3limited