

ShareGate

The State of **Microsoft 365**

2026

Migration, governance,
and the AI reckoning in
2026



What's in the report


01	Confidence ≠ control	3
	Confidence is a feeling. Incidents are facts. In 2026 they don't line up.	
	Confidence up, exposure up	4
	The pattern holds	6
	Already handled it?	8
	The geography of the gap	10
	The state of AI usage	14
	The concerns teams name	17
	What AI can reach	20
	What "no incidents" means	23
	Who owns it	26
	What to do	32

02	The AI reckoning	33
	Three quarters know what they want AI to do. Half can't see what it costs.	
	What AI is costing	34
	AI licensing tops the cost list	36
	The AI cost visibility problem	38
	Why AI ROI is hard to measure	40
	What teams do about cost	43
	Governance is the key to AI ROI	44
	Who is most confident in ROI	44
	Geography of spend and return	48
	Cost in the broader estate	50
	What to do	56

03	Governance	57
	77% had an incident. 1% use a tool built to prevent one.	
	What actually breaks	58
	When the incident is the alert	60
	What teams are governing with	63
	What's blocking improvement	65
	How teams rate themselves	68
	Links, guests and workspaces	71
	Guest access reviews	74
	What to do	77

04	The estate	78
	Creation is automatic. Retirement is manual. That's the whole problem.	
	The 68% in between	79
	What cloud-complete reports	81
	What the estate is made of	84
	SharePoint end of support	85
	The third that stayed	87
	Self-service postures	89
	Workspace lifecycle	91
	What holding the estate costs	95
	What to do	100

05	Compliance	101
	Not a chapter about restriction. A chapter about evidence.	
	A stop sign	102
	Jurisdictions	104
	What it consists of	107
	The approval cycle	109
	The auditing gap	111
	What holdouts need	113
	Fully cloud, still gated	115
	Sovereignty and AI	117
	What to do	124
06	Migration	125
	Projects end. This doesn't — 94% moved something, 2.5 types in two years.	
	An ongoing condition	125
	How long it takes	129
	Why teams move	130
	What actually moves	132
	What teams use	133
	What they'd change	133
	What predicts buying	135
	Identity	137
	The parts that go wrong	139
	Who gets called in	143
	What to do	146
07	The operating model	147
	Seven moves that turn the gap between belief and proof into a control loop.	
	The gap	147
	The operating loop	148
	The people layer	149
	Seven moves	150
	The cadence	158
	Sequencing	159
	Where to start	160
	About the research	164
	Who we surveyed, how to read the data, and how to cite it.	
	Index by topic	169
	Seventeen topics, with every page each one appears on.	



01 Confidence ≠ control

CHAPTER 1 IN BRIEF

93% are confident their governance framework is ready for AI. Yet 29% have already had it surface something it shouldn't have had access to.

71%

say their governance workload has grown since deploying AI

67%

of tenants run two or more AI tools, and 28% run three or more

48%

run continuous automated monitoring of AI usage

WHAT THAT MEANS

Confidence and exposure point opposite ways.

IT leaders overwhelmingly report feeling ready for AI. The exposure data doesn't match. The teams doing more governance work report more incidents and more concern; the teams doing less report less of both. Not because they're safer. Because they're not looking.

WHY IT MATTERS MORE NOW

Full Copilot deployment nearly doubled in a year.

Full Copilot deployment went from 29% to 56% between 2025 and 2026, and 93% of organizations now have it in production. Deployment isn't the interesting question anymore. What's underneath it is.

WHAT HASN'T CHANGED

The environment AI inherits.

53% of respondents don't have full visibility into what their Copilot agents can access, and permissions on Microsoft 365 content don't expire on their own. The tools moved fast. The environment underneath them didn't.

THE GAP

Concern is concentrated where monitoring already is.

Half of IT leaders (49%) are very or extremely concerned about AI reaching content nobody has reviewed. But that concern isn't evenly spread: it falls at every step away from active monitoring, from 54% among teams monitoring continuously down to the teams doing nothing at all, without reversing once. The teams with the least visibility are the least worried about it, which is the wrong way round.

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 1 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

See how confident you should be. The M365 Governance Index scores your estate against the 943 IT pros in one study and your AI governance against the 851 in the other.

Get your score →


Confidence is a feeling. Incidents are facts. And in 2026, they don't line up the way you'd expect them to.

This chapter is where we open up the gap between the two. What the confidence looks like when you break it down. How many incidents there actually are when you count them. How both change depending on the country you're in. And why we think the two are so important to look at together. Not because IT teams are wrong about what they've built, but because the surface they're managing keeps getting harder to see all at once at scale.


THE GAP, UP CLOSE

Confidence went up. So did exposure.

We asked 851 IT leaders how confident they are that their organization's current Microsoft 365 governance framework is sufficient to support Copilot and other AI initiatives responsibly. Half of them (51%) said "very confident", while another 42% said "somewhat confident".



The next question asked whether Copilot or another AI tool has ever surfaced sensitive internal data that should not have been accessible. 29% said yes. Another 8% said they weren't sure.



Two questions, same respondents, seemingly contradictory results. One picture of readiness. One picture of what's already happened.

More than a third of organizations either have a known AI exposure incident on the books or can't say whether they have, from a population that overwhelmingly reports feeling confident that they're supporting Microsoft Copilot and other AI initiatives responsibly.


“Interestingly, according to IBM, 77% of C-level technology executives say that AI adoption is already outpacing governance capabilities. So this confidence may be higher in IT professionals versus IT executives.”



RICHARD HARBRIDGE
MICROSOFT MVP

These incidents weren't a case of AI surfacing old, outdated information that slowed end user productivity, which is also a possible outcome of not properly governing your tenant for AI. They're incidents of AI surfacing the kind of content teams are typically asked to prove they control.

No obvious pattern in what AI reached: customer records, personal data, HR, financial documents, IP. Everything sensitive was roughly equally exposed. When AI overreaches, it doesn't fail selectively across the categories teams normally treat as separately governed.




The pattern holds across frameworks, permissions, and remediation

And it isn't a one-off. The same survey asked how confident respondents are that Copilot's access reflects current and appropriate permissions across SharePoint, Teams, and OneDrive. Again that magic number: 93% confident, 48% of them "very". Asked whether their IT team has the skills and tools to remediate AI-related governance issues if something goes wrong, it held a third time: 93% confident, 50% "very".

93% are confident the permissions are right

WE ASKED How confident are you that Microsoft Copilot's access reflects current and appropriate permissions across SharePoint, Teams, and OneDrive?




Nearly as many are only somewhat confident as are very.

AI GOVERNANCE SURVEY · N=851

Confident they have the skills and tools to fix it

WE ASKED How confident are you that your IT team has the skills and tools needed to remediate AI-related governance issues?



93% again, this time on having the skills and tools to remediate an AI governance problem. The third near-identical confidence number in the same study.

AI GOVERNANCE SURVEY · N=851

Altogether, 88% are confident that the framework is right, the permissions are right, and the team can fix what needs fixing. The paradox gets sharper when you isolate the strongest confidence group.

35%

"Very confident" on all three questions

framework · permissions · remediation · 297 of 851

YET

37%

of that 35% still had AI surface data it shouldn't

n=297 · +8 pts vs the 29% sample average

AI GOVERNANCE SURVEY · N=851 · n=851

Confidence and exposure aren't moving in opposite directions here. They're moving in the same direction, together. Microsoft MVP Eric Overfield puts a name on the mechanism.

“ Confidence is overstated and incidents are understated, in opposite directions. People say they're ready because the controls exist. They under-report the exposures because nobody volunteers those. What you're looking at is the gap between having controls and checking whether they worked.”



ERIC OVERFIELD
MICROSOFT MVP

WHERE SHAREGATE FITS

ShareGate Protect was built to close the gap Eric is talking about, then keep it closed. You can see who has access to what, where oversharing is happening, and what Copilot can actually reach. Apply automated policies to keep the environment clean as it changes, and fix issues right where you found them. No scripts, no admin-centre hopping.


[See ShareGate Protect →](#)

Whatever the 29% are looking at when they answer the incident question, it isn't visible in any of the confidence questions around it.

Are they confident because they've already handled it?

You might reasonably think that teams with incidents are confident *because* they've dealt with them: earned survivorship, new mechanisms in place, rather than a disconnect. The data doesn't support it. Respondents who reported an AI incident are nearly twice as likely to be concerned about AI accessing content that hasn't been recently reviewed for permissions:


Very or extremely concerned about AI reading unreviewed content



AI GOVERNANCE SURVEY · N=851

They're more than twice as likely to say governance workload has increased significantly since enabling AI:


"Governance workload has increased significantly since enabling AI"



AI GOVERNANCE SURVEY · N=851

And they rate their own governance maturity at roughly the same level as everyone else:

Rate their own governance "highly automated and continuously monitored"



Effectively identical. An incident doesn't dent how mature a team believes it is.

AI GOVERNANCE SURVEY · N=851

That isn't a "we fixed it and feel better" profile. It's a "we know it's still possible, we're working harder, and we're still worried" profile, held at the same time as saying they're confident in their framework, their permissions, and their ability to fix things. The confidence and the concern coexist.

The geography of the confidence gap

The averages hide something worth surfacing. Those confidence figures and the 29% exposure rate don't hold up evenly across countries.

Confident / very confident by country, with exposure alongside							
Country	n	Framework	Permissions	Remediation	Very conf. all 3	Incident	Not sure
United States	219	96% / 61%	94% / 62%	95% / 62%	49%	37%	9%
United Kingdom	204	97% / 60%	94% / 55%	94% / 58%	42%	26%	9%
Canada	106	92% / 45%	96% / 42%	92% / 45%	30%	26%	9%
France	108	91% / 42%	93% / 44%	88% / 34%	25%	29%	8%
Germany	104	90% / 45%	91% / 35%	96% / 39%	21%	23%	5%
Ireland	54	89% / 30%	89% / 32%	91% / 43%	24%	20%	13%
Netherlands	56	86% / 38%	86% / 32%	89% / 39%	20%	30%	5%
Sample average	851	93% / 51%	93% / 48%	93% / 50%	35%	29%	9%

AI GOVERNANCE SURVEY · N=851

Each confidence column shows two numbers: the share who were confident at all, then the share who were very confident. So 96% / 61% means 96% confident, and 61% of the whole group very confident.


No two countries match on both confidence and incidents. Four are worth examining more closely, each a distinct shape of what the paradox looks like in practice.

The US: where the paradox is sharpest


US respondents lead the sample on every "very confident" dimension and on the comprehensive measure: 49% "very confident on all three", 14 points above average. IT pros in the US also reported the highest exposure rate in the study. Inside that most-confident US cohort, the incident rate climbs to 43%. The strongest confidence in the sample sits in the same country as the highest exposure, and within it the two track together more tightly still.

United States: "very confident", against everyone else


Framework is sufficient




Permissions are appropriate



We could remediate




All three




AI GOVERNANCE SURVEY · N=851 · United States n=219 · rest n=632

United States: exposure, against everyone else

AI surfaced something it shouldn't



Can't say whether it has




AI GOVERNANCE SURVEY · N=851 · United States n=219 · rest n=632

The UK: US-level confidence, sample-level incidents


The UK shares the confidence but not the exposure. Effectively tied with the US across all three questions, 7 points behind on the comprehensive measure, and 12 points below on incidents. Whether that reflects better controls, different reporting norms, a smaller AI footprint, or something not visible here such as stricter regulation, we can't say. What it does say is that comprehensive confidence and comprehensive exposure aren't a one-to-one match. The association appears in this sample, but our survey can't establish the cause.

United Kingdom: "very confident", against everyone else


Framework is sufficient




Permissions are appropriate



We could remediate




All three




AI GOVERNANCE SURVEY · N=851 · United Kingdom n=204 · rest n=647

United Kingdom: exposure, against everyone else

AI surfaced something it shouldn't



Can't say whether it has




AI GOVERNANCE SURVEY · N=851 · United Kingdom n=204 · rest n=647

Germany: asymmetric confidence, low incidents


Germany shows a different shape. German respondents sit mid-table on whether their framework is sufficient and their permissions appropriate, but on remediation they come out highest in the sample. Less certain the setup is right, more certain they can fix it if it isn't. Only 21% are "very confident" on all three, one of the lowest in the study, and the exposure rate is the second-lowest at 23%. Where comprehensive confidence softens, the incident count softens with it.

Germany: "very confident", against everyone else


Framework is sufficient




Permissions are appropriate



We could remediate



All three




AI GOVERNANCE SURVEY · N=851 · Germany n=104 · rest n=747

Ireland: lowest incidents, highest uncertainty


Ireland points somewhere else entirely: the lowest incident rate in the sample, and the highest "not sure" rate, four points above average. What Ireland describes isn't necessarily a safer environment. It's one where a meaningful share of respondents genuinely don't know whether an incident has already happened. Small base (n=54), so read it as directional. But the shape is distinct.

Ireland: exposure, against everyone else

AI surfaced something it shouldn't



Can't say whether it has



AI GOVERNANCE SURVEY · N=851 · Ireland n=54 · rest n=797

Four countries, four different stories. What links them isn't a shared pattern in how confidence and exposure relate. It's that they don't relate cleanly in any of them. Whatever confidence is tracking, it isn't a smaller exposure surface. It may be a narrower view of one.

WHERE DO YOU SIT?

The country tables above are averages, and the averages hide more than they show. On the M365 Governance Index's 100-point scale, the entire spread between the highest-scoring country and the lowest is 11 points. The spread between the middle half of respondents—25th percentile to 75th—is 24. Which country you're in explains far less about your AI governance than where you sit inside it. The Index scores you across cleanup, ownership, monitoring, exposure, and maturity against the 851 IT leaders in this study, in about five minutes. [Take the M365 Governance Index →](#)


THE DEPLOYMENT BEHIND IT

The state of AI usage

This confidence and exposure paradox isn't happening in a small number of pilot environments. The AI these leaders are governing, worrying about, and being surprised by is already deployed at scale. Copilot is almost everywhere now, and in most tenants it's one of two or three AI tools IT is managing at the same time. That combination is what makes visibility the question the rest of this chapter is built around.

Copilot has landed

WE ASKED Have you deployed Copilot in your organization?




Data

- **► Yes, we've fully deployed it** 56%
- Yes, we've partially deployed it 37%
- No, but we plan to in the next 12 months 5%
- No, we don't plan to 2%

93% already have Copilot in production. What these leaders are governing is already live.

AI GOVERNANCE SURVEY · N=851

A year earlier, the biggest group was still piloting



Data

- Yes, we have a pilot program in place 33%
- **► Yes, we have fully deployed Copilot** 29%
- Yes, we have partially deployed Copilot 20%
- Still evaluating 14%
- No, not planning to 4%

Full deployment roughly doubled in twelve months, 29% to 56%. Any deployment went from 82% to 93%.

IT OPERATIONS SURVEY · N=943

· **M365 Decoded 2025.**

93% of respondents have Copilot already in production, and 98% are either deployed or actively planning to deploy within a year.

LAST YEAR → THIS YEAR



+11 PTS

Deploying / using Copilot

82%
2025


93%
2026

Any Copilot deployment went from 82% to 93%. Full deployment nearly doubled underneath it, 29% to 56%.

Compare that to the same question a year ago: full deployment has roughly doubled, from 29% to 56%, and any deployment—pilot, partial, or full—has gone from 82% to 93%. A year ago the largest group was still piloting. Today the largest group is fully deployed.

The 2026 answer set has no "pilot" option, so 2025's pilot respondents would most likely pick "partially deployed" today.


And it isn't just Copilot. Most organizations are running more than one AI tool at once.



Among the organizations running or planning Copilot, 73% also run at least one other named AI tool. Barely a quarter of Copilot organizations are Copilot-only.

Copilot is rarely alone in a tenant

WE ASKED What AI tools is your organization currently using or planning to deploy?



73% of Copilot organizations also run at least one other named AI tool. Barely a quarter are Copilot-only, so "what can Copilot reach" is the wrong question. The question is what any of them can reach.

IT OPERATIONS SURVEY · N=943

· Among organizations using or planning Copilot (n=728).

That matters for everything that follows. Permissions, monitoring, and remediation have to work across the whole set, not just the one with Microsoft's name on it. The visibility question in the next section isn't "what can Copilot reach", it's "what can *any* of these reach", and for two thirds of organizations that's a question about several products at once.


The concerns teams name themselves

Deployment is one question. What teams are actually worried about is another, and the IT operations survey put it to a different 943 respondents.


Same top two worries, two years running

WE ASKED What are your biggest concerns about using Microsoft Copilot? (choose up to three)


Data quality, relevance, retention




Security and access




Lack AI governance expertise




Realizing Copilot's ROI



Cost control



Governing agents, plugins, automations




Deployment doubled in the same twelve months. The worry list didn't reorder.

IT OPERATIONS SURVEY · N=943 · 2026 n=943 · 2025 n=623


Why the 2026 numbers above sit low compared to 2025

Both years capped you at three picks. 2026 offered one more thing to pick from, so every number on the chart above sits a little low. More options, same three slots. Here is a quick visual of what we mean:



In other words the “Lack of AI governance expertise” shows 37% in our chart, but if we account for the extra option the more honest answer is a range: 37% at the low end, 42% at the high end. Up a point on last year, or up six. It went up either way. How far up, we can't say.

Every other concern on this chart fell, whichever way you count it, with one other exception.



Copilot ROI is the one that could have gone either direction, which you can see in the visual explainer here.


The top two are tied, and neither is about the technology. Data quality and retention is a content governance problem. Security and access management is a permissions problem. Both existed before Copilot arrived, and both decide what it can surface.

Last year those same two led the list, at 58% and 57%. They still lead, in the same order, at slightly lower intensity. That's the part worth sitting with: full Copilot deployment roughly doubled in the same twelve months, and the concern set didn't reorder. Whatever teams were worried about before they scaled, they're still worried about now. Most of the drop in the levels is the extra option on this year's list rather than a change of mind, which the note under the chart explains. What's left of it invites two readings: normalization, familiarity reducing alarm; or the pattern this chapter keeps finding, that awareness travels with practice. As deployment broadens from early adopters to everyone, the average respondent is newer to governing AI and has looked at it less closely. Neither is comforting.

Then there's the bottom of the list, which turns out to be the most interesting entry on it. Governing AI agents, plugins, and automations ranks *last* as a concern, named by 19%. But the same survey asked what would have the biggest impact on improving governance and security, and the top answer, ahead of executive buy-in, automated remediation, better tenant visibility, more expertise, and more budget, was better controls for managing AI agents, plugins, and automations.

What would help most? AI controls.

WE ASKED What would have the biggest impact on improving your team's Microsoft 365 governance and security?



The governance ask moved to AI. The tooling mostly hasn't.


IT OPERATIONS SURVEY · N=943

Same respondents. Two questions. Agents come last as a worry and first as a want. That's less contradictory than it first looks, because the two questions measure different things. Concern is largely retrospective: what has already gone wrong or is visibly straining. "What would help most" is prospective: the capability teams can see they're missing. Agents rank low on the first because for most organizations the agent layer hasn't hurt them yet. They rank first on the second because the control gap is already obvious.

Which makes this the one place in either study where teams appear to be ahead of the problem rather than behind it, and worth acting on for exactly that reason. The agent surface is the fastest-growing part of a tenant and the least governed, and the window in which this concern stays hypothetical is unlikely to stay open for long.

What AI can reach, and what IT teams can see

The AI running in a typical Microsoft 365 tenant has permissions inherited from the tenant, access to content the tenant already contains, and no built-in limit on what it can surface if the underlying permissions and content aren't clean. There's a difference between knowing AI is active in your tenant and knowing what it can touch. Just under half of IT leaders have full visibility.



53% of respondents have less than full visibility into what their agents can access, and the largest group inside that 53% is "partial". These teams aren't flying blind, but they don't have the specificity that would let them say, in front of an auditor or a legal team, *this agent, in this business unit, can reach this data source, at this sensitivity level.*

“ Copilot follows the permissions you already have. If you don't understand your permissions model, you don't know what Copilot can reach. Teams spend the rollout on licensing and training, and visibility is the part that actually prevents the surprises.”



LAURA ROGERS
MICROSOFT MVP

That's the challenge partial visibility creates: it doesn't scale. As tenants add agents, data sources, and sensitivity classifications, the gap between "we know agents exist" and "we know what they can each reach" widens.

And it's not just how much they can see—it's also how often they look.

Half aren't watching in real time

WE ASKED How does your organization monitor how Microsoft Copilot and other AI tools are being used?



Only 48% monitor AI usage continuously; the rest find out later.

AI GOVERNANCE SURVEY · N=851

48% run continuous automated monitoring of AI usage. Another 37% run periodic monitoring. The remaining 15% do manual spot checks, only investigate after an incident, or don't monitor at all.

“ Most teams monitor Copilot adoption. Far fewer monitor what it's touching. Adoption tells you whether people are using it. Only data security monitoring tells you whether that's safe.”



JASPER OOSTERVELD
MICROSOFT MVP


Periodic beats spot checks, but it still only catches what's happening at the moment of the review. Between checks, things get away from you. Every partial-visibility and periodic-review respondent is running an environment where an audit question, an incident review, or a regulatory request would require an investigation to answer.

How worried teams are

Permissions on Microsoft 365 content don't automatically expire or reset. If a document was shared with someone in 2019 and nobody has reviewed those permissions since, it's still shared with them, which means AI can read it. We asked how worried IT pros are about exactly that.

Half are worried about what AI can already read

WE ASKED How concerned are you about AI tools accessing content that has not been recently reviewed for permissions?



49% are very or extremely concerned. Set that beside the 93% who say their governance is ready.

AI GOVERNANCE SURVEY · N=851

Half of AI leaders are very or extremely concerned. Add "moderately concerned" and it's 79%. Only 7% aren't concerned at all. That headline doesn't tell the full story, though. Cross it against how often teams actually monitor and a pattern emerges that runs the wrong way.

Concern about AI reading unreviewed content, by monitoring posture		
AI monitoring approach	n	Very or extremely concerned
Continuous automated monitoring	408	54%
Periodic monitoring	318	48%
Manual spot checks	77	40%
Only investigate after incidents	33	36%
No monitoring in place	15	13%
Sample average	851	49%

AI GOVERNANCE SURVEY · N=851

You'd expect the teams doing the least monitoring to be the most concerned. They're the ones with the least visibility into what AI is doing. The data shows the opposite, and it's monotonic: every step away from active monitoring produces a lower concern rate. The teams with less visibility are also the teams least likely to be worried about it.

That complicates how we read the 49% headline. Half the sample is very or extremely concerned, but that concern isn't evenly distributed. It's concentrated where monitoring is more active, in the 85% running continuous or periodic monitoring. Where governance discipline is thinner, so is the worry. The teams most at risk are the ones least likely to see it. This isn't the only place in the data where awareness travels with practice.


What "we haven't had an incident" actually means

Before AI ever reaches into content, most organizations have already done something about the content itself. We asked whether teams had run a structured content or permissions cleanup ahead of deploying or scaling AI.

Only half had done an organization-wide cleanup. Not ideal, but not unexpected.

Half cleaned up before scaling AI

WE ASKED Has your organization conducted a structured content or permissions cleanup in preparation for AI deployment or scaling?



Half ran an org-wide content and permissions cleanup first. A third did it in pockets, and one in seven has it on a list.

AI GOVERNANCE SURVEY · N=851

The interesting part shows up when you cross the cleanup data against exposure: how many respondents said AI had surfaced sensitive content it shouldn't have.

AI exposure, by whether the estate was cleaned up first				
Cleanup status	n	Had an incident	No incident	Not sure
Yes, organization-wide	433	40%	56%	4%
Yes, limited to certain departments	301	20%	70%	10%
Planned but not yet completed	84	11%	70%	19%
No cleanup	26	8%	69%	23%
Sample average	851	29%	63%	8%

AI GOVERNANCE SURVEY · N=851

On the surface this looks backwards. The teams that did the most cleanup—the ones you'd expect to have the fewest exposures—report the highest incident rate. Organization-wide cleanup produces a 40% incident rate. No cleanup produces 8%.


But the "not sure" column tells you what's actually going on. The teams that skipped cleanup entirely have a "not sure" rate nearly three times the sample average. The teams that did organization-wide cleanup have a "not sure" rate of half the average. That's the pattern. It isn't that cleanup produces incidents. It's that cleanup

produces *visibility*. Teams that audited their content and permissions can tell you whether AI has surfaced something it shouldn't have. Teams that didn't audit are much more likely to say they simply don't know.

Which means the low incident rate in the no-cleanup group isn't a signal of safety. It's a signal of not knowing. Those aren't fewer incidents. They're incidents that haven't been detected, or ruled out, either way. The concern data said the same thing from another angle. Doing the work produces awareness. Skipping it produces the appearance of safety, plus a meaningful share of respondents who genuinely can't answer what's going on in their tenant when asked.

What triggers real AI governance work

For the teams whose current operational state proves they take AI governance seriously—the 48% running continuous automated monitoring—we asked what first prompted the shift.



Regulators often require monitoring by law, so it tracks that a regulated industry would push AI governance up the list. If you can't demonstrate what your AI is doing, you don't just have a governance problem—you have a compliance problem. Microsoft MVP David Drever describes the specific moment governance becomes real:

“ Teams run a Copilot pilot and start finding information they never knew they had access to. That's the moment the fear arrives, and it's why my calendar is full.”



DAVID DREVER
MICROSOFT MVP


The triggers explain why teams started. What comes next is what that work actually looks like day to day.

Who owns it

Triggers explain why the work started. Ownership determines whether it continues.

Who owns AI governance

WE ASKED Which statement best describes AI governance ownership in your organization?



Fewer than half have an owner whose policies are actually enforced consistently. For a quarter there's an owner but enforcement varies by department. That's a name on a page, not authority.

AI GOVERNANCE SURVEY · N=851

Just under half of organizations report a clearly defined owner with formal policies and consistent enforcement. Another 26% have a defined owner, but enforcement varies across departments. 22% describe shared ownership across multiple teams with no formal governance structure, and 3% have no clearly defined ownership at all.

The middle two rows are where most of the risk sits. A named owner whose enforcement stops at the edge of their department produces exactly the pattern the exposure data shows: policy that's real in some parts of the tenant and theoretical in others. AI doesn't respect that boundary. It reaches wherever permissions allow, including into the business units where the policy was advisory. And shared ownership sounds collaborative and usually isn't. When four teams share accountability, the practical answer to "who decides whether this agent can access finance data" is often nobody, and the decision gets made by default, which is to say by whatever the permissions already allowed.

There's a capability question sitting underneath the authority question, and the other study puts a number on it: 37% of teams name lack of internal expertise on governance for AI as a top Copilot concern. Last year the same figure was 36%. A year of intensive Copilot deployment, in which full deployment roughly doubled, moved the AI governance skills gap by a single percentage point. The tooling arrived. The adoption arrived. The expertise to govern it didn't. That matters for ownership specifically, because naming an owner doesn't create capability. An owner without the skills to evaluate what an agent can reach is accountable for a decision they aren't equipped to make, which is how well-intentioned governance ends up ratifying whatever the permissions already allowed.


The pattern to watch for is a committee that advises and an admin who implements, with nobody who can say no. AI moves faster than that arrangement can respond.

For most teams, more AI means more work

The reason most organizations deploy AI is efficiency. The pitch has consistently emphasized time saved, tasks automated, work removed. But the governance side of AI is its own kind of work, and someone has to do it.

71% say AI made governance heavier, not lighter

WE ASKED Since enabling AI tools like Copilot, how has your governance workload changed?



One in ten got lighter. Those are the teams worth asking what they did differently.

AI GOVERNANCE SURVEY · N=851

Why the increase? Part of the answer is that the work itself isn't new. Microsoft MVP Susan Hanley sees governance that was always needed, now much harder to put off:

“ In some ways it hasn't increased as it was always there, it's just been raised in awareness and importance. The requirement for good governance hasn't changed. When Delve and other Microsoft innovation started surfacing content that wasn't secured properly, you could switch it off, and many customers did. There is no easy shutoff switch for AI, so now you have to do all the things you should have been doing all along.”



SUSAN HANLEY
MICROSOFT MVP

What has changed is the pace, and who is being asked to keep up with it. Drew Madelung, also a Microsoft MVP:

“ The rate of change has become unattainable for M365 admins. The people who used to organize SharePoint are now sitting on an AI steering committee, and they may not even know much about AI yet.”



DREW MADELUNG
MICROSOFT MVP

83 respondents said governance workload went *down* after they enabled AI. A small group, and worth being careful with. Three habits are more common among them than across the study, though not by enough to call any one of them the reason. The row that actually separates them is the last one.

Teams whose governance workload decreased after enabling AI, vs the sample		
	Workload decreased	Sample average
Run continuous automated monitoring	58%	48%
Have a clearly defined owner with consistent enforcement	55%	48%
Have done organization-wide or departmental content cleanup	89%	86%
Reported an AI exposure incident	16%	29%
Describe their governance as operationalized or better	77%	63%

AI GOVERNANCE SURVEY · N=851

n=83 of 851.

Continuous monitoring means the surveillance is running whether or not anyone's actively checking. Defined ownership means the accountability question is answered before an issue lands. Both are more common here, and neither gap is large enough to carry an explanation on its own. The cleanup gap is smaller still, close enough to the study average to be noise.

The incident row is the one to read carefully, because it is the easiest to misread. This group does report AI exposure well below the sample. But an exposure incident is itself governance work: among teams who had one, 88% say their workload went

up. A group defined by workload going *down* was always going to contain fewer of them. That row describes what being in this group means, not what the habits above bought.

What does hold up is the last row. Three quarters of this group describe their governance as operationalized or better, against 63% of the study, and that gap is the one still standing when you account for all three habits at once. It isn't a checklist. Most of this group is missing at least one of the three. It's whether the whole thing runs as an operation.

EVERY READER WANTS TO BE IN THE 10%

The M365 Governance Index tells you how close you actually are. Compare your environment against the 851 IT leaders in the AI-governance study, filtered by your region, size, and industry, for a peer-ranked score on cleanup, ownership, monitoring, exposure, and maturity, plus the dimensions where the gap between you and the teams whose workload got lighter is widest. [Get your score →](#)

That's the operating model. Not "deploy AI and hope". Not "deploy AI and add headcount". Deploy AI into an environment where monitoring and ownership are handled and defined, and the tools start giving back time instead of taking it.

What to do about it

This week

QUICK WIN

 EFFORT

Answer the "not sure" question for your own tenant

In this study, 8% of leaders couldn't say whether AI had surfaced something it shouldn't have. Among the teams who hadn't finished a content cleanup, that figure was 20%, three times the rate among teams who had. "Not sure" is the most actionable answer in the whole dataset, because it's the one you can resolve without a project.

Pull a list of every AI agent and Copilot deployment currently active. For each one, note what data sources it connects to and who approved it. Flag anything unowned, unknown, or connected to sensitive sources. You're not remediating yet. You're establishing whether you can answer the question at all. Most teams find the exercise takes an afternoon and changes the conversation, because "we have full visibility" and "we can produce the list" turn out to be different claims.

AI governance survey · n=851

This quarter

MEDIUM LIFT

 EFFORT

Move to continuous monitoring, and name one accountable owner

These two go together: monitoring without ownership generates alerts nobody acts on, and ownership without monitoring is accountability for something you can't see.

On monitoring: 48% run continuous automated monitoring today. If you're in the 37% doing periodic reviews, you don't need to jump straight to full coverage. Start with the signals that change fastest: new agents appearing, changes to what an existing agent can reach, and sensitivity-label changes on content in scope for AI.

On ownership: if you already have a defined owner with consistent enforcement, this quarter's work is coverage, making sure enforcement is genuinely consistent across departments rather than strongest where the owner sits. If you're in the 26% where enforcement varies, or the 22% with shared ownership, the work is naming one person with the authority to say no.

AI governance survey · n=851

This year

STRATEGIC INVESTMENT

EFFORT

SHAREGATE PROTECT

Built for exactly this: tenant-scale visibility into permissions and sharing, automated policies that keep the environment clean, in-context remediation, and impact metrics you can hand to an auditor or a CFO. See it, fix it, keep it that way.

[Discover Protect →](#)

Close the loop between cleanup, visibility, and evidence

The cleanup data makes a specific argument: organization-wide cleanup didn't produce fewer self-reported incidents, it produced fewer *unknowns*. The teams that had done the work could answer the question. The teams that hadn't were three times more likely to say they didn't know.


That's the target state, and it's worth being precise about what it means. Not "no incidents", which no environment can promise. An environment where you would *know*.

Concretely: content and permissions cleaned at tenant scale rather than departmentally, workspace lifecycle automated so inactive content stops accumulating, continuous monitoring across every AI tool rather than just the Microsoft one, and a defined owner who can evidence all of it to an auditor.

The 10% whose governance workload went down after deploying AI are the proof this is achievable. Not as a checklist: most of them are missing at least one of the three habits this chapter walked through, and the survey never asked what order anyone did things in. What three quarters of them share is a posture, governance they describe as operationalized or better, against 63% of the study. That is the one difference that holds up once the individual habits are accounted for. Then re-measure, and treat the result as the evidence, rather than the confidence, that the first part of this chapter was about.

AI governance survey · n=851

CH 01 · END



02 The AI reckoning

CHAPTER 2 IN BRIEF

86% expect measurable ROI from their Microsoft 365 AI investments within 18 months. For 51%, what blocks the math is a cost they can't see.

4%

say AI has had no measurable impact on their IT budget

73%

of Copilot users are also paying for ChatGPT, Claude, or Gemini

2x

Teams who cleaned up estate-wide are twice as likely to be very confident of AI return

WHAT AI COST COVERS

Not just Copilot seats.

AI cost isn't just Copilot seats. It's the other AI tools running beside them, the storage and compute they generate, and the governance work needed to make any of it defensible.

WHY IT MATTERS MORE NOW

22% say AI hit a fifth of the IT budget.

The last chapter was about what AI can reach. This one is about what it costs to let it. AI licensing is now the most-cited cost challenge in Microsoft 365, named more often than storage growth. And 22% of organizations say AI has hit more than a fifth of their annual IT budget in twelve months.

WHAT CHANGED

July price rise, E7, agents to govern.

Microsoft raised prices in July and launched E7, folding Copilot into the enterprise licence. At the same time, a fifth of teams now name governing Copilot agents and plugins among their top concerns, a surface that barely existed when these budgets were set.

THE GAP

The bottleneck isn't strategy or budget.

Three quarters of organizations say they know exactly what they want AI to achieve. But half can't see what it costs. So the bottleneck isn't strategy and it isn't budget. It's that the environment underneath the invoice was never mapped, and you can't do arithmetic on something you can't see.

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 2 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

Governance maturity was the clearest predictor of confidence in an AI return anywhere in this research. The M365 Governance Index scores your tenant's governance, compares you to your peers, and gives data-backed advice on what to tackle first.

[Get your score →](#)

Adoption happened fast. Most teams are running Copilot and other AI tools in environments they still can't fully see. Read more about that in

[CH 1 Confidence ≠ control →](#).

The permissions sprawl that makes AI risky is the same sprawl that makes it expensive. And the governance work that contains the risk is the same work that makes the spend explainable. IT teams tend to experience these as separate conversations, one with security and one with finance. But the data suggests they're the same conversation with two audiences.

IT leaders are already finding that AI costs are taking up a meaningful share of the IT budget, and most teams can't yet trace them well enough to say whether they're working. That's not a Copilot problem. It's a visibility problem in a trench coat.

This chapter of [*The State of Microsoft 365*](#) looks at what AI is actually costing, why those costs are hard to pin down, how governance work changes investment confidence, and where AI spend sits next to the other things quietly draining the budget, like storage growth, unused licenses, duplicate tenants nobody consolidated after the acquisition.


Let's dig into the insights we collected from 1,700+ IT pros and expert advice from Microsoft MVPs to explain why these costs are so hard to trace and what to do about it for a quick win, a medium lift, and a strategic investment.

What AI is costing

AI costs show up in licensing, in storage, in compute, and in the quiet expense of creating content and running tools nobody is tracking. Start with the size of it.

What share of the IT budget AI is taking

WE ASKED Approximately what percentage of your total IT budget has been impacted by AI-related infrastructure and licensing costs (e.g., Copilot licensing, storage, compute) over the past 12 months?



The largest group puts AI infrastructure and licensing at 6-10% of total IT spend. But 22% are already past a fifth of the budget.

AI GOVERNANCE SURVEY · N=851

Asked what percentage of their total IT budget has been affected by AI infrastructure and licensing over the past 12 months, the answers cluster firmly in the middle of the range. 59% put it between 6% and 20%. Another 22% are already past a fifth of the annual IT budget. Only 4% report no measurable impact at all. For most teams this isn't a pilot-sized number. It's a line finance has noticed. AI governance survey · n=851

“ AI is currently coming out of the IT budget, and IT is still treated as a cost center rather than an innovation center. The business asks for AI. IT pays for it.”



DREW MADELUNG
MICROSOFT MVP

And the bill scales with how far into deployment you are.

What AI takes out of the IT budget, by how far Copilot has been deployed			
Copilot status	n	AI over 10% of budget	AI over 20% of budget
Fully deployed	477	59%	32%
Partially deployed	313	40%	9%
Not deployed	61	21%	12%
Sample average	851	49%	22%

AI GOVERNANCE SURVEY · N=851

Across all organizations, 22% say AI has taken more than a fifth of the budget. Among teams that have fully deployed Copilot, that rises to 32%. So if you are mid-deployment, your current costs are a preview rather than a total.

And if you haven't deployed at all, your floor still isn't zero. A fifth of that group reports AI taking more than a tenth of the budget, because readiness work, storage growth and other AI tools land on the same budget line whether or not Copilot has been switched on. Note the one place the table stops being tidy: at the higher threshold, teams with no Copilot deployment (12%) edge past teams that are partway through one (9%). The base is small and directional, but the direction is real. Not deploying isn't the same as not paying.

None of which means the answer is to deploy less. It means being deliberate about who gets a seat.

“ Not everyone in your organization needs AI. Find the people who will actually use it, give them real headroom, and build a feedback loop so you can see what came back.”




DREW MADELUNG
MICROSOFT MVP

AI licensing tops the cost list

The operations survey asked a different population about cost pressures across the whole Microsoft 365 estate. AI licensing didn't replace the familiar costs of storage, duplicate tenants and on-prem infrastructure. It passed them.

AI licensing is the new top cost pain

WE ASKED What are the biggest cost-related challenges in your Microsoft 365 environment? (choose up to two)




Copilot/AI licensing edged out storage and unused licenses.

IT OPERATIONS SURVEY · N=943

Storage bloat is a structural problem teams have managed for years. It accumulates quietly, it's nobody's fault in particular, and it has had years to become the thing everyone complains about. Copilot licensing is a line item that has existed for a year or two at most organizations, and it's already top of the list. On a question that only let teams name two things, 41% spent one of their two on it, ahead of storage, ahead of the tenants left over from the last acquisition, ahead of every license nobody is using. IT operations survey · n=943

Five hundred employees is where AI licensing bites



Respondents could name up to two cost challenges, so this is 41% of a two-item budget across the whole sample.

IT OPERATIONS SURVEY · N=943

Five hundred employees is where this stops being a line item. Below 250, 30% of teams name AI licensing as a top cost challenge. Through 499 it creeps to 35%. At 500 it jumps eight points, and then it just sits there. Everything above 1,000 lands within half a point of 45%. Whatever happens at that size happens fast, and it doesn't get worse.

Third place on the list is a different kind of answer. A quarter of teams say a lack of visibility into what they're spending versus what they're actually using is one of their top challenges.

That isn't a cost. It's not knowing what the costs are.

The AI cost visibility problem

Nobody we surveyed is confused about whether AI is costing them. Even the teams that haven't deployed Copilot have AI taking a chunk of the IT budget. The problem is that most teams can't connect the cost to specific outcomes.


Where the concern shows up in Copilot specifically

We asked IT leaders what worries them most about Copilot. Data quality and access management came back effectively tied at the top: 463 respondents against 461, out of 943.


Same top two worries, two years running

WE ASKED What are your biggest concerns about using Microsoft Copilot? (choose up to three)


Data quality, relevance, retention




Security and access




Lack AI governance expertise




Realizing Copilot's ROI



Cost control



Governing agents, plugins, automations




Not on the 2025 questionnaire – first year asked.

Governing AI agents is new to this list. Only ROI and cost control aren't governance problems.

IT OPERATIONS SURVEY · N=943 · 2026 n=943 · 2025 n=623


Why the 2026 numbers above sit low compared to 2025

Both years capped you at three picks. 2026 offered one more thing to pick from, so every number on the chart above sits a little low. More options, same three slots. Here is a quick visual of what we mean:



In other words the “Lack of AI governance expertise” shows 37% in our chart, but if we account for the extra option the more honest answer is a range: 37% at the low end, 42% at the high end. Up a point on last year, or up six. It went up either way. How far up, we can't say.

Every other concern on this chart fell, whichever way you count it, with one other exception.



Copilot ROI is the one that could have gone either direction, which you can see in the visual explainer here.

Microsoft MVP Susan Hanley hears the data-quality half of that tie from users, in different words.

“ Every user I talk to says search doesn't work. Search isn't the problem. Content management is. The old content never gets removed, so the right answer doesn't surface on the first page. Unfortunately, AI can't necessarily tell which content is correct and the outcome is that people don't trust the AI answer.”



SUSAN HANLEY
MICROSOFT MVP

Realizing the ROI of Copilot is a top concern for 34% of IT leaders, a clear six points above cost control. Teams are more worried about proving Copilot is worth it than about what it costs. A price is a number you can look up. A return is something you have to demonstrate, and demonstrating it means tracing spend to outcomes across an environment most teams can't fully see.

The two governance concerns at the top of this list are the subject of Chapter 1, and they're worth seeing again here [CH 1 Confidence ≠ control →](#). Because the thing teams say is hardest about proving AI's value isn't the pricing model. It's visibility, and visibility is a governance capability. IT operations survey · n=943

Cost control falling while AI licensing tops the Microsoft 365 cost list isn't a contradiction. One question asks what is expensive across Microsoft 365; the other asks what worries you about Copilot. A line item can be the biggest on your bill without being the thing that keeps you up.


Read the five slopes together and the shape is the point. Four fall. Only lack of internal expertise on AI governance rises, the one worry a year of Copilot deployment was supposed to fix.

Governing AI agents has no 2025 line above it because it wasn't on last year's list at all. In its first year of being asked, one in five named it.

What makes AI ROI hard to measure


Ask what makes AI return on investment hard to measure and the list splits in two. Cost visibility, governance complexity and unreliable data lead. Then a real gap, and below it the answers you would expect to top a question like this: evaluation

windows that are too short, goals that were never clear, benefits too diffuse to attribute.




That ordering is a finding in itself. The classic reasons ROI is hard to prove—fuzzy objectives, not enough time—are the ones teams named least. What they named instead were three versions of the same problem: they can't see clearly enough to do the arithmetic.

The overlap makes it sharper.



88% of teams named at least one of the three clarity problems. A third named nothing else at all—no complaint about goals, no complaint about timelines—which is three times the number who blamed goals and timelines on their own. Now set that against how the same population describes its own AI strategy. AI governance survey · n=851



Three in four say their AI business goals are well or extremely well defined. Of those same teams, more than half name cost visibility as one of the biggest things blocking AI ROI.

AI GOVERNANCE SURVEY · N=851

That's the sharpest juxtaposition in this chapter, and it's worth stating plainly. Three quarters of organizations know what they want AI to achieve. More than half of those same teams can't see what it costs. The bottleneck isn't strategic clarity. It's visibility.

Permissions, content sprawl, lifecycle, ownership—the reason the numbers won't reconcile is that the environment underneath them was never mapped. That's the uncomfortable part of this chapter. The cost question can't be answered without going back to the basics.

The teams watching most closely report the biggest visibility problem

We expected teams doing more monitoring to find cost visibility easier. They don't. The rate is the same across every group that monitors at all.


Naming cost visibility a top ROI blocker, by monitoring posture		
AI monitoring approach	n	Names cost visibility a top ROI blocker
Continuous automated monitoring	408	52%
Periodic monitoring	318	51%
Manual spot checks	77	52%
Only investigate after incidents	33	46%
No monitoring in place	15	33%
Sample average	851	51%

The two smallest groups are directional given their base sizes, but the direction is the opposite of the intuitive one. Teams with no monitoring at all are the least likely to say cost visibility is a problem.

This is the third time this pattern has appeared in the data. Chapter 1 found that concern about AI reaching unreviewed content rose with monitoring maturity, and that teams which had done cleanup reported more incidents because they could see them (CH 1 **Confidence ≠ control** →). Here the same shape shows up in cost. The likeliest explanation isn't that these teams' costs are clearer. It's that nobody has looked. Reporting a problem requires noticing it first, so if anything, the 51% headline understates how many organizations have a cost visibility problem, because the teams least equipped to detect one are also the least likely to report one.

What teams are doing about cost today

Knowing the bill is hard to read hasn't stopped teams from trying. We asked how IT pros have approached Microsoft 365 costs—not just AI costs—over the past twelve months.



Half of organizations are actively working the problem. But a third are watching the number rise without acting on it, and one in nine has seen costs increase and done nothing yet. The gap between monitoring and optimizing is the same gap that runs

through the governance chapter between periodic review and continuous control
CH 3 Governance →. Watching a number isn't the same as having a lever attached to it. IT operations survey · n=943

WHAT THE CONFIDENT 39% HAVE IN COMMON

Governance is the key to AI ROI

So far: AI is taking a real share of the budget, it tops the list of cost challenges in Microsoft 365, and proving ROI is difficult. So we asked how confident organizations are that a measurable return is coming in the next eighteen months, and whether AI-related governance work changes that confidence.

“ Good AI at work is AI that uses your data. If it isn't using your data, go use a free tool. Which means your data governance is your AI strategy.”




ERIC OVERFIELD
MICROSOFT MVP

Who is most confident in AI ROI

86% expect measurable ROI from Microsoft 365 AI initiatives within 18 months. Only 39% are very confident of it, and 13% are uncertain or not confident. Goals are clear and expectation is high. Conviction is what's thin, and what stands between expectation and conviction is, by the sample's own account, the ability to see the costs and manage the governance complexity.

Which raises the question worth chasing through the rest of this chapter. Who are the 39%, and what makes them different? The first answer runs against the intuitive assumption that bigger bills make people nervous.

The bigger the AI bill, the surer they are of the return



From 21% at the bottom of the range to 77% at the top, in a clean line. The teams under the most cost pressure are the most convinced it's worth it.

AI GOVERNANCE SURVEY · N=851

Confidence rises with spend, in a clean line, from 21% at the bottom to 77% at the top. The two extreme bands have small bases and should be read as directional, but the gradient across the well-populated middle is unambiguous. There are at least two readings and the data doesn't separate them. Organizations spending more may be seeing more, having pushed past pilot scale into workloads where value is measurable. Or conviction may drive the spending rather than follow it. Most likely both.


None of this means cost anxiety is misplaced. In the operations study AI licensing is the most-cited cost challenge in Microsoft 365, and cost control is a top Copilot concern for more than a quarter of teams. The pressure is real. What this gradient adds is that the teams under the most pressure are the most convinced it's worthwhile.

How governance work shapes AI investment confidence

Note what the question measures: how much, not which way.

Governance work moves the AI investment case

WE ASKED To what extent do AI-related governance activities (permission audits, cleanup, lifecycle management) impact AI investment confidence?



78% say permission audits, cleanup and lifecycle work shape their confidence to invest further in AI, a third of them significantly. The question asks how much, not which way.

AI GOVERNANCE SURVEY · N=851

Asked to what extent AI-related governance activities—permission audits, cleanup, lifecycle management—affect their confidence to invest further in AI, 78% said those activities have a significant or moderate impact. 32% said significant. Fewer than one in twenty say they make no difference. As beliefs go, that's close to consensus. AI governance survey · n=851


Governance isn't a tax on the AI budget. It's a precondition of the AI investment case. That reframes the internal argument. The usual framing puts governance and AI in competition for the same money, with governance cast as the responsible-but-boring option. The data says the opposite: the organizations that can prove they have some control over their tenants are the ones that can justify the next batch of AI investment, and the ones that can't are the ones whose AI programs stall at the business case.

Beliefs are cheap, though, and a survey is a comfortable place to hold one. The more useful question is whether the teams who actually did the work feel differently from the teams who didn't.

Three different measures of “did you do the work”. Same slope every time.

By self-described governance maturity


automated → reactive



AI governance survey · n=851

By how far the cleanup went


org-wide → none



AI governance survey · n=851

By cost posture (operations study)

optimizing → inactive



IT operations survey · n=943

The first two panels report the share very confident of a measurable AI return within 18 months. The third is a different study measuring a different outcome—the share running proactive governance monitoring—which is exactly why it's worth reading alongside. Two independent studies, two different framings, one finding: the organizations that can see and control their environment are the organizations that can see and control what it costs.

Confidence drops as you move down each ladder, but the steps aren't even. On maturity, the sharpest fall comes between governance that's consistently enforced and governance that exists as process without being operationalized, a 14-point drop, against barely a point between the two rungs below it. On cleanup the effect is stronger still: teams that cleaned up estate-wide before scaling AI are more than twice as likely to be very confident of a return as those that cleaned up only in places, and more than five times as likely as those that only have a plan.

Before anyone builds a business case on that gradient, be clear about what it proves. These are correlations. The data can't tell us whether cleanup produces returns, whether the organizations disciplined enough to run an estate-wide cleanup are also the ones disciplined enough to measure and realize returns, or some mixture. What it does establish is that governance maturity and confidence in AI value move together, consistently, across two independent measures. Nothing here supports treating them as competing priorities.

There is a reason only 51% of teams managed an org-wide cleanup. It's hard to do, especially with scripting or native tooling alone.

WHERE SHAREGATE FITS

ShareGate Protect surfaces stale permissions, oversharing and inactive workspaces across every team, site, group and OneDrive at once, not just the twelve you had time for. One place, no admin-centre hopping. And with ShareGate MCP you can connect Protect to Copilot, Claude or ChatGPT and clean up the tenant without leaving the chat.

[See how Protect works →](#)

“ The CFO conversation changes the moment you stop pitching governance as insurance and start showing it as the thing that unlocks the next phase of AI spend. Nobody wants to fund a cleanup project. Everybody wants to fund the reason the Copilot rollout can safely double. It is the same work and the same invoice. The difference is entirely in whether you can show what it buys.”




RICHARD HARBRIDGE
MICROSOFT MVP

The geography of AI spend and return

The budget picture isn't evenly distributed across countries, and it doesn't map neatly onto confidence.

AI budget share and ROI conviction, by country



The shaded corridor is one standard deviation either side of the fitted trend. Four countries sit inside it — their conviction is about what their spend predicts. Germany and Netherlands sit below it, and United States above.

AI GOVERNANCE SURVEY · N=851

Spend and conviction broadly travel together, which is the corridor. What the corridor is for is the three countries that leave it. The US isn't simply high on both. Its conviction runs further ahead than even its spend predicts. Germany goes the other way, and hardest: the second-highest AI budget share in the sample paired with one of the lowest confidence levels, 55% against 26%. The Netherlands sits in the same quadrant on a smaller base.

High investment, low certainty of return. The rest of the German data explains it.

← BELOW SAMPLE AVERAGE ABOVE →

Category	Survey Result (N=851)	Sample Average (n=104)
AI over 10% of IT budget	55%	49%
Cleanup completed org-wide	45%	51%
Governance highly automated	26%	37%
Continuous automated AI monitoring	36%	48%
Very confident of ROI	26%	39%

One bar above the line, four below. Spending more and having built less.

AI GOVERNANCE SURVEY · N=851 · n=104


Is this just national modesty? Some of it, probably. German respondents rate themselves lower on every confidence scale in the study, which is exactly what a conservative answering style looks like. Here is why it doesn't explain the gap. Three of the five measures above aren't opinions. Budget share, completed cleanup and monitoring approach are things a team either did or didn't do. Germany is spending more and has done less. No amount of modesty moves those.

Cost in the broader estate

ShareGate

Which of these can you actually move?

WE ASKED What are the biggest cost-related challenges in your Microsoft 365 environment? (choose up to two)



Cut the AI bill and you buy fewer seats, which means less AI. The four in the middle cost you time rather than capability. The last one isn't a cost at all. It's not knowing what the costs are.

IT OPERATIONS SURVEY · N=943

There is a practical difference between AI costs and the rest of this list, and it's worth being precise about. You can reduce your AI bill by buying fewer seats. But that means less AI, which is rarely what anyone wants six months into a rollout. Reducing storage, reclaiming unused licenses, consolidating duplicate tenants: those usually don't cost you capability. They cost you time and effort. That's a different kind of decision, and historically it's the one teams chose not to prioritize.

Storage growth is a content sprawl problem ([CH 4 The estate →](#)). Duplicate environments from M&A are a migration problem ([CH 6 Migration →](#)). Unused licenses are a lifecycle problem ([CH 3 Governance →](#)). They arrive as invoices, but they originate as governance debt, and each gets more expensive the longer it waits, because storage grows, renewals recur, and duplicate tenants accumulate their own sprawl.

There is an honest complication on that chart too. 16% of teams name the cost of third-party tools needed to fill gaps in native capability as one of their biggest cost challenges. Buying something to do this work is itself a cost, and any argument for it has to clear that bar rather than pretend it isn't there. What the calculation usually leaves out is the other side of the ledger.

“ Teams look at the systems and capabilities they already own and ask why they'd pay for another tool. The question they skip is what their own time going from system to system is worth.”



NOOREZ KHAMIS
MICROSOFT MVP

What's changed is the size of the payoff. AI creates more content in the tenant and then crawls everything it can reach, whether that reach is warranted or not, so the cleanup that used to buy back some storage now buys back control over what AI can see.

Same work, better return.


Cleanup, consolidation, and lifecycle automation aren't hygiene projects competing with strategic initiatives. They're now security issues and how you free the budget that can help offset your new AI invoices.

What tips a cleanup into a migration

At some point tidying the estate stops being enough and teams start changing its shape: consolidating tenants, retiring what is left on-prem, moving workloads somewhere cheaper to run. Cost is already driving those structural decisions.

88% say cost is steering the roadmap

WE ASKED To what extent are cost pressures (licensing, storage, duplicate tenants) driving your organization's migration or consolidation decisions?




A driver for almost everyone, the primary one for barely a third.

IT OPERATIONS SURVEY · N=943

88% say cost pressures are a moderate or major driver of their migration and consolidation plans, with 36% calling it a major driver. Only one team in fifty says it plays no part. But look at where the weight sits: half call it one consideration among several, and only about a third say it's a primary reason they're moving.

What matters here is the direction of causation. Teams aren't migrating and then discovering savings. They're migrating because the cost of not doing so became visible. IT operations survey · n=943

Only one cost problem actually changes the migration decision



Share calling cost a major driver of migration or consolidation.

IT OPERATIONS SURVEY · N=943

Among teams carrying duplicate environments or tenants from M&A activity, the share calling cost a major driver jumps to 44%, well clear of everyone else. No other cost problem on the list moves the number more than three points. That

makes sense once you look at what a duplicate tenant actually is. Storage growth and unused licenses are costs you can attack without changing anything structural. Two tenants running in parallel can't be optimized away. The only way to stop paying for the second one is to consolidate.

What to do about it

This week

QUICK WIN

 EFFORT

Get a single view of what AI is actually costing you

Before the next renewal conversation, not during it. Half of teams named cost visibility as one of the biggest things blocking AI ROI. That's half the room heading into a renewal without the number they need. Capture:

- Fixed license commitments, and assigned-versus-active use.
- Consumption-based agent, API, model, and connector costs.
- Third-party AI subscriptions.
- Supporting security, governance, and operational costs where material.
- The owner or cost centre for each spend category.
- Known gaps in attribution.

You aren't optimizing yet. You're establishing whether you can answer "what does this cost us" with a figure rather than an estimate.

This quarter

MEDIUM LIFT

 EFFORT

Reclaim and consolidate the three cleanup-shaped costs

- Unused or underused licenses. E3, E5, the newer E7 tier that bundles Copilot and Agent 365 outright, and whatever AI seats got bought on someone else's budget.
- Storage growth (39%).
- Duplicate M&A tenants (22%).

AI seats are typically the least audited, because they're the newest and because they're spread across vendors. Two thirds of organizations run or plan more than one AI tool, two named products each, on average. Among Copilot organizations, 73% also have ChatGPT Enterprise, Claude for Work, or Gemini for Workspace in use or on the roadmap. Some of those seats are already billing. Some are about to.

That overlap may well be deliberate; different tools genuinely suit different work, and nothing in this data says otherwise. But it's rarely audited, it's often bought outside IT, and with Copilot now bundled into a top-tier SKU rather than sold purely as an add-on, an idle seat is a bigger number than it used to be.

Start with license reclamation because it's the fastest to evidence, then scope tenant consolidation, which is slower but usually larger. Track what you recover. That number funds the strategic item below, and it's the most persuasive thing you can bring to a budget conversation.

This year

STRATEGIC INVESTMENT

EFFORT

SHAREGATE PROTECT

That's the job Protect was built for: continuous visibility into who can reach what, automated lifecycle so inactive workspaces stop quietly accruing storage and risk, and reporting you can hand to finance without translating it first.

[See how it works →](#)

Instrument governance so it produces evidence, not just cleaner tenants

Teams that ran an estate-wide cleanup before scaling AI are more than five times as likely to be very confident of a measurable return as teams that have only planned one: 56% against 10%. And 78% say governance activities shape their confidence to invest further in AI.

That's the argument to make internally, and it inverts the usual one: you aren't asking finance to fund governance *instead of* AI, you're asking them to fund the thing that makes the AI spend defensible.

Concretely, that means continuous visibility into what AI can reach, automated lifecycle so content stops accumulating cost and risk in parallel, and reporting that ties governance activity to spend avoided.

Every priority use case needs five things defined:

- Business outcome and baseline
- Adoption and usage
- Quality or effectiveness
- Risk or control indicators
- Total cost with a useful unit measure

The last one is where most business cases fall apart. "We spent \$400,000 on Copilot last year" isn't a number anyone can act on. "Copilot costs us \$18 per active user per month, against a support case that used to take 40 minutes and now takes 12" is a number that survives a budget review.

Pick a unit and hold it steady year over year. The absolute figure will always be argued with. The trend won't.

Risk and control indicators are the hardest of the five to produce, and they're the ones governance instrumentation hands you almost for free. You can't report exposure you can't see. By investing in it, the ROI conversation becomes less dependent on anecdotes and more grounded in a repeatable set of cost, adoption, outcome and risk measures.

CH 02 · END



03 Governance

CHAPTER 3 IN BRIEF

77% had a governance incident last year. Only 1% have a tool built to prevent one.

30%

find out about governance problems only when someone tells them, if at all

1

setting predicts incidents better than anything else we measured

2

changes to it take the most common setup from 82% incidents to 43%

WHAT GOVERNANCE COVERS

Access, ownership, apps, evidence.

Microsoft 365 governance is what stops a tenant from accumulating access nobody reviewed, content nobody owns, apps nobody approved, and gaps nobody can explain to an auditor.

WHY IT MATTERS MORE NOW

71% say AI grew their governance workload.

Governance has always been on the to-do list, but it was always easy to let it slip. AI changed that. [CH 1 Confidence ≠ control →](#)

In our AI governance survey, 79% are at least moderately concerned about AI reaching content nobody has reviewed the permissions on, and 71% say their governance workload has grown since they turned AI tools on.

WHAT HASN'T CHANGED

Still 1% using purpose-built tooling.

The number of people who use purpose-built tooling to govern their tenant stayed flat from 2025 to 2026. Over the same period, most teams in our IT operations survey had at least one governance incident.

THE GAP

The blocker is executive buy-in, not budget.

Asked what would most improve their governance, one of the top answers was executive buy-in to prioritize governance. Not more budget or expertise. So teams know it's important, but they can't convince their leadership to dedicate resources to it.

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 3 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

The M365 Governance Index tells you where you stand against your peers and what separates the 23% without an incident from the rest.

[Get your score →](#)

The last two chapters were about AI: what it can reach

[CH 1 Confidence ≠ control →](#),

and what it costs [CH 2 The AI reckoning →](#). This one is about the environment underneath both, and a finding that reframes them. The AI exposures in chapter 1 are one part of a broader set of governance problems.

A lot of teams already did this work. In the run-up to Copilot, permissions got reviewed, oversharing got audited, cleanup got funded and finished. Then the project closed, attention moved on, and the tenant kept changing every day afterwards.


That's the shape of the governance problem in 2026. Not neglect—completion. Governance is the one kind of IT work that degrades the moment you stop doing it, because it was never a state you arrive at. It's the ongoing job of keeping access, ownership, content and evidence from drifting away from whatever you last decided. An environment that was defensible eighteen months ago isn't defensible now, and nothing in it will tell you when that changed.

And unlike AI, none of it's new. These are the same controls, the same sprawl, and in one case the same statistic we published a year ago. What makes this chapter worth sitting with isn't that governance is hard. It's that everything changed except how teams handle it.

What actually breaks, and how teams find out

Governance failures don't announce themselves. Someone keeps access after they leave. A site nobody owns keeps collecting files. An auditor asks a question and the answer takes three weeks to assemble. That's what makes them easy to underrate, and it's why we asked about them directly rather than asking whether IT leaders felt exposed.

Chapter 1 framed its confidence paradox around one specific kind of failure: AI tools surfacing sensitive content to the wrong people. 29% of the 851 IT leaders in that study said it had happened to them [CH 1 Confidence ≠ control →](#). The IT operations survey asked a different 943 people something wider: not "did AI expose something", but "what governance events has your organization experienced across the estate in the past 12 months".



IT OPERATIONS SURVEY · N=943 · AI GOVERNANCE SURVEY · N=851

The two surveys aren't measuring the same thing, and the difference runs in a direction worth stating plainly. The AI question had no time limit. It asked whether this had ever happened, across however long a team has been running Copilot. The estate question was capped at twelve months. It still came back more than twice the size.

The like-for-like comparison is narrower and more useful. Oversharing on its own—sensitive content reaching people who shouldn't have it—affected 26% of organizations in a single year, against the 29% who have *ever* had AI surface something it shouldn't. Measured per year, the non-AI number is almost certainly the larger of the two.

That reframes chapter 1 rather than contradicting it. AI didn't introduce a category of risk that wasn't already there. It gave teams a faster and more articulate way to run into one that was.

“ Copilot is mostly exposing old governance problems rather than creating new ones. Oversharing isn't a new problem at all. SharePoint search has always been capable of returning content that a user technically had permission to see but arguably shouldn't have. What Copilot changes is how easy that content is to discover.”




LAURA ROGERS
MICROSOFT MVP

The most common governance incidents

Ask what actually went wrong and the answers are ordinary.

More than three in four hit an incident this year

WE ASKED In the past 12 months, has your organization experienced any of the following?



77% had at least one in the last 12 months. Stale access leads: a former employee or partner who kept access they shouldn't have. Audit gaps are close behind.

IT OPERATIONS SURVEY · N=943

The question offered a fifth answer that 41% picked: people couldn't find content, or Copilot returned stale results. It's excluded here. Nothing was exposed and nobody kept access they shouldn't have. It's what the other four produce, not a fifth kind of incident. Counted in, the headline would read 82%.

Two of the four are the same failure wearing different labels. Stale access is permission that outlived the person. Oversharing is permission that outlived its purpose. Both are what happens when access is granted once and never revisited, and they're the first and third most common incidents in the study.


Stale access leads, and it's the most mundane item on the list. A contractor finishes. An employee moves teams. Nothing is revoked, because revoking access is rarely anyone's named responsibility. It's a task that only exists if someone builds a process to generate it.

The audit and compliance gap in second place is a different kind of failure. Not that something was reachable, but that nobody could produce evidence about it when asked. That problem gets its own chapter ([CH 5 Compliance →](#)). Shadow IT and shadow AI round out the list, and they're the only entry that arrives from outside IT rather than accumulating quietly inside it.

None of this is dramatic. But all four share a property that makes them harder to stay ahead of than their simplicity suggests: none of them generates an alert. Nothing in a tenant announces that a former agency still has access to a site, or that a workspace lost its owner eight months ago. Which raises the question of how teams noticed any of it.

When the incident is the alert

Governance issues get discovered in a few different ways. Here's how it broke down across the 943 IT leaders we surveyed:



Just over a third (35%) run proactive monitoring with automated alerting. Another third (35%) rely on periodic scheduled reviews. 27% wait to be told: by a user, by an auditor, or by the incident itself. A final 3% have no consistent way of finding out at all, and we come back to them below.

Periodic review sits in an awkward middle. It's a real practice, and it beats waiting for a complaint, but a quarterly audit means a governance issue can exist for up to three months before anyone looks. The detection lag isn't a rounding error. It's the window in which stale access accumulates, an anonymous link circulates, or a Copilot agent quietly indexes a site nobody remembered.

Splitting the sample by discovery method reveals a pattern worth naming. Teams using proactive monitoring and automated alerting had 9 percentage points more clean-year responses than everyone else combined, and they reported fewer incidents on average.

Incident rate in the past 12 months, by how teams discover governance issues			
How teams discover governance issues	n	Had an incident	Clean year
Proactive monitoring and automated alerting	332	71%	29%
Periodic scheduled reviews	328	81%	19%
Reactive (user reports, audit findings, or after an incident)	256	83%	17%
No consistent discovery process	27	30%	70%
Sample average	943	77%	23%

IT OPERATIONS SURVEY · N=943

The bottom row is 27 organizations. Read it as directional, and as a measurement artifact rather than a result. A team with no discovery process isn't reporting that nothing happened; it's reporting that nothing reached them.

Monitoring adoption and lower incident counts move together. The data can't tell us whether monitoring causes fewer incidents, whether organizations that invest in monitoring also invest in the broader governance practices that reduce incidents independently, or some combination of both. The correlation is real. The causal story would need more than this study can give.

There's also a reading that runs the other way, and it's the one chapter 1 already surfaced: teams that monitor more are teams that *know* more

CH 1 Confidence ≠ control →. Some of the 9-point gap may be detection rather than prevention.

The bottom row of the table makes that case on its own. The small group with no consistent way of discovering governance issues reported the fewest incidents of anyone, by a wide margin: 70% said nothing went wrong all year, against 23% across the sample. That isn't a safety record. It's a missing instrument. A team without a discovery process isn't telling you nothing happened—it's telling you nothing reached them.

Because those 27 teams sit inside the "everyone else" column, the 9-point gap is if anything understated. Remove them and it widens to 11 points. Both readings point at the same operational conclusion: the teams that are looking are the teams who can answer the question at all.

“ The reason organizations are flying partly blind is that they don't have the people. It's not indifference. There's no human capital left to do more than they're already doing.”



JOANNE KLEIN
MICROSOFT MVP

THE NUMBER THAT DIDN'T MOVE

What teams are governing with

Ask what an IT team still does by hand in a Microsoft 365 tenant and the list gets shorter every year. Backups run on a schedule. Endpoint protection runs as a service. Patches deploy on a cycle. Each of those was somebody's manual job once, and none of them wait for a person to remember anymore.

Governance never made that transition. It's the one operational discipline in Microsoft 365 that most teams still run by hand, from reports someone has to think to pull, on a cadence someone has to maintain. And it isn't for lack of options. Software built to do this has been on the market for roughly as long as SharePoint has existed. Which is a strange result for a platform this mature.

“ Microsoft 365 is the best toolset there has ever been for working together. That's exactly the problem. Teams automated the collaboration and left the controls manual.”



DAVID DREVER
MICROSOFT MVP

Given 77% of organizations had a governance incident in the past year, you'd expect IT leaders to be shopping. So far, they aren't.


A year on, purpose-built governance tooling is still 1%

WE ASKED How is your organization currently handling governance and security in Microsoft 365?


Built-in Microsoft tools




Manual / internal policies



Purpose-built tool



Not addressing it



More teams are addressing governance than last year, but with the same built-in tools and manual policies. The purpose-built share hasn't moved.

IT OPERATIONS SURVEY · N=943 · 2026 n=943 · 2025 n=650

96% of organizations are governing Microsoft 365 with built-in Microsoft tools, manual internal policies, or some mix of the two. Another 4% aren't addressing governance and security at all. That leaves six respondents out of 943—1%, rounded the same way the 2025 report rounded it—using something purpose-built. Six people is too small a group to profile, cross-tabulate, or draw any conclusion from beyond its size.

We asked the same question of a different sample last year and got the same answer. Last year we called it the 1% gap. A year later, after a year in which full Copilot deployment went from 29% to 56% AI governance survey · n=851, it's still 1%. The gap didn't close while the surface being governed got substantially larger and substantially more reachable.

What did move is the split above it. Reliance on built-in Microsoft tools grew six points, manual policies fell four, and the share not addressing governance at all fell from 5% to 4%, the only movement pointing anywhere good, and a small one. The year's activity was lateral. A few teams stopped ignoring the problem, and a larger group consolidated onto tooling they already owned.

Nothing here suggests indifference. It suggests a ceiling.

To be fair to the built-in tools

Built-in tools aren't bad tools. They're capable within the scenarios they were designed for. The challenge is that many were built to administer individual services, not to continuously govern a Microsoft 365 environment at the pace sprawl now accumulates. Microsoft's native capabilities can support continuous controls, but coverage, configuration, licensing and cross-tool visibility vary.

The real test is whether the combined control set gives your team the coverage, context and response time it needs. Which is where the discovery data comes back: finding a setting when you know where to look isn't the same as being alerted when it changes, or having a workspace with an anonymous sharing link left active since 2023 brought to your attention.

“Microsoft's tools are a toolkit. Purpose-built tools are a solution. A toolkit hands you options and expects you to build the thing you actually needed. Purpose-built gives you what you need without having to build it.”



DREW MADELUNG
MICROSOFT MVP

“SharePoint Advanced Management got genuinely decent, and it arrives with Copilot licensing. A lot of teams are starting there to see whether it's enough. That's a reasonable first move. It isn't a finished one.”



JASPER OOSTERVELD
MICROSOFT MVP

What's blocking governance improvement

The incidents are real, the workload is growing, and software built for this has been on the market for decades. So why hasn't the number moved? If a team knows it has a governance problem and hasn't bought anything to fix it, the usual assumptions are budget or expertise. We asked what would have the biggest impact on improving a team's Microsoft 365 governance and security, and unlike most questions in this survey, allowed one answer only.

The ask isn't money. It's a mandate.

WE ASKED What would have the biggest impact on improving your team's Microsoft 365 governance and security?



Executive buy-in outranks budget seven to one. Budget finishes last.

IT OPERATIONS SURVEY · N=943

Chapter 1 covered the top answer, better controls for AI agents and automations **CH 1 Confidence ≠ control →**. Set that aside and the rest of the ranking is a finding in itself. The largest remaining answer was executive buy-in to prioritize governance, at 20%. More budget for tooling or headcount finished last at 3%, with more internal expertise just above it at 6%. Everything in between—automated remediation, better tenant visibility, delegated reviews—is a request for capability rather than cash.

So the thing teams most want isn't money, and it isn't knowledge. It's for this to matter to someone above them.

That's easy to dismiss as managers wanting their bosses to care more. But the data doesn't support the dismissal. 97% of respondents are a manager, director or VP, and the demand for executive buy-in climbs with seniority rather than falling.

What would have the biggest impact, by seniority			
What would have the biggest impact	Manager (459)	Director (363)	VP (94)
Better controls for AI agents, plugins, automations	36%	35%	29%
Executive buy-in to prioritize governance	17%	21%	33%
Automated remediation	20%	15%	20%
Better tenant visibility	9%	11%	6%
Delegated reviews to end users	8%	10%	9%
More internal expertise	7%	6%	2%
More budget for tooling or headcount	3%	3%	1%

IT OPERATIONS SURVEY · N=943 · n=916, excluding 27 respondents who selected "Other". The VP column is 94 respondents.

If this were just people passing blame upward, managers should name executive buy-in most often, since they have the most layers to point at. In this group they name it least. It's the VPs who name it almost twice as often.

The likeliest explanation is that governance loses on merit, every cycle, and the people who watch it lose know why. It has no deadline. It generates no revenue. Nothing fails visibly until something fails badly. Against an AI program with a slide deck, projected profits and a launch date, that's not a fair fight. That's why the ask isn't for money. Money wouldn't change the shape of the argument. A mandate would.

But waiting isn't a neutral position. Governance doesn't hold the level you left it at.

“Governance is not a one-time thing. Say you get to ninety percent and you're happy with it. If you stop working at it, that ninety becomes sixty or seventy. What you are looking at is a lagging indicator of how ready you used to be.”



MIKE MAADARANI
MICROSOFT MVP

“ Two years ago I could not get away from Copilot readiness and oversharing reviews. That was all we did for months. Now those are less frequent, because teams believe they finished. Readiness at a point in time is not readiness.”



DREW MADELUNG
MICROSOFT MVP


That changes what the 1% tooling number is measuring. A team that ran a cleanup two years ago and hasn't been back isn't sitting where they finished. They're somewhere below it, and nothing in the environment reports the drift. The number that stayed still was a number about purchasing. The thing it was meant to protect has been moving the whole time.

How teams rate their own governance

This chapter has been running on the IT operations survey. It's worth crossing to the AI governance study for a few questions, because it asked 851 IT leaders to grade their own governance, and asked separately about the specific controls they run, which means the summary and the specifics can be set side by side within the same sample.

AI governance maturity, self-rated


WE ASKED Which statement best describes your organization's current Microsoft 365 governance maturity?



Only 37% call themselves highly automated and continuously monitored. Over a third describe governance that is not operational at all.

AI GOVERNANCE SURVEY · N=851

63% describe themselves as operationalized or better. 37% picked the top option outright. Only 2% call themselves reactive.




Now hold that against what the same 851 people said elsewhere. 42% have fully centralized visibility into external sharing. 29% manage external sharing permissions through automated monitoring and policy enforcement. Both are well short of 63%.

The easy conclusion is overconfidence, like chapter 1. But look at what the five maturity options actually describe: highly automated, consistently enforced, mostly manual. Every one is a statement about how the process *runs*. None asks how much of the estate it runs across, a limit of the question we wrote, and also the harder thing to answer. A team can know precisely how disciplined its process is and still not know what sits outside it.

So we took the 318 respondents who picked the top rating and looked at what they'd told us elsewhere. Every difference runs in the direction you'd expect: this group reviews access more often, automates more of its sharing enforcement, and sees more of its estate than the sample average. The rating is sorting people correctly.

84% claim monthly or quarterly access reviews

WE ASKED How frequently does your organization conduct formal access reviews across Microsoft 365 (SharePoint, Teams, OneDrive)?



A remarkably disciplined-looking picture: only 2% review less than once a year. Worth pressure-testing against what a review actually covers.

AI GOVERNANCE SURVEY · N=851

Visibility is where the gap is biggest. 57% of the self-rated most mature have a fully centralized view of external sharing activity, against 42% across the sample, which still leaves 43% of the most confident band without a complete view of who their tenant is sharing with. In the other two cross-tabs the margins are narrow: eight points ahead on monthly access reviews, effectively identical on quarterly, five points ahead on managing external sharing permissions. A majority of them, 51%, describe that work as a combination of automated tools and manual review rather than automated enforcement.

Count how many of the three each of them actually holds—fully centralized visibility, monthly access reviews, automated sharing enforcement—and the band splits almost in half. 8% hold all three. 36% hold two. 34% hold one. 22% hold none.

For the 44% holding two or three, the earlier finding stands: they have a real process, and what separates them from everyone else is mostly that they can see more of what it covers. The rating is accurate and the gap is a perimeter.

“ The confidence isn't fake. It's uninformed. The organizations that built the information architecture and the security architecture underneath it years ago are in a much better position today. The ones who never learned what information architecture really is and built their IA without any deeper thought processes are likely confident about something that likely won't work.”



NOOREZ KHAMIS
MICROSOFT MVP

The other 56% hold one or none. That isn't a narrower perimeter, it's a thinner process. Set the bar at quarterly reviews instead of monthly and the split moves to 65/35, but the shape doesn't change: a substantial share of the most confident band is using the language of automation without much of it underneath. Which is worth naming plainly, because the two groups need different things. A team with a strong process and limited reach needs visibility. A team describing automation it doesn't have needs the automation.

The three-capability split and the top-band cross-tabs above were derived from the respondent-level survey data rather than read off a published summary. AI governance survey · n=851

External sharing: links, guests, and the workspaces nobody owns


Stale access and oversharing were two of the four most common incidents. This section is about where they come from. Three controls account for most of it: who can create anonymous links, whether guest accounts get reviewed, and what happens to a workspace once nobody needs it. None of the three fails suddenly. All three accumulate quietly.

Anonymous sharing links

One tenant setting decides how far a shared link can travel.

A quarter leave the door open

WE ASKED How does your organization handle anonymous sharing links in Microsoft 365?



25% allow anonymous links tenant-wide with no expiration.

IT OPERATIONS SURVEY · N=943

70% of organizations allow anonymous links across the whole tenant. A quarter allow them with nothing set to make them expire. Every one of those links is a door that stays open until somebody actively closes it, and by design nobody has to be signed in to walk through.

Cross that against the incident data and the gap is one of the widest in the study.

Incident rate in the past 12 months, by anonymous sharing link policy			
Anonymous sharing link policy	n	Had an incident	Clean year
Blocked tenant-wide	74	43%	57%
Restricted to specific sites only	204	64%	36%
Allowed tenant-wide, with expiry	428	82%	18%
Allowed tenant-wide, no expiry	237	89%	11%
Sample average	943	77%	23%

IT OPERATIONS SURVEY · N=943

The blocked band is 74 organizations, a hair under the n<75 mark, so read it as a direction rather than an exact figure. What holds either way is the ordering: every step toward more open links comes with a higher incident rate.

The gradient runs in one direction across all four bands: 43% of organizations that block anonymous links reported an incident last year, against 89% of those that allow them tenant-wide with no expiry. Stale access is where the ratio is largest: it runs from 9% at the blocked end to 52% at the open end, close to a sixfold difference, where the overall incident rate roughly doubles. That ordering makes

mechanical sense as well as statistical sense. An anonymous link with no expiration isn't a route to stale access. It is stale access—a credential that works indefinitely, held by someone who never had to sign in.

The more useful detail is where in the gradient the jumps happen, because they aren't the same size. Blocked to site-restricted costs 21 points. Site-restricted to tenant-wide costs 18. Dropping expiry on top of that costs 7. The first two are about where links can be created at all; the third is about how long they survive once made.

We split the sample 18 ways to test this: by organization size, by how teams discover problems, by how much self-service they allow, and by country. The step from site-restricted to tenant-wide pointed the same way in every slice where both bases were large enough to report: 17 of 18, with France and Japan having no reportable site-restricted base either way. The expiry step didn't hold up the same way. It reversed in two slices, which is roughly what you'd expect from chance when the underlying effect is only 7 points wide.

So restricting *where* links can be created is associated with lower incident rates in every population we tested. Setting an expiry is associated with lower rates on average, but the effect is small enough that we can't find it reliably in smaller groups. That ordering is the reverse of what most organizations have actually done. The single most common configuration in the data—45% of respondents, more than any other—is links allowed tenant-wide with an expiry policy attached. It's the responsible-looking version of the permissive option: free, one switch, and it only ever applies going forward. On these numbers it's the smaller half of the fix.

What this doesn't prove

None of this makes the relationship causal. A tenant setting isn't randomly assigned, and organizations that restrict anonymous links differ from those that don't in ways this survey doesn't capture.

One counter-explanation is worth walking through, because it runs in the finding's favour rather than against it. Suppose the causation is backwards. Teams that suffered an oversharing incident responded by locking their link policy down. Those teams would then appear in the blocked and site-restricted bands *with an incident on record*, raising the rate in precisely the bands where it's lowest. The gradient would look shallower than the truth, not steeper. We can't test that directly; nothing in either survey asks when a policy was set versus when an incident happened. But it means the 46-point spread is more likely to understate the relationship than to exaggerate it.

What the data can't settle is whether blocking outright is the right endpoint.

“ A full block tends to push users toward email attachments, personal drives, or shadow IT, which relocates the risk rather than removing it.”



RICHARD HARBRIDGE
MICROSOFT MVP


Only 8% block tenant-wide, and this study has no way to test whether the other 92% are making a mistake or a trade-off.

Guest access reviews

Anonymous links let someone in without an account. Guest accounts are the opposite: a named identity, created on purpose, with a person attached to it. Which ought to make them easier to keep track of. 38% of organizations had a former employee or external partner keep access longer than they should have, and guest accounts are the mechanism for the second half of that.

Two-thirds review guests on a schedule

WE ASKED Does your organization have a formal process for reviewing and removing guest accounts?



Data

Yes, at least quarterly	65%
Ad hoc only	28%
No process	5%
Not sure	2%

65% review and remove guest accounts at least quarterly. 38% still had someone keep access they should have lost. The schedule exists. The outcome does not follow.

IT OPERATIONS SURVEY · N=943

Two thirds of organizations review guest access on a schedule, which is genuinely better than the folklore suggests. But what it buys is smaller than you'd hope. Organizations with a regular review reported stale access 38% of the time. Organizations reviewing ad hoc reported it 44% of the time. A six-point difference.

A third either review ad hoc or not at all, and the two smallest groups are the familiar ones. The 43 organizations with no process report stale access at 23%, and the 24 who don't know report it at 4%. That's the second and third time in this chapter that the teams doing least have reported least, and it means what it meant the first time. Both bases are far below the $n < 75$ threshold, and both are far more likely to be a visibility artifact than a safety record.

Split the review question by link policy, though, and the difference between regular and ad hoc reviewers gets genuinely useful. Where anonymous links are confined to specific sites, review cadence makes no measurable difference: 20% against 21%. Where links are open across the tenant, it's worth 10 to 12 points: 41% against 51% with an expiry set, 51% against 63% without one.

So the process works in proportion to what the configuration lets through. That's useful in both directions. If your link policy is tight, a quarterly guest review is probably enough. If it isn't, no review cadence available in this data gets you below 51%.

Bases for the split above: 122/62, 269/138, 174/51. The blocked-tenant band had too few ad hoc reviewers to compare, and the 62 and 51 cells are directional. IT operations survey · $n=943$

Links and guests are the two controls in this chapter that don't need a budget line, an owner, or a mandate. They need a decision.

What to do about it

This week

QUICK WIN

— — — EFFORT

Set anonymous links to expire, then find out which sites still allow them

A quarter of organizations have no expiry at all, so start there. Expiry is a single tenant-level setting and it costs nothing to change. The one caveat is that it only applies to links created after you set it, so anything already circulating needs finding separately.

Require Anyone links to expire, review whether they should stay available tenant-wide, then apply stricter defaults to sensitive sites. Before shortening the maximum duration at tenant level, assess the impact on active links and communicate the change to affected owners. Fourteen days is a common starting point. Prioritize existing links for review using sensitivity, age, activity and business ownership, not age alone. Done well you end the week with a significant risk reduced and the backlog visible.

IT operations survey · n=943

This quarter

MEDIUM LIFT

— — — EFFORT

Narrow where links can be created, give every guest an owner, move one control to continuous

The quick win produced two lists: sites with Anyone sharing switched on, and the links already out there. This is where you decide what to do with them. Decide which sites keep Anyone sharing, change the tenant-wide setting, and handle those sites as the exception. Tell the people whose workflow you're changing before you change it.

Guest accounts are next, and the problem isn't the review. It's that IT pros are usually the wrong people to make the decision. Somebody in marketing invited an agency contact two years ago; IT has no way of knowing whether that agency still works here, so the account survives another cycle. Give every guest a named sponsor, usually whoever invited them, make the sponsor re-attest on a fixed cadence, and define what happens when a sponsor leaves or doesn't respond. It converts an unowned admin chore into an owned decision.

The third move is the small one. A third of teams find out about problems through periodic review, which can leave something sitting in the tenant for three months. You don't have to fix that everywhere. Pick the highest-risk signal you have—external sharing changes is usually the right first choice—and put automated alerting on it rather than waiting for the next cycle. Record where you started and where you finished.

IT operations survey · n=943

This year

STRATEGIC INVESTMENT

EFFORT

Decide what your governance is actually made of, and get someone above you to agree to it

Work out which combination of native capabilities, automation, internal processes and purpose-built tools meets your requirements. Run a capability assessment against whether your current set-up actually covers:


- Current, centralized visibility.
- Risk-based policy enforcement.
- Alerting tied to response procedures.
- Automated remediation for low-risk, reversible actions.
- Human approval for destructive or high-impact actions.
- External-access lifecycle.
- Evidence retention and reporting.

96% of organizations are running governance on built-in tools or scripting and 1% on anything purpose-built, in a year when 77% had an incident. Whatever the right answer is for your environment, the current distribution isn't a considered choice. It's an inherited default.

This is also the move that isn't yours alone. Asked what would most improve their governance, the largest answer after AI agent controls was executive buy-in. So what you bring upstairs isn't a spending case. It's the before-and-after from last quarter. Then re-measure. The point isn't to complete a cleanup, it's to reach a state where you'd notice the next problem without anyone having to report it.

IT operations survey · n=943

CH 03 · END



04 The estate

CHAPTER 4 IN BRIEF

28% of organizations are fully in the cloud. 31% deliberately chose an on-premises future.

37%

say technical complexity is what's blocking their SharePoint migration, not budget

65%

have no automatic way to remove a workspace once it stops being used

2 pts

separate the most locked-down tenants from the most open on shadow IT

WHAT THE ESTATE COVERS

Cloud, on-premises, and what accumulates in both.

The estate is what you're actually governing: the data in your cloud tenant, whatever is still running on-premises, and the workspaces that accumulate in both.

WHY IT MATTERS MORE NOW

68% have to apply every control twice.

Every control in the last chapter gets applied to a surface [CH 3 Governance →](#). This chapter is about the surface. 68% of organizations keep data in both the cloud and on-premises, so most controls have to be applied twice. And 65% have no automatic process for removing a workspace once it falls out of use, which means the surface those controls cover keeps getting larger.

WHAT CHANGED

Nearly a third with no migration plan, down to 4%.

Microsoft ended support for SharePoint 2016 and 2019 in July 2026, and 73% of organizations were off legacy SharePoint before the date. A year ago 29% had no firm migration plan. Today it is 4%. Full cloud adoption moved from 22% to 28%.

THE GAP

The work that moved was the work with a date.

A date did in twelve months what years of migration guidance had not. Nothing else in the estate comes with one. No support notice arrives for permission sprawl, for a workspace whose owner left, or for an exception nobody revisited, so a team that wants that work done has to set the date itself.

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 4 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

You've just read where 943 IT teams sit. Answer a short set of questions and the M365 Governance Index scores you on your estate and your AI governance, then places you against the teams in this study and against your own industry, country and org size.

[Get your score →](#)

Everything in this report happens somewhere. This chapter is about that somewhere.

The governance gaps in the previous chapter did not appear from nowhere

CH 3 Governance →

. They accumulated in an environment that is still half-built, and the shape of that environment explains a good deal about why they persist. It's tempting to treat the estate as settled background—the thing the interesting problems happen on top of. The data says otherwise. The Microsoft 365 estate most teams are governing in 2026 is still partly on-premises, still accumulating content faster than anyone retires it, and broader in composition than the file-and-mail picture most planning documents assume.

An estate like that changes for one of two reasons. Someone inside decides to change it, or someone outside sets a date. This year we got a good look at the second kind.

Microsoft ended support for SharePoint 2016 and 2019 in July 2026, which handed the industry something it rarely gets: a hard date attached to a piece of infrastructure work. Teams responded to it. What they didn't all do was respond by moving to Microsoft 365. Leaving unsupported software and arriving in the cloud turn out to be two different projects, and a good part of this chapter is about the space between them.


Nothing else in the estate came with a date. How much employees can create without asking, what happens to a workspace after the project that needed it ends, how much storage accumulates along the way: none of it generates a support notice, and none of it fails visibly enough to force a decision. That's the asymmetry worth carrying through the rest of this chapter. The estate got exactly one deadline, and it only ever covered one part of it.

Fully cloud, fully on-premises, and the 68% in between


Hybrid can be a real architecture choice, especially where it's set up for hybrid search and taxonomy. But sometimes it's just what ends up happening when an organization starts working in the cloud while holding onto its on-premises

environment. Some of these teams are actively working toward being 100% in the cloud. Others have decided to keep maintaining two environments.

Let's start with the question everyone assumes is settled.



Only 28% of organizations are fully in the cloud. 3% are fully on-premises. Everyone else is somewhere in between.



Cloud-complete has moved from 22% to 28% year on year—real progress, and slower than the industry narrative suggests. At that rate the majority of organizations will be running hybrid estates for years, not months.

Of the roughly seven in ten that still run on-premises infrastructure, half describe what's left as a small amount. The rest split their data evenly or have more on-premises. Hybrid covers all of those, which is part of why the word doesn't tell you much on its own.

What a split estate costs you

The number matters because of what sits downstream of it. A hybrid estate means two permission models, two audit trails, two backup and retention regimes, and two places a piece of sensitive content can live. It also sets a ceiling on AI: Copilot reaches cloud content CH 1 Confidence ≠ control →. On-premises content is outside native Microsoft 365 grounding unless it's brought into scope through a configured connector. For many organizations that means what's on a file server is invisible to Copilot—sometimes a feature, sometimes a findability problem, and teams rarely have a clear view of which.

The self-perception gap is the interesting part. Most of these organizations don't describe themselves as hybrid. They describe themselves as cloud organizations with some legacy left over. That framing is comfortable, and it quietly removes the leftover from the governance plan.


“ Almost nobody says "we're a hybrid organization" anymore. They say they've moved to the cloud, and then mention the file server, the SharePoint farm, and the line-of-business app almost as an aside. The problem isn't the on-premises footprint. It's that once you've decided you're a cloud shop, the leftovers stop appearing in the architecture diagram, and things that aren't on the diagram don't get governed.”



RICHARD HARBRIDGE
MICROSOFT MVP

Cloud-complete respondents self-reported more governance problems

A file share doesn't generate an oversharing alert. Neither does a SharePoint farm, and neither shows up in the Microsoft 365 unified audit log. Which is worth knowing before looking at how governance incidents break down by cloud footprint, because on paper the result reads backwards.



Both charts are percentages of all 943 respondents in each band, so the two gradients are directly comparable.

Sort the sample by how much of the estate sits in the cloud. The teams with the most data in the cloud report the most incidents. They can also see the most of their own tenant. Two gradients, same direction, same four groups.

There are two readings.

The first is surface area. A fully cloud estate is a bigger, more connected, more shareable thing—more sharing links, more guests, more cross-tenant collaboration, more of everything that produces the incidents on the governance chapter's list

CH 3 **Governance** →

The second is detection, and it's the pattern that recurs throughout both studies. Cloud-complete organizations can see more. The file share we opened this section with may be producing exactly the same problems, unobserved and uncounted.

Two of the four move with the footprint

Split the incident types apart and they stop behaving alike.

Which incidents, by cloud footprint					
Cloud footprint	n	Oversharing	Stale access	Audit or compliance gap	Shadow IT or shadow AI
Completely migrated to cloud	268	32%	44%	41%	25%
Small amount on-premises	335	26%	38%	32%	24%
Roughly equal	240	23%	35%	34%	24%
Majority or all on-premises	100	13%	28%	34%	27%
Sample average	943	26%	38%	35%	25%

IT OPERATIONS SURVEY · N=943

Each column is the share of that band selecting the incident type. Respondents could select more than one, so the columns don't sum.

Two of the four types move with the footprint and two don't. Oversharing runs 32% down to 13%, stale access 44% to 28%. Audit gaps and shadow adoption sit flat across every band.

Which is the tell. The two that move are the two a connected tenant produces more of. The other two happen everywhere.

So how much of it is just better eyesight?

Most of it is real.

Hold visibility steady. Compare only the 435 teams who can see exactly what their agents reach, and the gap narrows from 19 points to 11. Better eyesight explains about two fifths of it. The rest is surface area. A cloud-complete estate isn't just better at spotting problems. It has more of them.

Which settles the migration question

Completing a cloud migration does not reduce your governance workload. It changes it, and usually enlarges it. Teams that treat migration as the project and governance as the phase that follows have the sequence backwards.

So what does help?

Which raises the more useful question. If where your data sits works against you, what works for you?

Share of each group that had a governance incident, by whether they watch the tenant				
Cloud footprint	n	Watches the tenant	Everyone else	Points lower
Completely migrated to cloud	107	80%	88%	8 pts
Small amount on-premises	115	68%	80%	12 pts
Roughly equal	76	67%	76%	9 pts
Majority or all on-premises	34	59%	70%	11 pts
Sample average	332	71%	80%	9 pts

IT OPERATIONS SURVEY · N=943

"Watches the tenant" is the 332 teams who find governance issues through monitoring and automated alerting. Everyone else is the remaining 611. The n column counts the watchers in each band.

One answer turns up in every group. Teams who watch their tenant, with monitoring and automated alerts, had fewer incidents than everyone else. Nine points fewer, whether they are fully in the cloud or barely started.


And they are not simply missing things. They can see about twice as much of their estate as everyone else can. Looking harder normally turns up more problems. It turned up fewer.

What the estate is actually made of

The composition question is worth asking separately from the location question, because assumptions about it are usually a decade out of date. When organizations move a tenant, here's what has to come with it.

Tenant-to-tenant touches everything

WE ASKED What workloads were (or will be) included in this tenant-to-tenant migration?



The 'mailbox move' is now a whole-estate move.

IT OPERATIONS SURVEY · N=943 · the 310 running a tenant-to-tenant move

Power Platform tops the list, ahead of SharePoint, Teams, and Exchange. That ordering should reset a few assumptions. The mental model of a Microsoft 365 estate as mail plus documents is obsolete. What teams are actually governing includes a substantial layer of apps, flows, and automations that were built by people outside IT, hold their own permissions and connections, and are now the most commonly in-scope workload when an estate changes hands. It's also the layer least likely to have an owner who still works there.


The migration mechanics are covered in their own chapter [CH 6 Migration →](#). The point here is what the list says about the surface area: the estate is broader than the org chart of the people who think they own it.

WHAT A DATE DOES TO WORK THAT WOULD OTHERWISE DRIFT

The SharePoint end-of-support moment moved almost everyone

Microsoft ended support for SharePoint 2016 and 2019 on July 14, 2026. Nothing switched off. It wasn't a power cut. A SharePoint 2019 environment runs today exactly as it did in 2025, except now there are no more security updates, no more bug fixes, and no more support calls.

Support ended on a Tuesday. Only 21% still had SharePoint 2016 or 2019 running on the Wednesday.



So 73% were off legacy SharePoint before the date. The 21% still carrying it split into 17% mid-migration and 4% with no plan at all.



Versus last year, comparing like with like: in 2025, 66% of organizations were still actively running SharePoint 2016 or 2019, and 29% had not yet formed a firm migration plan. Today, roughly 21% still have those versions in play (17% mid-migration, 4% with no active plan), and the share with no plan at all has fallen to under 4%. That second pairing is the cleanest measure of what changed: nearly a third of organizations had no plan a year ago; now almost everyone does.

That's a twenty-five point move in twelve months, and it's worth being honest about what caused it. Not advocacy. Not architectural conviction. A date. Microsoft's July 2026 end-of-support deadline achieved in twelve months what years of migration guidance did not.

There's a lesson in that worth carrying into every other section of this report. The governance work in the previous chapter and the AI visibility work in chapter 1 have no deadline attached. Nobody sends an end-of-support notice for permission sprawl. Teams that want that work to happen have to manufacture the forcing function themselves, because the environment will not supply one

[CH 7 The operating model →](#).

The third of the market that chose to stay

One number in that chart deserves more attention than it usually gets. 31% did not move to the cloud. They moved to SharePoint Server Subscription Edition—which is to say they deliberately chose a supported on-premises future.

That is not procrastination, it's a strategy, and it's nearly as common as completing the cloud migration. For those organizations the driver is usually a compliance or sovereignty requirement they cannot satisfy in a shared cloud, and the decision is often correct [CH 5 Compliance →](#). But it has a governance consequence that tends to go unplanned: they've committed to running two governance models indefinitely rather than transitionally. Everything in the governance chapter about continuous monitoring and lifecycle now has to work in two places, permanently, and native cloud tooling covers only one of them.

The teams who moved to Subscription Edition are more hybrid than average, not more on-premises.

Where the Subscription Edition cohort keeps its data, against the full sample			
Cloud footprint	n	Moved to SPSE	All respondents
We have completely migrated to the cloud	266	21%	28%
We still have a small amount of data stored on-premises	266	43%	36%
We have roughly equal amounts on-premises and in the cloud	266	28%	25%
We have the majority of data on-premises, with only some in the cloud	266	5%	7%
All of our data is stored on-premises	266	3%	3%

IT OPERATIONS SURVEY · N=943

43% have only a small amount left on-premises, seven points above the sample. 28% run roughly half and half. Just 8% are on-premises-heavy. And then there's the whopping 21% who say they're already fully in the cloud, while also saying they migrated their SharePoint 2016/2019 to Subscription Edition.

We had trouble squaring that. There are a few possible explanations. One is that you can run Subscription Edition on cloud infrastructure, so these teams could have migrated to it and then run it on IaaS. We didn't ask, but based on Microsoft's documentation that path doesn't look popular enough to account for one in five of the Subscription Edition migrators. It's also technically possible they migrated to Subscription Edition and then to the cloud, but then why not say they moved from SharePoint 2016/2019 to the cloud, which was one of the options?

The likelier reading is the one Richard Harbridge described earlier in this chapter. The server didn't stop existing. IT leaders stopped counting it. For most of this group Subscription Edition isn't the estate. It's one workload they've decided to keep while everything else runs in the cloud.

Two governance models, and it shows

Of the three SharePoint paths, this one had the most trouble: 87% had at least one governance incident in the past 12 months, against 72% of the teams who moved SharePoint to the cloud. Some of that is who's looking. Not all of it.

The Subscription Edition incident gap, before and after controlling for detection				
Sample	SPSE n	Moved to SPSE had an incident	Moved to SharePoint Online had an incident	Gap
No control	266	87%	72%	15 pts
Teams with full visibility	106	86%	74%	12 pts
Teams running proactive monitoring	70	77%	69%	8 pts

IT OPERATIONS SURVEY · N=943

Each row compares two groups. The n column is the Subscription Edition side; the SharePoint Online side is 363 with no control, and smaller in the two controlled rows.

Compare only teams who monitor the same way and the gap narrows from 15 points to 8. It doesn't close. Eight points separate two groups finding problems the same way, and the difference between them is that one of them is doing it in two places.

What held the stragglers back

21% of organizations still had SharePoint 2016 or 2019 in play when support ended. We asked them the primary reason they hadn't finished, and the blockers are not primarily financial.

What's holding the holdouts

WE ASKED What is the primary reason your organization has not completed the migration?



The deadline moved almost everyone. For the few still stuck it isn't money. It's complexity and missing hands, which no end-of-support date fixes.

IT OPERATIONS SURVEY · N=943 · the 180 still on SharePoint 2016 or 2019

Complexity and capacity together account for nearly two thirds. Budget is a distant third. The teams still on legacy are not underfunded so much as under-resourced against a genuinely hard technical problem, usually one involving customizations somebody built a decade ago and nobody has owned since.

Prioritization by leadership came in last, which is worth holding against one of the data points from the governance chapter [CH 3 Governance →](#). Asked what would most improve their Microsoft 365 governance and security, executive buy-in was the second-most selected answer. Governance work stalls for want of a mandate. Migration work has the mandate and stalls anyway.

WHERE SHAREGATE FITS

Complexity is what ShareGate Migrate is built for. It takes the manual work out of the hard parts: recreating broken permission inheritance, mapping users and groups automatically, and cleaning up legacy hierarchies during the move rather than after it.

[See how Migrate handles complex SharePoint moves →](#)

Self-service and workspace lifecycle: how the estate grows

Somebody decides they need a team or a site, and it exists, with a name, an owner, a purpose, and a line in a log. When the project ends, nothing corresponding happens. No request arrives, nobody signs anything, and the workspace carries on

holding whatever it held. Every control in the governance chapter still applies to it


CH 3 **Governance** →

What the three self-service postures get you

How much can people do without asking IT? Less than you might expect.

Self-service is still rare

WE ASKED How much can employees do in Microsoft 365 without going through IT?




Only 1 in 10 give employees broad autonomy in M365.

IT OPERATIONS SURVEY · N=943

Only one in ten organizations gives employees broad autonomy in Microsoft 365. Nearly half route workspace creation through IT entirely. This is usually presented as a governance strength, and in one sense it is: fewer uncontrolled workspaces, fewer orphaned Teams, a cleaner tenant. The standard counter-argument is that tight control simply relocates the demand, pushing people toward personal drives, unmanaged SaaS, and AI tools nobody approved.

We tested that. It isn't what the data shows.

Shadow IT or shadow AI, by how much employees can do without IT



The sample rate is 25%. The full spread across all three postures is 2 points.

IT OPERATIONS SURVEY · N=943


Locking it down doesn't help. Opening it up doesn't hurt. A quarter of organizations hit shadow IT either way, and the three postures land within two points of each other. Guided self-service is marginally the tidiest, for what it's worth.

That's worth reporting precisely because the opposite is so widely assumed—including by us, before we looked. Whatever produces shadow adoption is happening somewhere this setting can't reach: unmet needs, procurement speed, or tools that arrive through a browser without asking anyone's permission.

Which is useful, because it means the self-service decision can be made on its own merits, as a question about admin load and user experience, rather than as a governance gamble. The organizations doing it well tend to sit in the middle band—guided self-service with guardrails, templates, and automatic lifecycle attached—because that combination keeps the tenant tidy without making IT the bottleneck for every new project.

What happens to a workspace when nobody needs it

The most common way an inactive workspace gets dealt with is by hand, during a review somebody scheduled.



35% have something that runs on its own. The other 65% depend on a person. Every workspace that is never retired stays in the estate permanently: holding permissions, consuming storage, and remaining reachable by both people and AI.

Periodic manual cleanup works when the surface is small. It stops working when a tenant holds hundreds of workspaces, thousands of files inside each one, and permissions assigned years ago by people who have since left.

“ Copilot honors every permission you have set. What it removes is obscurity. Organizations leaned on content being hard to find for years, and an AI agent doesn't care that nobody remembered the site existed.”



DAVID DREVER
MICROSOFT MVP

So every stale workspace nobody has reviewed is a workspace AI can still reach

CH 1 **Confidence ≠ control** →.

What automating lifecycle can tell you

Teams with automated retirement report the most AI exposure of any group. That isn't a reason to avoid automating it.

AI exposure and monitoring, by how inactive workspaces are handled				
Lifecycle approach	n	Had an AI exposure incident	Can't say	Had one or can't say
Automatically archived or deleted	300	38%	5%	43%
Reviewed periodically, archived by hand	472	24%	9%	33%
Handled manually when issues surface	60	27%	17%	44%
Sample average	851	29%	8%	37%

AI GOVERNANCE SURVEY · N=851

The 15 respondents with no formal lifecycle process and the 4 who were unsure are too small a base to report.

The incident column isn't a safety ranking. The automated group finds the most because it is looking hardest: 66% of them monitor AI use continuously, against 40% of the teams who review periodically and 23% of the reactive ones. Add the first two columns together and the two ends land in nearly the same place, 43% and 44%, either having had an incident or being unable to say. They got there from opposite directions. One found its problems. The other can't tell you whether it has any.

A low incident count on a manual process isn't evidence that less is happening. It's evidence that less is being found. It's a floor, and the teams who can't see their tenant have the lowest floor of anyone. What separates these groups isn't how much went wrong. It's how many of them can tell you what went wrong, and what they did about it.

“ Organizations underestimate how easily good governance intentions come apart in day-to-day SharePoint use, and they don't really know what content they have. They think they know. They think the security is in place. In most cases it's loose, with a lot of loose ends.”




MIKE MAADARANI
MICROSOFT MVP

Getting a workspace is hard. Losing one is nobody's job.

The front door and the back door, side by side.

What stands between creating and retiring workspaces



Two surveys with two sets of respondents, so this is not a within-organization comparison. The creation figure is a request queue (47%) plus guardrails and templates (43%).

IT OPERATIONS SURVEY · N=943 · AI GOVERNANCE SURVEY · N=851

90% of organizations put something in the way of creating a workspace: a request queue, or a guardrail. 35% put anything in the way of one outliving its purpose. That's the mechanism connecting this chapter to every other one. The estate doesn't get messier because teams are careless. It gets messier because creation is gated and retirement is not.


“ For many organizations, the content management model has been really good at content checking in and really bad at deleting what is no longer needed. It's the roach motel or Hotel California, whichever analogy you like. People are much better at adding content than they are about removing content.”



SUSAN HANLEY
MICROSOFT MVP

They bought a better view. They didn't buy much more help.

318 organizations rate their own governance "highly automated and continuously monitored"—the top band of five. Measured against the sample, control by control, the claim only half holds.



The first three are controls that let a team see the estate; the last three act on it without a person. Top band n=318 of 851.

AI GOVERNANCE SURVEY · N=851

On the controls that let a team see its estate, the top band runs 13 to 15 points ahead of the sample. On the controls that act without a person, 5 to 8. Automatic workspace retirement sits in the second group, 43% against 35%. The gap between those two clusters is the whole self-assessment: the top band has bought itself a much better view of the tenant and roughly the same amount of help.

“ Having the visibility and acting on it are two different projects. Plenty of teams can produce the report. Far fewer have restricted access control or content discoverability switched on because of what the report told them.”



NOOREZ KHAMIS
MICROSOFT MVP

Which means the most sophisticated governance operations in the study are still, in the majority of cases, retiring workspaces by hand.

WHERE SHAREGATE FITS

ShareGate Protect finds inactive and orphaned workspaces automatically, shows how much storage each one is consuming, and delegates the keep-or-archive decision to the people who actually know the answer. The estate stops growing while nobody's looking.


[See how Protect automates workspace lifecycle →](#)

What the estate charges you for holding it

The estate has a cost profile, and it grows on its own. Storage costs growing faster than expected affect 39% of organizations, second only to AI licensing in the cost rankings. Maintaining on-premises infrastructure alongside cloud affects 14%.

Four of these seven start inside the estate

WE ASKED What are the biggest cost-related challenges in your Microsoft 365 environment? (choose up to two)



Four of these seven start inside the estate, and none of them are a purchase anyone argued for. Storage grows because content is created faster than it's retired. Duplicate tenants exist because an acquisition closed and consolidation did not. Unused licenses are people who left. On-premises alongside cloud is a migration that stopped partway.

IT OPERATIONS SURVEY · N=943

Storage growth is the purest expression of what an ungoverned estate does over time. Nobody decides to spend more on storage. Content accumulates, nothing gets retired, and the invoice reflects it. The cost chapter covers the economics in full **CH 2 The AI reckoning →**; the point here is that this cost is a direct output of the estate's lifecycle practices, or the absence of them.

One cost is flat. The other tracks your footprint.

Read the two cost lines down the footprint bands.

Storage growth and on-premises upkeep as cost challenges, by cloud footprint			
Cloud footprint	n	Storage growing faster than expected	Cost of on-premises alongside cloud
Completely migrated to cloud	268	41%	12%
Small amount on-premises	335	40%	13%
Roughly equal	240	40%	17%
Majority or all on-premises	100	34%	17%
Sample average	943	39%	14%

IT OPERATIONS SURVEY · N=943

Respondents could choose one or two answers, so a percentage here is a share of a two-item budget.

Storage pressure sits around 40% in every band except the most on-premises one, which comes in at 34%. That band has the least in the cloud to begin with, which is the whole explanation. Self-service posture doesn't move it either, running between 35% and 41%.

The cost of running on-premises alongside cloud runs the other way, rising as the footprint does, and it's the smaller number at every point on the chart.

Chose two environments, counts one

Now split the same costs by the SharePoint route teams took.

Estate cost challenges, by what a team did about SharePoint			
SharePoint 2016/2019 path	n	Cost of on-premises alongside cloud	Storage growing faster than expected
Still migrating	146	20%	34%
Migrated to SharePoint Online	363	14%	43%
Migrated to SharePoint Server Subscription Edition	266	8%	45%
Sample average	943	14%	39%

IT OPERATIONS SURVEY · N=943

Respondents could choose one or two answers, so a percentage here is a share of a two-item budget.

Ask the Subscription Edition cohort what's expensive and they'll tell you about storage. 45% name it, more than any other path in the study. Ask about the on-premises half they deliberately kept, and 8% mention it. The lowest of anyone, from

the teams most committed to running it—76% of them hold data in both places.

A footprint somebody chose on purpose stops registering as an expense.

What to do about it

This week

QUICK WIN

— — — EFFORT

Build a disposition-based inventory of the estate

Most teams can produce a rough cloud-versus-on-prem split from memory. Far fewer can say, workload by workload, what's on-premises and why. For each major workload or platform, record:

- Owner
- Location
- Support status
- Sensitivity and criticality
- Key dependencies
- The reason it remains where it is
- The plan for it: retain, retire, migrate, modernize or replace
- A next review date

Include Power Platform solutions, connectors, service accounts and externally indexed sources, because that layer is the one most likely to be missing from the last inventory anyone made.

The reason-it-remains column is the whole exercise. Some of those reasons are real constraints: a compliance requirement, an integration dependency, an operational resilience argument, a product limitation, the economics of moving. Constraints need a plan to satisfy them ([CH 5 Compliance →](#)). The rest is work that hasn't been scheduled, and writing it down is what gets it scheduled.

IT operations survey · n=943

This quarter

MEDIUM LIFT

— — — EFFORT

Make the sanctioned path the fast path, and put a date on the remaining legacy

Look at where your self-service sits. If nearly everything routes through IT, work out what the request queue is costing in both admin time and user experience. Moving the most common request types—a standard team, a project site—into guided self-service with automatic lifecycle attached usually reduces both.

At the same time, assign an owner and a target quarter to migrate or archive anything still on unsupported or soon-unsupported infrastructure. The end-of-support data shows the effect a date has. Use it.

SHAREGATE MIGRATE

With Migrate you don't have to write scripts or remap permissions. You get a full analysis of your source tenants, and you can move identity, mailboxes and content in the same tool.

[See how it works →](#)

IT operations survey · n=943

This year

STRATEGIC INVESTMENT

EFFORT

Automate lifecycle so the estate stops growing by default, and govern hybrid as one environment


Only 35% of organizations automatically retire inactive workspaces. Until that runs on its own, every other control you build is applied to a surface that keeps expanding, and in a hybrid estate it usually only expands on one side, because the automation you can buy for the cloud half doesn't reach the on-premises half.

Use a common set of control objectives, owners, measures, and evidence across cloud and on-premises workloads, while allowing enforcement to vary by platform.

The second half matters just as much and gets less attention. If roughly seven in ten organizations are running meaningful on-premises alongside cloud, and 31% have deliberately committed to a supported on-premises future, then governance that only covers the cloud half covers most of the risk but not all of it. Treat hybrid as the operating model it demonstrably is, rather than a transitional state that will resolve itself, and make sure your visibility spans both.

AI governance survey · n=851

CH 04 · END



05 Compliance

CHAPTER 5 IN BRIEF

34% delay or avoid migrations over compliance risk, up from 20% just one year ago.

56%

name access control policies as a top compliance challenge

49%

of on-premises teams need proof Microsoft 365 meets their regulatory requirements before they'd move

80%

say digital sovereignty is influencing their AI governance decisions

WHAT CHANGED THIS YEAR

The total held. The composition didn't.

The share of teams who let compliance shape a migration decision barely moved, from 82% to 87%. Underneath it the share who delay or avoid outright went from 20% to 34%, and the validate-first middle gave up the difference. Confidence in the platform stayed flat across the same period, so this isn't teams losing faith in Microsoft 365. It's teams finding it harder to prove their own use of it.

THE DISTINCTION TO CARRY

A requirement that won't move, or the time to assemble an answer.

Some of what gets called a compliance blocker is a requirement that will never move. Some of it is the weeks it takes to put an answer together by hand. They look the same on a roadmap and they don't respond to the same fixes. Most of this chapter is about the second group, because that's the one you can do something about.

WHERE SOVEREIGNTY COMES IN

It stopped being a hosting question.

Compliance in Microsoft 365 has meant knowing where data lives for most of the last decade. Digital sovereignty turns that into a question about what can reach into the data, which is why it now shows up in AI governance decisions rather than hosting ones. A Copilot business case increasingly has to answer where the processing happens as well as what it costs

[CH 2 The AI reckoning →](#)

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 5 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

The Microsoft 365 Governance Index scores your visibility and content control against 1,700+ IT pros and ranks your widest gaps.


Compliance is the chapter everyone expects to be about restriction. It isn't. It's about evidence.

Teams are focused on whether they can show, on demand, that they're meeting their obligations. And on what happens to a project if they can't.

This chapter is about what that looks like in practice: the specific things IT pros name when asked what's holding up a decision right now, where data has to sit, who signs off before anything moves, and what teams still running on-premises say they'd need before they'd move.

Compliance became a stop sign for Microsoft 365 migrations

Before a migration, someone has to establish that the destination satisfies the requirements, that the move itself won't create exposure, and that both can be evidenced afterwards. We asked IT pros how heavily that weighs on their migration decisions, and the change from last year was toward more caution.




From validate-and-move to don't-move

Against last year the combined figure moved modestly, from 82% to 87%. The composition underneath it shifted much harder.


A gate that became a stop sign

WE ASKED How do compliance concerns impact your migration decisions?


Significantly, we delay or avoid migrations




Moderately, we validate before migrating



Slightly, a consideration not a barrier



Not at all



The overall figure moved five points. The group that stops work outright moved fourteen.

IT OPERATIONS SURVEY · N=943 · 2026 n=943 · 2025 n=650

Japan and Australia are new this year and both sit above average, so some of the rise is the sample changing. Not much: across the five countries asked in both years the top answer still went from 20% to 32%, and it rose in four of them. [Compare by region](#) to see each one.

Four countries moving the same way points to something shared rather than local, and the composition is where it shows: the band that stops work outright took everything the validate-first middle gave up.

“ Compliance teams are under real pressure from the board now. It may not be perfect, but they’re trying to make visible progress against risk, and that pressure is new.”



JOANNE KLEIN
MICROSOFT MVP


Newness is the operative part, because the platform assessment underneath it didn't move.

Nobody is doubting Microsoft

Asked directly, teams don't hedge.

Almost nobody actually doubts M365


WE ASKED How confident are you in Microsoft 365's ability to meet your organization's compliance needs?



96% are confident, only 3% are not. Yet 87% say compliance still shapes their migration decisions: caution without doubt.

IT OPERATIONS SURVEY · N=943

LAST YEAR → THIS YEAR



'Very confident' M365 meets compliance needs


A verdict on Microsoft's platform rather than on their own tenant, and it barely moved.

A year ago it was 60% very confident. So across the same period that produced a 14-point rise in teams blocking migrations, the platform assessment underneath it barely moved.

Those two facts are the chapter. Near-total confidence in what the platform can do, sitting beside compliance as one of the largest brakes on actually using it. Teams aren't losing faith in Microsoft 365. They're finding it harder to prove their own use of it, and confidence in Microsoft's capability is not the same thing as confidence in your own ability to demonstrate it **CH 1 Confidence ≠ control →**. The platform being compliant, and you being able to show that your configuration of it is compliant, are two different problems, and only one of them is Microsoft's to solve.

The moderate band is the one to notice. Requiring strict validation before migrating isn't resistance, it's a process requirement, and at 53% it's still the majority position. Validation is real and it's widespread. What varies is whether it's scheduled or discovered.

Compliance stops or defers migrations for half of IT teams in Japan and a fifth in Canada. The 34% average sits between two quite different pictures.



What the spread does establish is that the national picture matters. Two teams with comparable estates and comparable tooling can be working under materially different constraints. Japan is also the market least confident in the Microsoft platform itself, at 44% very confident against 57% for the study average, and it's the one place where the two measures move together.

“ In some jurisdictions an unauthorised disclosure of a single document containing a passport number is enough to trigger a reportable breach. That's why organizations have started asking where their data actually sits.”



DAVID DREVER
MICROSOFT MVP

Industry varies too, but less.

← BELOW SAMPLE AVERAGE ABOVE →

Sector	Strong (%)	Weak (%)	Sample Average (%)
Energy / Utilities (n=48)	48%	34%	34%
Retail / E-commerce	41%	34%	34%
Public Sector (n=41)	39%	34%	34%
Financial Services	38%	34%	34%
Education (n=70)	36%	34%	34%
Healthcare	35%	34%	34%
Technology / Software	32%	34%	34%
Professional Services (n=57)	28%	34%	34%
Manufacturing	23%	34%	34%


IT OPERATIONS SURVEY · N=943

SEVEN CONSTRAINTS, UNEVENLY CARRIED

The obligations behind the word "compliance" aren't interchangeable. Some describe a state the tenant has to be in. Others describe a process the team has to go through. They don't respond to the same fixes, and they don't cost the same amount of time.

56% name access control, the top compliance challenge

WE ASKED What are the biggest compliance challenges impacting your IT decisions right now? (choose up to three)



Three of the top four aren't about where data can live. They're about whether you can show what happened to it.

IT OPERATIONS SURVEY · N=943

IT pros selected 2.5 challenges on average, so these stack rather than substitute. The three the chart groups are access control, auditing and reporting, and approval processes. That's an evidence problem, and evidence problems respond to tooling rather than to intent.

Access control leading is consistent with what the same survey found elsewhere. Chapter 3 reported stale access as the most common governance incident of the year at 38%, with oversharing at 26% [CH 3 Governance →](#). The access surface producing those incidents is the same one compliance policy gets written against. The likeliest reading is that the compliance requirement and the governance backlog are the same work arriving under two different names.

SEVEN CONSTRAINTS, UNEVENLY CARRIED

The M365 Governance Index scores ten dimensions of your estate and your AI governance, then names the three where you sit furthest behind the field. Ranked by distance rather than by lowest score, so the weakness everyone shares doesn't crowd out the one that's actually yours.

[Get your score →](#)

The industry breakdown holds one result worth pausing on.

Top compliance challenges, by industry					
Industry	n	Data residency	Access control	Regulatory delay	Auditing needs
Technology / Software	190	56%	53%	43%	33%
Healthcare & Life Sciences	101	50%	50%	38%	36%
Retail / E-commerce	109	46%	60%	38%	34%
Manufacturing	164	38%	57%	43%	36%
Financial Services / Insurance	138	36%	57%	46%	41%
Sample average	943	44%	56%	41%	37%

IT OPERATIONS SURVEY · N=943


Technology and software organizations report the highest data residency concern in the study, at 56%, twenty points above Financial Services. That's the opposite of the intuitive ranking, and the likely explanation is the sovereignty pressure covered later in this chapter: technology companies are disproportionately multinational, disproportionately European in this sample, and disproportionately the ones being asked where their customers' data physically sits.

Of the five sectors with bases above 100, Financial Services leads on the two evidentiary measures, auditing needs and regulatory delay, which is closer to what you'd expect. Across all ten sectors those two are led by Energy and by Education, both on bases too small to rank.

The requirement is fixed. The approval cycle isn't.

Regulatory approval processes are named by 41% of IT pros, and they're the clearest dividing line in the study between teams who delay migrations and teams who don't.

Approval is what separates stalling teams



Access control is the most-named challenge in the study, and it barely tells these two groups apart: 53% against 57%.

IT OPERATIONS SURVEY · N=943

Each row compares the 320 teams that delay or avoid migrations with the 623 that don't. People could pick more than one challenge, so the two figures on a row are separate shares rather than halves of a split.

Approval processes split the two groups by nine points, the widest gap in the list. Certifications come next at seven. The M&A row is the same size and points the other way, which is the one worth stopping on.

Labelling is identical at 25% either way. Access control barely moves, and what movement there is runs backwards: 53% of the teams that stall name it, against 57% of the teams that don't. The most-named compliance challenge in the study is not what separates them. The queue is.


Teams whose compliance challenge is managing multiple tenants during a merger are less likely to delay, at 9% against 17%. The survey doesn't test why, but it could be that deal-driven migrations arrive with a contractual deadline, and a deadline someone else set doesn't leave room to wait.

The approval cost doesn't show up where you'd look for it. Teams naming approval processes don't report longer projects. If anything they report slightly shorter ones: 19% finished inside three months against 13% of everyone else, and the share running past a year is identical at about 10%. That looks backwards until you notice

who the question reaches. It only asks teams who ran a project. A migration still sitting in an approval queue has no timeline to report, so the teams that got through the queue are the only ones the question can see.

The auditing gap is a tooling gap

Two numbers, from opposite ends of the same problem.



Two ends of one problem, from two different questions: the biggest compliance challenges here, and how teams are handling governance and security in Microsoft 365 in [CH 3 Governance →](#).

Those two numbers describe the same problem from opposite ends. A meaningful share of teams are being asked to produce audit-grade evidence about permissions, retention, and access history, using tooling that was designed to administer a tenant rather than to report on it [CH 3 Governance →](#).

The usual result is a manual export, a spreadsheet, and a week nobody planned for. When an auditor asks who had access to this site in March, the honest answer for most environments is that it would take some work to find out.

Having been caught moves the number. Chapter 3 reported that 35% of organizations had an audit or compliance gap in the past 12 months. Among those teams, 43% name auditing and reporting tools as a compliance challenge, against 33% of teams with no finding. The teams who have been through a finding are the likeliest to name the reporting gap.

“ Compliance is very often not in IT, and when they start freaking out about something, IT has to react. That doesn't mean IT is gung-ho to solve the problem. They're gung-ho to check the box.”



MARC ANDERSON
MICROSOFT MVP

Given a single choice about what would most improve their governance and security, 25% of teams who delay or avoid migrations choose executive buy-in to prioritize governance, against 17% of everyone else. Only 1% of them ask for more budget, so this isn't a money problem. It's about getting leadership to prioritize governance work at all, even though it isn't glamorous and doesn't have a concrete deadline like other projects. Chapter 4 found the mirror image on migrations: prioritization came last among the reasons a legacy migration hadn't finished

CH 4 The estate →. Governance work stalls for want of a mandate. Migration work has the mandate and stalls anyway.

WHERE SHAREGATE FITS

ShareGate Protect is built for the evidence problem. It gives you a tenant-wide view of what users and Copilot agents can reach, refreshed continuously, with an action log on every change and a history behind every access decision. So you can show your work without assembling it first.

[Discover ShareGate Protect →](#)

Two constraints that connect to other chapters

Sensitivity labelling and data classification (25%) looks like a middling concern until you read it alongside chapter 1 **CH 1 Confidence ≠ control →**. Labels are one of the primary levers for controlling what AI can surface, because they're how the platform knows which content deserves the strictest defaults. A quarter of teams naming labelling as a compliance challenge is, in practice, a quarter of teams telling you their AI exposure controls rest on an incomplete foundation.

Managing compliance across multiple tenants during M&A is the smallest entry on the list at 14% across the sample, and one of the hardest problems on it. Chapter 6 finds that 23% of organizations say their most significant migration was an M&A move **CH 6 Migration →**. During those integrations, two tenants run in parallel with

different policies, different labels, and different retention rules, while the compliance obligation applies to both as one organization. The teams who named it are not overreacting.

PROOF, AND THE KEYS

What on-prem teams need before migrating to Microsoft 365

Not every organization keeps its data where Microsoft 365 can reach it. File shares, on-premises SharePoint and local servers still hold a meaningful share of the corporate record, and the teams running them carry the same obligations as everyone else.


Chapter 4 covers how large that footprint is and how much of it is deliberate

CH 4 **The estate** →

The question here is narrower: for teams still holding data on-premises, what would have to be true before they'd be confident moving it. Their answers are the most direct statement in the study of what compliance assurance actually has to look like.

What moves on-prem data to the cloud? Proof.

WE ASKED If your organization stores data on-prem, what would need to be true for you to feel confident moving it to the cloud? (choose up to two)



The two ways of having somebody else vouch for it come last, at 12% and 8%. Nobody's waiting to be told it's fine.

IT OPERATIONS SURVEY · N=943

· the 675 who still store some data on-premises

Every leading answer is an artefact. Proof. Key control. A validated pilot. Residency guarantees. Nobody is asking to be reassured, and the two options that amount to someone else vouching for it, an auditor endorsement and a third-party assessment, come last by a wide margin.

Two of those conditions are less about where data sits than about who can read it. 45% want to hold their own encryption keys. Residency guarantees sit fourth overall at 34%, but that average hides two different worlds: Japan at 49%, Australia at 42%, Germany at 41% and France at 39%, against the US and Canada at 28% and the UK at 23%. The markets that want a residency guarantee are the ones about to ask the same question of their AI, which is where this chapter goes next. Japan's cut stands on 71 respondents and France's on 59, so read both as directional.

The alternative to delaying your migration is moving without the artefacts. And finding out afterwards that you missed something.

“ I've never thought compliance gets in the way. Compliance is how you define and govern your data. I would far rather carry the delay than finish a migration and discover I lost the crown jewels. I've seen that happen.”



DAVID DREVER
MICROSOFT MVP

Every condition on this list is a way of making sure that doesn't happen. Which is unusually good news, because artefacts can be produced on a schedule. A validated pilot migration is a project with a start date. A control-mapping document

against your specific regulatory obligations is a deliverable. Customer-managed keys are an architectural, licensing, and operating decision. None of this requires the regulator to change position or the platform to add a feature.

“ A validated pilot with the compliance team in the room from day one does more to unblock a roadmap than another year of vendor assurances, because at the end of it you own an artefact instead of an opinion. And you can reuse that artefact on the next twelve workloads.”



RICHARD HARBRIDGE
MICROSOFT MVP

WHERE SHAREGATE FITS

A validated pilot is the one condition on that list you can start this week. ShareGate Migrate makes it easy and secure: your source is only ever copied, so nothing you're protecting is at risk, and you finish holding an artefact rather than an opinion. A per-item record of what moved and what didn't, yours to hand to whoever signs this off.

[Discover ShareGate Migrate →](#)


Fully cloud, still gated

As we mentioned in chapter 4, we were surprised to find that 72% of organizations still have data on-premises (CH 4 **The estate** →). We thought compliance might be the explanation. The data doesn't support that, and the way it fails is more interesting than the assumption.

Finishing your migration doesn't end the compliance delays

WE ASKED

How do compliance concerns impact your migration decisions?



43% of the teams that already finished still delay or avoid a move over compliance. That's the highest of any group.

IT OPERATIONS SURVEY · N=943

If compliance were what kept data on-premises, this would climb to the right. The two groups on the right are the smallest, 68 people and 32, so the upturn there is too small to rank.

So compliance isn't the main reason organizations are keeping data on-premises. Chapter 4 asked the 180 organizations still carrying SharePoint 2016 or 2019 why they hadn't finished, and compliance came fourth at 13%, behind migration complexity at 37% and a lack of internal resources at 26% [CH 4 The estate →](#).

Read the other way round, the same crosstab is starker still.

Cloud completeness, by how much compliance shapes migration decisions		
Compliance impact on migration	n	Fully in the cloud
Significantly delays or prevents moves	320	36%
Moderately — we validate before moving	503	24%
Slightly — a consideration, not a barrier	91	24%
Not at all	29	38%
Sample average	943	28%

IT OPERATIONS SURVEY · N=943

Gating and cloud completeness move together at the top of that table. The 29 teams who say compliance doesn't affect them at all are too few to read, which is why that row is marked.

The likeliest explanation is volume. Fully-migrated organizations are the most migration-active group in the study, reporting 2.6 migration types completed, underway or planned in the past 24 months against a sample average of 2.4, and only 2% report none at all. Among organizations with all their data on-premises, a third ran no migrations in the period. More migration work means more occasions for compliance to stop one.

Volume doesn't account for all of it. Teams with a small on-premises remainder report almost as much activity, at 2.5 types, and gate 11 points lower. But it reframes what the top row means.

Reaching the cloud doesn't end the migration programme. Chapter 6 makes the case that migration is a standing capability rather than a project [CH 6 Migration →](#), and the compliance validation travels with it. Compliance doesn't stop movement. It adds a stage to it, and that stage is producing the evidence, for an auditor, a regulator, or a board.


FROM WHERE IT SITS TO WHAT CAN REACH IT

Sovereignty is now an AI governance question

Residency asks where data is stored. Sovereignty asks the broader question: whose jurisdiction does it sit under, and who ultimately controls the systems it passes through? What's changed this year is how far sovereignty reaches, because it's now shaping decisions about AI, not just storage.

80% say digital sovereignty is shaping AI decisions

WE ASKED To what degree is digital sovereignty influencing your AI governance decisions?




Data

Major influence	27%
Moderate influence	53%
Minor influence	15%
No influence	4%
Not sure	1%

For a third it's a major influence, not just a consideration. Only 4% say it has none.

AI GOVERNANCE SURVEY · N=851

Sovereignty requirements are written by national governments and regional blocs. If the pressure followed the policy, it would concentrate where the policy is written. It doesn't.



Two questions from the same study, on one scale. Ask the same pair by company size instead and the difference closes, 15 against 18: this is about where you operate, not how big you are.

“ In highly regulated organizations I see people digging in hard on where their data sits. The political climate is part of it. But it varies enormously from one company to the next, far more than you would expect.”



JOANNE KLEIN
MICROSOFT MVP

So the useful question isn't which jurisdiction you're in. It's how hard your own compliance function is pushing.


NATIONAL AVERAGES MAKE POOR SUBSTITUTES

The M365 Governance Index scores your estate and your AI governance against the 1,700+ IT pros in these studies, filtered to your industry and your country. You get a percentile instead of an impression, and your gaps ranked by how far behind the field they actually are.

Get your score →

Most teams say they'd prioritize EU platforms

Asked how much they prioritize EU-based or French-developed platforms over US platforms such as Microsoft 365, a majority of IT leaders said they lean European.



59% lean toward EU-based or French-developed platforms, 31% toward US platforms, and 10% are neutral. Read that as a stated preference rather than a procurement plan: chapter 1 found Copilot deployed or being deployed in 93% of these organizations [CH 1 Confidence ≠ control →](#). Preference and installed base are pointing in different directions.


Unlike the sovereignty requirement, the platform preference does follow the map. European countries prefer EU-based platforms most, followed by the UK, the US, and Canada, with a 42-point spread between France at 82% and Canada at 40%. Ireland and the Netherlands sit inside that European cluster on 54 and 56 respondents, so treat both as directional.

Half of US-based IT leaders leaning toward EU platforms is the part worth pausing on, and the policy question underneath it points the same way. 46% report being significantly or moderately affected by French government directives encouraging EU-based technology platforms, and that includes plenty of respondents nowhere near France. Only 27% report no effect.

Part of that is the sample: this study is weighted toward European markets and technology-sector organizations, so read these figures as that audience rather than as a global average. But the question also asks about the organization rather than the individual, and the likeliest reading is that multinationals with European subsidiaries or European customers are answering for the whole entity. The rest of their answers support it. US respondents reporting significant impact favour EU platforms at 83%, against 44% of other US respondents, and 95% say sovereignty is influencing their AI governance, against 81%. Whatever they're describing, they describe it consistently. That group is 41 respondents, so it's directional, but it's internally coherent.

And then the pattern that has run through this entire report

Ask whether Microsoft has it covered, and the doubt disappears again.




Eighty percent say sovereignty is shaping their AI decisions. Fifty-nine percent lean European. Ninety percent still trust Microsoft 365 to meet their sovereignty needs. Those three numbers are not contradictory, but they describe an unstable position: teams are hedging toward EU platforms while expecting the incumbent to remain adequate. Something eventually resolves that tension, and it will probably be resolved by procurement rather than by IT.

The country split is the figure above. 34 points separate the UK at 59% very confident from the Netherlands at 25%, and the Netherlands anchors that low end on 56 respondents, so the width of the spread is directional even though its shape isn't.

Doubt is the sharper measure, because it doesn't depend on anyone reaching for a top box. 17% of French respondents say they're uncertain or not confident, along with 16% in the Netherlands and 11% in both Germany and Ireland. In the UK and the US it's 4% each. Canada belongs to neither group: just 5% express doubt, but only 35% are very confident, with a clear majority sitting in the middle.

It would be tidy if platform preference and platform doubt were the same thing, if leaning European meant expecting Microsoft to fall short. The data doesn't support it.



Each row is the share of that group who are uncertain or not confident, against 8% for everyone. Read doubt rather than "very confident": people who picked a strong answer on platform preference also picked the strongest answer on unrelated questions in this study, in all seven countries, which is a habit of answering rather than a view about sovereignty.

Respondents who prioritize EU platforms most strongly are no less likely to doubt Microsoft 365 than anyone else. Preference and doubt are running on separate tracks.

A sovereignty requirement used to shape an infrastructure decision. On this evidence it now shapes an AI one, which means a Copilot business case has to answer where the processing happens as well as what it costs

CH 2 The AI reckoning →. The teams that already know where their AI processing happens will spend the next two years answering the question rather than researching it.

What to do about it

This week

QUICK WIN

EFFORT

Turn your constraints into a requirements register

Not a list of the top three in this report. Access control, residency and approval timelines are the three most-named constraints in the data, and this chapter's own industry table shows they're the top three for almost nobody in particular: Technology leads on residency and sits below sample on access control, Financial Services leads on approval delay and sits eight points below on residency, Retail leads on access control and sits below on approval. The delay-vs-everyone cut goes further: the most-named challenge in the study is the one that *least* distinguishes teams that stall from teams that don't.

So build the register instead. For each requirement record the obligation and its source, the applicable data and workload, the control objective, the required evidence, the decision owner, the current status or exception, and a review date. Chapter 4's estate inventory is where these surface [CH 4 The estate →](#); this is where they get specified.

Then identify the three unresolved requirements most likely to delay work in flight. Precision matters, because vague constraints block more than specific ones. "We can't put regulated data in the cloud" stops a project. "Category 3 data requires customer-managed keys and EU residency, with a six-week approval cycle" is a scope item, and it's the version you can hand to a vendor.

IT operations survey · n=943

This quarter

MEDIUM LIFT

EFFORT

Run a validated pilot migration, and keep the evidence pack

37% of on-prem holdouts say a successful pilot with compliance validation is what they need, and 49% want proof of regulatory fit. A pilot produces both at once.

Choose a workload that's genuinely representative rather than conveniently easy, involve the compliance team in defining the validation criteria before you start, and document the controls, the testing, and the outcome as you go. The artefact is as much the point as the migration. Done once properly, it becomes the template that unblocks everything behind it.

The pilot is also the thing to take upstairs. Executive buy-in to prioritize governance is the ask that rises most among teams who delay or avoid migrations: 25% pick it against 17% of everyone else, which moves it past automated remediation into second place behind AI agent controls. Budget comes last in both groups, at 1% among the teams that stall. A completed validation with a compliance sign-off attached is a considerably better ask than a request for money nobody is asking for.

IT operations survey · n=943

This year

STRATEGIC INVESTMENT

EFFORT

Negotiate key control and residency up front, and build a reusable evidence set

45% of on-premises holdouts say control of their own encryption keys is a precondition. That's a contractual and architectural decision best made at the start of a program rather than discovered in the middle of one.

Alongside it, build the evidence set as a maintained asset rather than a per-project scramble: control mappings to your specific obligations, audit and reporting coverage, residency documentation, and the pilot validation record. Across the whole sample, 41% say regulatory approval processes are delaying deployments and 37% say they lack adequate auditing and reporting tools.

SHAREGATE PROTECT

Protect gives you a tenant-wide view of what users and Copilot agents can access, ranked by severity, then lets you close what shouldn't be open and keep it closed with automated policies.


[Know what your AI can reach →](#)


The requirement is fixed. The approval cycle usually isn't, and it runs long because the evidence gets rebuilt rather than maintained. A standing evidence pack, produced by tooling rather than by hand, is what turns each approval from a research exercise into a review.

And expect sovereignty to keep escalating from an infrastructure question into an AI question. The teams that already know where their AI processing happens will spend the next two years answering the question rather than researching it.

IT operations survey · n=943

CH 05 · END





06 Migration

CHAPTER 6 IN BRIEF

94% of organizations moved something in the past two years. Only 7% would do it the same way again.

35%

of the biggest migrations were driven by org change, not technology strategy

83%

want identity prep inside their migration tool. 11% have it

2x

Shadow IT doubles between teams running one migration and teams running three

WHAT THIS CHAPTER
COVERS

Migration as a standing condition, not a project.

Most teams are running more than one migration at a time and starting the next one before the last finishes. This chapter is about what that costs, what teams wish they'd done differently, and which step reliably hurts the most.

WHY IT MATTERS MORE
NOW

A quarter of the biggest projects came from a deal.

Cloud migration has a finish line. Corporate change doesn't. Mergers, divestitures and restructures keep starting tenant-to-tenant work, and they arrive on someone else's timetable with a deadline already attached.

THE GAP WORTH
CLOSING

Identity is the step teams most want inside the tool.

83% want identity preparation handled inside their migration tool. 11% have it there. That's the widest want-versus-reality gap in this data, and it sits on the step teams already name as the hardest part of a tenant move.

WHAT PREDICTS BUYING
A TOOL

How much you move, not how big you are.

Third-party tool use runs from 41% among teams doing one kind of migration to 79% among teams doing three or more. The same pattern holds inside every company size band. The teams doing the most are also the most likely to say they should have bought sooner.

Every stat is tagged with the study it came from: the AI governance survey (n=851) or the IT operations survey (n=943). [Full methodology →](#)

This is Chapter 6 of *The State of Microsoft 365*, a report on how organizations run, secure, and migrate in 2026. [Learn more →](#)

SEE WHERE YOU STAND

Migration keeps arriving. The M365 Governance Index scores the estate it lands in, from visibility and tenant control to incident resilience, against 943 IT pros. Then it tells you what the most effective teams do differently.

[Get your score →](#)

Migration is the most misunderstood work in this report, because almost everyone still describes it as a project. Projects end. This doesn't.

Once compliance clears the path, the work becomes movement

CH 5 Compliance →


. And the movement does not stop. Organizations acquire, divest, restructure, consolidate, and modernize on a rolling basis, and each of those events produces a migration. The data shows teams working through several, finishing one and starting another, and consistently wishing they'd treated it as a capability the first time round.

Migration as an ongoing condition

Start with how much of it is actually happening.

Everyone's migrating something, and usually more than one thing

WE ASKED What types of migrations has your organization completed, started, or planned in the past 24 months?



Teams run ~2.5 migration types every two years. One-and-done is a myth.

IT OPERATIONS SURVEY · N=943

94% of organizations migrated something in the past two years, and those that did averaged 2.5 migration types completed, underway, or planned across the 24-month period. 6% stood still. Of the teams that moved anything, 81% ran two or more kinds of migration inside the same 24 months and half ran three or more. That's the case for treating migration as an operating capability rather than an event: the average team isn't *between* migrations, it's working through two or three of them.

Laura Rogers sees the same pattern from inside the projects.


“ Most teams still treat migration as something with a finish line. Move the content, move on. The successful ones build the ownership and the processes that make the next change easier, because there's always a next change.”



LAURA ROGERS
MICROSOFT MVP

A third of the biggest moves are organizational

WE ASKED Which of the following best describes your organization's most significant migration project in the past 24 months, the one that consumed the most time, budget, or attention?



Data

On-prem to cloud	49%
On-prem to on-prem	15%
► M&A or divestiture	24%
► Consolidation or restructure	11%
Other	1%

Organizational change (24% + 11%) **35%**

Half are still the classic lift to the cloud. But a quarter of the most significant projects were a merger, acquisition, or divestiture. That's the move that restarts the migration, governance, and cost clocks at once.

IT OPERATIONS SURVEY · N=943 · the 889 who ran a migration


Nearly a quarter of organizations say their biggest migration was driven by a merger, acquisition, or divestiture. The cost chapter found that cost pressure is a moderate or major driver of migration and consolidation for 88% of organizations, and duplicate M&A tenants are among the largest cost lines in the estate

. The two findings describe the same cycle from different ends. Add consolidation and restructure work and 35% of the most significant projects were driven by organizational change rather than technology strategy.

This is the structural reason migration never ends. Cloud migration has a finish line, at least in principle. Corporate change does not. As long as organizations keep buying, selling, and reorganizing, tenant-to-tenant work keeps arriving. It arrives on someone else's timetable, with a deal deadline attached.

Constant motion leaves a trace. Teams that ran three or more kinds of migration in the past two years are more likely to report a governance incident than teams that ran one: 83% against 68%. The sharpest split is shadow IT and shadow AI, where people pick up tools, apps or automations outside IT's line of sight. That one doubles.

Shadow IT doubles as the migration count climbs



Sixteen percent of teams running a single migration reported it. Thirty-two percent of the heaviest movers did.

IT OPERATIONS SURVEY · N=943

Each bar is the share of that group who reported a shadow IT or shadow AI incident in the past year, meaning someone picked up tools, apps or automations outside IT's visibility. The three groups are 166, 279 and 442 teams.

The survey can't tell you which way that runs. Repeated migration may leave sprawl behind, or the teams generating the most sprawl may be the ones restructuring most often. Both are plausible, and the pattern shows up in all five company size bands. What it does establish is that heavy migration activity and governance debt turn up in the same places.

We asked the MVPs why that debt accumulates rather than clears. David Drever put it down to change management rather than tooling.

“ The organizations that do this well take the time to get compliance-ready, then move into a controlled environment and keep it that way. Where it fails is where compliance was pushed onto people without anyone explaining why. Then all that work quietly decays, because nobody understands what they're maintaining.”



DAVID DREVER
MICROSOFT MVP


Decay is the right word for it. 77% of organizations reported a governance incident in the past year, and the teams moving most often are the ones furthest into that number [CH 3 Governance →](#).

How long it takes

Planning benchmarks matter when the timeline is set by a deal rather than by IT.

A third run past six months

WE ASKED How long did this project take from kickoff to completion?



Half land in the three-to-six month window. But 36% run longer than that and one in ten past a year, long enough for the estate to change underneath them.

IT OPERATIONS SURVEY · N=943 · the 889 who ran a migration


Most projects land in the 3 to 6 month band, but more than a third take seven months or more, including 10% that exceed a year. Read that against the M&A finding and the tension is obvious: a significant share of migrations are triggered by transactions whose integration deadlines are set in a boardroom, and the typical project takes half a year. Teams who haven't pre-built the capability are being asked to compress a six-month process into whatever the deal timeline allows.

Projects running past six months, by what drove the migration		
Most significant project	n	Ran seven months or longer
On-prem to cloud	435	34%
Merger, acquisition, or divestiture	214	39%
On-prem to on-prem	136	35%
Consolidation or restructure	96	41%
Other	8	75%
Sample average	889	37%

IT OPERATIONS SURVEY · N=943

Both deal-driven rows cover M365-to-M365 and Google Workspace-to-M365 moves. "Other" is 8 teams, too few to read.

Last year's report asked the same question of M&A migrations, so that's the one group where the two years compare. One band moved enough to mention.




Treat that as a direction rather than eleven exact points: last year's figure comes from a dedicated M&A module of 520 people, this year's from the 214 whose biggest project was a deal. The other bands drifted a few points each and none of them says much on its own. What thinned out is the quick one.

Why teams move

Teams could name up to three reasons. One of them is new this year.

Why teams migrate, and AI just joined the list

WE ASKED What drove the need for this migration project? (choose up to three)



Security still leads, but AI readiness is now a named reason to move (28%).

IT OPERATIONS SURVEY · N=943

Security and scalability remain the top two, as they were last year, though they've swapped places: in 2025 scalability led at 65% with security a point behind at 64%. Every driver scores lower in absolute terms this year, which likely reflects a change in how the question was asked rather than a collapse in motivation, so read the ranking rather than the levels. The newcomer worth watching is AI readiness, now a named driver for more than a quarter of organizations.

Two years ago AI readiness wouldn't have appeared on this list. It's now cited more often than end-of-support deadlines, and it connects this chapter directly to the first: content has to be somewhere Copilot can reach it, governed well enough that reaching it is safe [CH 1 Confidence ≠ control →](#).

Mike Maadarani described what that looks like at the moment of the move.

“ Migrate stale content into a tenant with Copilot switched on and it'll be indexed unless you have deliberately excluded it. You're not just moving old data, you're deciding what to feed your AI.”




MIKE MAADARANI
MICROSOFT MVP

What actually moves when a tenant moves

Scoping a tenant move starts with this list, and it runs longer than most teams expect.

Tenant-to-tenant touches everything

WE ASKED What workloads were (or will be) included in this tenant-to-tenant migration?



The 'mailbox move' is now a whole-estate move.

IT OPERATIONS SURVEY · N=943 · the 310 running a tenant-to-tenant move

Power Platform tops the list. That should reset the scoping conversation. The assumption that tenant-to-tenant work is mostly mail and files is out of date, and it's the assumption most likely to produce a mid-project surprise. Apps, flows, and connectors are the most commonly in-scope workload, they hold their own

permissions and connections, and they're frequently owned by someone outside IT who built them for a purpose nobody documented. The estate chapter covers what this says more broadly [CH 4 The estate →](#).


A third of these projects also involve Google Workspace, which is worth noting for anyone assuming cross-platform work is a niche case.

Power Automate flows are also the example named in the shadow IT question earlier in this chapter, the incident type that nearly doubles as migration volume rises. The workload teams most often have to move is the one they're least likely to have documented.

WHAT YOU BUY FOLLOWS HOW MUCH YOU MOVE

What teams use to move

Most teams don't pick a tool. They pick several.




Three of the options teams were offered, not all of them. Teams could pick as many as applied, which is why these don't total 100%.

That's the shape of a field where teams are stacking approaches rather than swapping them, which is what you'd expect from teams running two or three kinds of migration every two years. Native tools are the floor almost everyone builds on. What changes is what gets added on top, and how soon.

What teams would do differently

We asked the teams who'd already finished one what they'd change.



Only 7% would change nothing. The top three regrets are all the same shape: plan earlier, tool earlier, clean up earlier. Hindsight is the most honest advice in any survey, because the person giving it has already paid for the lesson. Three of the top four answers describe work that should have happened before the project started.

Microsoft MVP Marc Anderson names the most common version of that.

“ The biggest mistake is lifting and shifting, dragging all your prior mistakes along with you. Dump the shared drive into a SharePoint site and you've got a shared drive in the sky. You've solved nothing, except maybe your old server isn't going to die on you. Then in six months you do it again, and you've doubled your cost.”



MARC ANDERSON
MICROSOFT MVP

What actually predicts buying a tool

Not the size of the company. How much migration it runs.

The more you move, the more likely you are to buy




Run one kind of migration and it's a coin flip. Run three or more and it's close to a certainty.

Used at least one third-party migration tool.

IT OPERATIONS SURVEY · N=943

The obvious objection is that this is really a story about company size. Bigger companies run more migrations and buy more software. It isn't. The same pattern shows up inside every size band in the study, and it's widest in the largest one.

Same pattern, inside the biggest companies only



Among companies with 3,000+ employees, buying runs from 21% at one migration type to 73% at three or more. Those companies are 25% of the one-type group and 23% of the three-or-more group, so the mix barely moves.

IT OPERATIONS SURVEY · N=943

Only 42 teams this size ran a single migration type, so 21% is the softest number on this page, and a bigger sample could move it either way. The 52-point gap to the three-or-more group is wide enough that the pattern holds regardless.

The regret climbs the same way. Among teams that ran one kind of migration, 27% wish they'd bought tooling sooner. At three or more, 43%. Over the same range, the share saying nothing went wrong halves.

Tool use and regret, by how much migration a team runs				
Migration types run	n	Used a third-party tool	Wish they'd bought sooner	Nothing went wrong
Ran one migration type	166	41%	27%	10%
Ran two types	279	63%	36%	6%
Ran three or more types	442	79%	43%	5%
Sample average	943	63%	38%	7%

IT OPERATIONS SURVEY · N=943

Tool use is out of everyone in each group. The two hindsight columns are out of the 629 teams who answered that question – 113, 172 and 336 of them, group by group.

Split the same question by what teams actually bought and the answer lands in a place worth sitting with.

Regret about tooling, by what teams actually used		
Group	n	Wish they'd bought sooner
Native tools only	162	25%
Any third-party tool	427	45%
ShareGate users	201	47%
Sample average	629	38%

IT OPERATIONS SURVEY · N=943

Out of the 629 teams who answered the hindsight question, not all 889 who ran a project. Groups overlap – most teams used more than one tool.

The regret sits with the teams who bought, not the teams who didn't. Native-only teams are the least likely to wish they'd bought sooner. ShareGate users are the most likely, at the highest rate in the study.

Read plainly, that's hindsight rather than present-tense pain. The people who bought the tooling wish they'd bought it earlier. The people who never bought it don't report missing it. Both can be true, and only one of them is an argument. What the data supports is that the decision, once made, tends to feel overdue. What it doesn't support is any claim that native-only teams are quietly struggling. They don't say they are. They're also running the least migration, which is most of the explanation.

So the useful question was never *which tool*. It's how much migration you're about to run, and whether what you buy covers the part that actually hurts. The rest of this chapter is about that part.


HINDSIGHT, IN ADVANCE

Hindsight is the most useful data in this chapter and the most expensive way to get it. The M365 Governance Index ranks your widest gaps by how far they sit behind the field, so you can work the list before the next migration instead of after it.

See where your gaps are →

Identity as the hardest part

Ask what actually hurts in a tenant-to-tenant move and identity surfaces immediately.



36% find identity preparation difficult or the single hardest part. Another 40% call it manageable while noting it could be easier, which is a polite way of saying it works but costs more than it should. More than four in five want identity preparation

handled inside the migration tool rather than bolted on. Now the number that turns that preference into a gap.

83% want identity inside the tool. 11% have it.

Want identity prep handled inside the migration tool 83%



Use a migration tool that actually includes it 11%



The widest want-versus-reality gap in the migration data, and the reason the difficulty scores sit where they do.

Teams who called a unified workflow critical or very important, against teams whose actual method is a migration tool with identity prep built in. Both questions went to the 310 teams that ran tenant-to-tenant work.

IT OPERATIONS SURVEY · N=943

Four in ten prepare identities by hand

WE ASKED When preparing for this tenant-to-tenant migration, how did your organization handle identity preparation (creating users, mapping groups, assigning licenses in the destination tenant)?

Microsoft cross-tenant tools 41%



By hand in the admin centre ▶ 28%



PowerShell scripts ▶ 13%



Third-party tool with identity 11%



External partner 4%



Ad hoc, no defined process ▶ 2%



Not sure 1%



Microsoft's cross-tenant tools lead. But 43% still assemble identities manually, by script, or with no defined process at all. 83% say they want it inside the migration tool.

IT OPERATIONS SURVEY · N=943 · the 310 running a tenant-to-tenant move

That gap explains the difficulty scores above. More than four in ten teams are doing identity prep in a separate Microsoft tool, and another 41% are doing it by hand in the admin center or with scripts. The step teams find hardest is the step most often

handled outside the workflow that everything else runs in.

How hard identity prep feels, by how teams actually do it			
Identity preparation method	n	Difficult/Hardest	Easy
Using Microsoft's cross-tenant migration tools	128	34%	19%
Manually through the Microsoft 365 admin center or Azure portal	86	35%	33%
Using a third-party migration tool that includes identity migration	34	32%	38%
Using PowerShell scripts	42	41%	17%
Relying on an external partner / consultant	14	29%	29%
We don't have a defined process, it's ad hoc	5	80%	0%
Not sure	1	100%	0%
Sample average	310	36%	25%

IT OPERATIONS SURVEY · N=943

The difficulty does track with the method, though the bases get small quickly. Teams using PowerShell scripts report the most difficulty, while the small group with identity built into their migration tool reports the highest ease. That last group is only 34 people, so read it as a direction, not a decimal.

WHERE SHAREGATE FITS

ShareGate Migrate moves Entra ID in the same workflow as mailboxes and content. Conflicts surface before anything runs, and you can override any match by hand. Save the mapping plan and re-run it across waves.


[Discover ShareGate Migrate →](#)

The parts that do not go to plan

The pattern generalizes beyond identity.

What went wrong on migration projects

WE ASKED What were the biggest challenges during this project? (choose up to two)



Complexity, downtime and compliance validation lead. Permissions work sits well down the list, and it is the thing teams say blows up timelines. That gap is worth asking about.

IT OPERATIONS SURVEY · N=943 · the 889 who ran a migration project


Everyone who ran a project saw this list.

None of these is the bulk content move itself.

The survey routes the deal-driven projects, mergers and acquisitions and consolidations or restructures alike, to a separate challenge list, and theirs is more pointed still.

On a deal-driven migration, identity is the problem

WE ASKED What were the biggest challenges during this project? (choose up to two)



Identity and permissions take two of the top three spots on the deal-driven challenge list, on a timeline set by a transaction rather than by IT.

IT OPERATIONS SURVEY · N=943 · the 318 whose biggest project was a deal


The 318 teams routed here are those whose biggest project was a merger, acquisition or divestiture (214), a consolidation or restructure (96), or something else (8).

Identity and permissions occupy two of the top three. M&A projects also run longer. 39% pass the six-month mark against 34% of on-premises-to-cloud moves, on timelines usually set by a transaction rather than by IT.

The planned, happy-path portion of a migration, the part every tool demos well, rarely breaks a project. What breaks projects is everything that wasn't fully scoped: the customization nobody documented, the permission that doesn't map, the Power Platform flow that surfaces halfway through, the coexistence window that runs longer than planned.

The mailbox move breaks on more than downtime

WE ASKED What are the biggest challenges you face (or faced) specifically with the mailbox migration portion of this project? (choose up to two)



Downtime leads, but destination identity prep (38%) and coexistence (36%) come straight after it. Those are the two things nobody puts on the plan.

IT OPERATIONS SURVEY · N=943

· the 126 whose project included a mailbox move

The mailbox picture reinforces it: identity prep is the second-largest challenge, and coexistence, which is largely an identity and routing problem, is third. Which reframes what tooling is actually for. Automating the predictable steps has value,

but the predictable steps were never the risk. The tooling that earns its cost is the tooling that absorbs the exceptions, and it earns it repeatedly, because a team running 2.5 migration types every two years will meet those exceptions again.

“ Every migration I've seen go sideways went sideways in the last ten percent. The content moves fine. Then you find the workflow nobody documented, the group that doesn't map, the mailbox with twelve years of delegation on it. Teams budget for the ninety percent because that's the part you can estimate, and then the exceptions eat the timeline. The tooling question isn't "can it move the data". Everything moves data. It's "what happens when something doesn't fit".”



RICHARD HARBRIDGE
MICROSOFT MVP

BEFORE THE NEXT PROJECT

Most of what's on those lists is a condition of the tenant rather than an event in the project. Visibility and tenant control are two of the ten dimensions the M365 Governance Index scores. Answer the same questions these 943 IT pros answered and see which of yours sit furthest behind theirs.


[Get your score →](#)

Migration feels solved. It isn't.

One answer from each study, side by side.

Migration is a do-it-yourself job

WE ASKED Did you involve an external partner (IT consultant, VAR, or MSP) in this project?



Data


- **Entirely in-house** 77%
- Partner: planning only 5%
- Partner: implementation only 6%
- Partner: planning + implementation 8%
- Partner: all phases 4%

77% ran their most significant migration entirely in-house. Fewer than one in ten brought a partner into both planning and implementation.

IT OPERATIONS SURVEY · N=943
· the 889 who ran a migration

Almost everyone would bring a partner in for AI

WE ASKED How likely is your organization to engage an external partner to conduct an AI governance assessment before scaling AI use?



86% are likely to. Only 4% rule it out.

AI GOVERNANCE SURVEY · N=851

77% of teams handled their biggest migration entirely in-house. Ask the same kind of question about AI governance and 86% say they'd bring in a partner.

Two surveys, two questions, so this isn't anyone changing their mind. But it's a clean read on what teams think they've got covered. AI governance is new, so it gets help. Migration is familiar, so it doesn't.

Familiar is doing a lot of work in that sentence. Only 7% of teams would run their last project the same way again. 45% of the ones who bought tooling wish they'd bought it sooner. 29% of that same AI group has already had an exposure incident **CH 1 Confidence ≠ control →**. And when teams do bring a partner in, they rarely do it for one phase. It's planning and implementation together, or the whole thing, which is what you'd expect from work that turned out to be bigger than it looked.

What to do about it

This week

QUICK WIN

 EFFORT

List what you'll have to migrate again in the next 18 months

Not what you're migrating now. What's coming: the acquisition in diligence, the tenant you inherited and never consolidated, the file shares still sitting behind the last project, the workloads that hit end of support next.

Most teams have never written this list down, which is precisely why each migration feels like a surprise. A single page changes the conversation from "how do we get through this one" to "what capability do we need". For each likely migration, capture the trigger and target date, source and destination, workloads and approximate scale, identity and domain dependencies, compliance constraints, business owner, confidence in scope, and current planning stage.

IT operations survey · n=943

This quarter

MEDIUM LIFT

 EFFORT

Pull identity preparation into the migration workflow, and clean up before the move

83% want identity handled inside the tool and only 11% have it there. That's the widest want-versus-reality gap in this data. If your process still involves a separate admin center or script stack for identity mapping, that's the highest-value thing to consolidate, especially if deal-driven work is in your future, where identity and permissions take two of the top three reported challenges.

Pair it with content cleanup, which 36% of teams wish they'd done more of beforehand. Sorting before you move usually beats sorting after: you move less, you move cleaner, and you don't carry a decade of permission debt into a fresh tenant. It also serves the AI readiness that 28% of teams now name as a reason to migrate.

Just don't turn it into its own project. A cleanup that runs long will hold up a deal that has a date on it, and deleting isn't always yours to choose. Retention schedules, legal holds and records rules all have a say. Sort into three piles instead. Retire what nobody needs, archive what has to be kept but not used, and move what still earns its place.

IT operations survey · n=943

This year

STRATEGIC INVESTMENT

EFFORT

SHAREGATE MIGRATE

Migrate covers the whole arc, not just the move. Surface the identities, permissions and conflicts before anything runs, move Entra ID, mailboxes, SharePoint, Teams and OneDrive in one workflow with permissions and version history intact, then validate what landed and save the plan for the next wave.

[Discover ShareGate Migrate ->](#)

Build migration as a repeatable capability with verification built in


94% of organizations migrated something in the past two years and 35% of the biggest projects were driven by organizational change, which isn't going to stop. Treating each move as a bespoke project means paying the setup cost every time and relearning the same lessons.

A capability means: a standard approach documented and owned, tooling that handles identity and exceptions rather than only the happy path, the permissions and delegation layer in scope from the start rather than discovered mid-project, pre-migration cleanup as a standing step, and post-migration verification that proves what moved and what didn't. The verification piece is the one most often skipped and the one auditors ask about first.

The teams that build this stop treating the next acquisition as an emergency. It becomes a process with known stages, explicit assumptions, and a more defensible range for duration and risk. That's the difference between a migration capability and a migration crisis.

IT operations survey · n=943

CH 06 · END



07 The operating model

CHAPTER 7 IN BRIEF

The gap in Microsoft 365 is not between teams that care and teams that do not. It is between what teams believe about their environment and what they can currently prove.

That gap has shown up in the same shape in every chapter of this report.

The gaps that exist in Microsoft 365 today

WE ASKED 1,794 IT leaders across two independent studies, an IT operations survey and an AI governance survey, what they believe about their environment and what they can currently show.

	WHAT TEAMS BELIEVE	WHAT THEY CAN PROVE
CONFIDENCE	93% are confident their AI permissions are right	29% have already had AI surface something it should not have
COST	75% have clear AI goals	51% name cost visibility as the top blocker to proving return
GOVERNANCE	63% call their governance operationalized or better	1% govern with a tool built for it
ESTATE	90% put something in the way of creating a workspace	35% put anything in the way of one outliving its use
COMPLIANCE	96% are confident M365 meets their compliance needs	34% still delay or avoid migrations over compliance
MIGRATION	83% want identity preparation inside the migration tool	11% have it there

None of those pairs is a contradiction. Each is the distance between a belief and a proof.

That isn't a criticism of the teams doing the work. Microsoft 365 behaves less like a collection of products and more like an operating environment. It changes every day as people create, share, connect, automate, migrate, and now delegate work to AI. A control that was correct at deployment can drift without one dramatic mistake ever being made.

Drew Madelung on what that means for anyone trying to hold a tenant still.

“ It's technically impossible to say with confidence that you'll have zero overshared content in your tenant. The environment never stops moving. What we do as administrators is decrease risk, not eliminate it.”




DREW MADELUNG
MICROSOFT MVP

Closing the gap is therefore not a seven-step project. It's a control loop: see what exists, decide what matters, change it safely, operate the controls continuously, and prove that the controls and investment are working. Ownership and decision rights run through every stage.

The operating loop

It runs as five motions rather than a policy document.



The people layer

An operating model needs more than a sequence of activities. It needs names against decisions. Accountability can sit with one person without pretending that one person has all the expertise.

Most teams don't have six functions to assign this to, so this isn't an org chart. It's six questions. If you can answer each one with a name, you already have the model. A name only counts if it carries the decision and the consequence. A team, a distribution list or an admin role is not an answer.

1 Who owns the governance outcome?

One name, accountable for the scorecard, the priorities and the exceptions nobody has resolved. Usually an accountable service owner.

3 Who decides what compliant means here?

Someone who sets the control requirements, reviews the riskier changes and says whether the evidence holds up.

5 Who checks the money?

Finance or procurement, seeing cost attribution, realized savings and renewal decisions before the renewal rather than at it.

2 Who keeps the lights on?

The named people doing inventory, lifecycle, monitoring, remediation and evidence, whether or not it is in their job title.

4 Who confirms a workspace still has a reason to exist?

The business or data owner. IT can see that a site is inactive; only the owner can say whether it should be.

6 Where does a cross-team exception get settled?

A standing forum that assigns the action and records who accepted the risk. A recurring meeting is fine.

IF YOU CAN'T ANSWER THEM

1 Governance is a shared intention rather than an owned outcome, and it loses every argument with a deadline.

3 Compliance arrives at audit time as an opinion instead of a standard you were already meeting.

5 The first real conversation about AI spend happens after it is already committed.

2 The work happens in the gaps between other work, and stops the week someone is busy.

4 Nothing is ever retired, because the only people looking are the ones who cannot answer the question.

6 Exceptions become permanent by default, and nobody can say who agreed to them.

The minimum version is two names: one accountable owner, and a named owner for every workspace that matters. Everything else can start as an agenda item on a meeting you already hold. The structure isn't the point. Every question having an answer you could say out loud is the point.

THE CAPABILITIES INSIDE THE LOOP

Seven moves that turn pressure into control

These build the capabilities inside the loop. They're recognizable starting points, not seven gates every organization must complete in the same order.

01 See the estate

1 OF 7

Know what you are actually operating, on-premises and cloud.

EFFORT Low to start, ongoing to maintain BEGIN This month

Only 28% of organizations are fully in the cloud, only 35% automatically retire inactive workspaces, and Power Platform is the most commonly in-scope workload when a tenant moves. Most teams are governing a surface that is larger and more varied than they picture, and growing by default, because creation is automatic and retirement is manual.

IN PRACTICE

A current inventory: workspaces, sites, guests, sharing links, apps and flows, and the on-premises workloads that are still there and why. Not a one-time audit but a maintained view, because a picture that is six months old is a picture of a different tenant.

YOU ARE DOING THIS IF

You can produce a current list of the environment, its owners, and its known exceptions without launching a project to find out.

THE COST OF NOT DOING IT

Every other control is being applied to a surface whose edges are unclear.

EVIDENCE · CH. 4

[The estate. Creation is automatic. Retirement is manual. That's the whole problem. →](#)

02 Validate the constraints

2 OF 7

Treat compliance proof as a deliverable, not a blocker.

EFFORT Medium BEGIN This quarter

87% let compliance shape migration decisions and 34% delay or avoid moves because of it. But when asked what would unlock them, the answers were artefacts rather than assurances: proof of regulatory fit (49%), control of encryption keys (45%), a validated pilot (37%).

IN PRACTICE

Write your hard constraints down precisely, then produce the evidence set that satisfies them: control mappings, a validated pilot with its documentation, residency and key-management decisions made up front. Constraints you have documented are scope. Constraints you carry in your head are blockers.

YOU ARE DOING THIS IF

The compliance team can show why a decision was approved, rejected, or excepted without rebuilding the evidence first.

THE COST OF NOT DOING IT

You join the third of organizations whose roadmap waits on questions that have no assigned owner, evidence requirement, or decision date.

EVIDENCE · CH. 5

[Compliance. Not a chapter about restriction. A chapter about evidence. →](#)

03 Move with repeatability

3 OF 7

Make migration a capability, not a one-off project.

EFFORT High first time, low every time after BEGIN This year

94% migrated something in the past two years, running 2.5 migration types across them, and 35% of the most significant projects were driven by organizational change rather than technology strategy. Acquisitions and restructures do not stop, so neither does migration. The typical project runs 3 to 6 months, on a timeline usually set by someone else.

IN PRACTICE

A standard, owned approach: cleanup before the move, identity handled inside the workflow rather than beside it (83% want this, 11% have it), tooling chosen for the exceptions rather than the happy path, and verification after the fact that proves what moved.

YOU ARE DOING THIS IF

The next acquisition or restructuring would enter known stages with explicit assumptions and a defensible range for duration and risk.

THE COST OF NOT DOING IT

You pay the setup cost every time, discover the same dependencies late, and keep meeting the 38% regret about buying tooling too late.

EVIDENCE · CH. 6

[Migration. Projects end. This doesn't — 94% moved something, 2.5 types in two years.](#)
[→](#)

WHERE SHAREGATE FITS

ShareGate Migrate runs the whole sequence in one tool. Scan and clean before the move, migrate Entra ID, mailboxes and content, verify after, then save the plan and reuse it for the next wave. Split and merge inside the same tenant for restructures. That is what "low effort every time after" looks like in practice.

[Discover ShareGate Migrate →](#)

04 Govern continuously

4 OF 7

Purpose-built control that keeps up after the environment changes.

EFFORT Medium, with a low-effort start available BEGIN This week

77% had a governance incident in the past 12 months. 30% find out about problems reactively: when a user reports one, during an audit, after an incident, or through no consistent route at all. 57% run governance on built-in tools alone and 1% on anything purpose-built, a figure that has not moved since last year.

IN PRACTICE

Close the cheap gaps first: link expiry and scheduled guest review. Then move from periodic audits to continuous monitoring with automated alerting on the signals that change fastest. The teams with proactive monitoring reported fewer incidents and more zero-incident years. Whether that is prevention or detection, both are better than finding out from the person affected.

YOU ARE DOING THIS IF

You would learn about an oversharing incident from a system rather than from a person.

THE COST OF NOT DOING IT

The 77% governance incident rate is the baseline, and for 30% of teams the incident is the alert.

EVIDENCE · CH. 3

[Governance. 77% had an incident. 1% use a tool built to prevent one. →](#)

05 Monitor AI access

5 OF 7

Prove what AI can reach, and what it shouldn't.

EFFORT Medium **BEGIN** Now: deployment already happened

Copilot is deployed almost everywhere. Of the teams running agents, 53% have less than full visibility into what those agents can access. 48% monitor AI usage continuously. And the cleanup data showed something uncomfortable: teams that had done the work reported *more* incidents, because they could see them. Teams that skipped it answer "not sure" at nearly three times the sample average.

IN PRACTICE

Map agents to data sources before scaling, monitor continuously rather than periodically, and name one accountable owner with the authority to enforce rather than advise. Only 48% have that owner today.

YOU ARE DOING THIS IF

You can show which sanctioned agents exist, what each can reach and do, who owns it, and which unknowns remain.

THE COST OF NOT DOING IT

29% report that AI has already surfaced something it should not have, and a further 8% cannot say whether it has.

EVIDENCE · CH. 1

[Confidence ≠ control. Confidence is a feeling. Incidents are facts. In 2026 they don't line up. →](#)

WHERE SHAREGATE FITS

ShareGate Protect maps what Copilot and your agents can actually reach across Teams, SharePoint, Groups and OneDrive. Not the permissions you think you set, the access that exists today. "Not sure" stops being an answer once you can show the list, and automated policies hold the line as the tenant changes.

[Discover ShareGate Protect →](#)

06 Optimize cost

6 OF 7

Turn complexity into a line item you can explain and reduce.

EFFORT Low to medium BEGIN Before your next renewal

AI licensing is now the biggest cost pressure in the estate at 41%, ahead of storage growth at 39%. The top obstacle to proving AI's return is cost visibility (51%), followed by governance complexity (47%). And 78% say governance activities directly shape their confidence to invest further in AI.

IN PRACTICE

Establish an AI cost baseline covering fixed license commitments and assigned-versus-active use; consumption-based agent, API, model and connector costs; third-party AI subscriptions; material supporting security, governance and operational costs; the cost centre or owner for each category; and the known gaps in attribution. Then take savings in order of confidence. Reclaim inactive licenses, add budgets and alerts for variable consumption, apply lifecycle policies to high-volume storage, and build the case for consolidation where fragmentation is materially increasing cost. Track realized savings, avoided future cost, and capacity released separately: reclaimed capacity does not always become cash available to fund another program.

YOU ARE DOING THIS IF

You can answer "what is AI costing us" with a figure rather than an estimate.

THE COST OF NOT DOING IT

Roughly eight in ten organizations now spend 6% or more of IT budget on AI, while about half lack the visibility needed to defend or redirect that investment confidently.

EVIDENCE · CH. 2

[The AI reckoning. Three quarters know what they want AI to do. Half can't see what it costs. →](#)

07 Prove and improve

7 OF 7

Re-measure, so confidence rests on evidence rather than feel.

EFFORT Low, recurring **BEGIN** Set the date now

This is the move that makes the other six durable, and the one least likely to arrive with a natural deadline. Many of the gaps in this report persist because control status is not measured consistently. Others reflect capacity, process, skill, or tooling constraints. Measurement does not solve those by itself, but it makes the next decision harder to avoid.

IN PRACTICE

Create a small scorecard covering the estate, AI access, governance operations, compliance evidence, migration readiness, cost, and value. For each indicator, define the source, threshold or target, owner, review cadence, trend, and required response when it moves outside tolerance. Keep a decision and action log so the review produces change rather than another report.

YOU ARE DOING THIS IF

The next review is on the calendar, the indicators have owners, and an out-of-tolerance result creates an assigned action.

THE COST OF NOT DOING IT

Controls decay quietly and confidence remains unchanged until an incident, renewal, audit, or migration forces the question.

EVIDENCE · CH. 4

[The estate. Creation is automatic. Retirement is manual. That's the whole problem. →](#)

The compounding runs both ways, and it shows up years later. Noorez Khamis has watched it.


“ The clients who did the foundational work, information architecture, security architecture, labels, DLP, years before anyone said the word AI, are in a completely different position today. That investment aged extremely well.”



NOOREZ KHAMIS
MICROSOFT MVP

The operating cadence

Not every organization needs the same frequency. The cadence should follow risk, rate of change, and regulatory need.



A note on sequencing

Do not treat the seven moves as seven gates. Establish a thin baseline across the full environment, then deepen the capability tied to the most immediate risk.

An unknown AI connection, unrestricted sharing path, or unsupported platform may require containment now while the wider inventory continues. Compliance requirements should be identified alongside discovery—not before it or after it—because teams need enough information about the data and workload to know which requirements apply. Cost should be visible from the start and revisited as use scales. After every major deployment, migration, or policy change, run the loop again.

For organizations that have already deployed Copilot without mapping its knowledge sources and actions, move five isn't fifth. It's an immediate risk assessment inside the See motion. For organizations whose roadmap repeatedly stalls on compliance, move two should deepen early. The model is consistent; the emphasis changes with the constraint.

The forcing function has to be yours. The SharePoint end-of-support data showed what a date does to work that would otherwise drift CH 4 **The estate** →. Nobody will send an end-of-support notice for permission sprawl, stale workspaces, unresolved AI connections, or an ageing exception. The operating model has to create those dates.

“ Whether this is done with Microsoft tools, a third-party tool, or a combination, the goal has to be visibility and accountability at scale.”



LAURA ROGERS
MICROSOFT MVP

Where to start

Most teams already have parts of this model. Start where the evidence is weakest or the consequence of delay is highest.

IF THIS IS TRUE OF YOU		START HERE
You cannot produce a current estate, agent, or cost baseline	1	then the minimum registers for moves 5 and 6
AI is active but its knowledge sources or actions are unclear	5	Contain the immediate risk first, then move 5
Users are still the first people to report governance problems	4	connect the highest-risk signals to owners and runbooks
Compliance repeatedly delays decisions	2	define the requirement, evidence, owner, and decision date
Migrations keep arriving as emergencies	3	build the demand pipeline and playbook
AI cost or return cannot be explained	6	establish a cost and outcome baseline
Controls exist but no one can show whether posture improved	7	put the scorecard and next review on the calendar

If you're not sure which description fits, that uncertainty is itself useful evidence. Begin with move one, keep the first baseline thin, and let what you find determine where to go deeper.

The teams that end up in the best shape aren't the ones who bought the most or planned the longest. They're the ones who picked a date, measured honestly, and were willing to be embarrassed by what they found. The organizations that tend to make the most progress start by discovering something uncomfortable. The ones that stay stuck are the ones whose confidence never gets tested.

Buying something is rarely the hard part. Microsoft MVP Marc Anderson on what usually happens next:

“ When people get the tools, they think they've solved the problem, and they don't necessarily use them. It's not that the tools are hard to use. It's that they have to become part of your daily routine, and I don't know that they always do.”




MARC ANDERSON
MICROSOFT MVP

Find your starting point

The M365 Governance Index scores your estate and your AI governance against the organizations in these two studies, across the dimensions this report has been circling. It takes about five minutes and returns the two or three moves that would shift you most. This is what move seven looks like the first time you do it.

[Get your score →](#)



CONTRIBUTORS

Microsoft MVPs shaped every chapter.

A survey of 1,700+ IT pros tells you what's happening. It doesn't always tell you why. So before we published anything, we took the findings to many Microsoft MVPs, people who have spent years inside hundreds of tenants. We asked them what sits behind each number. Which findings matched what they see on the ground. Which ones surprised them. And which ones we should be careful about claiming.

We're grateful to all of them for their time, expertise, and candor. Their commentary runs through the whole report.



MICROSOFT MVP

Richard Harbridge

Principal Industry Advisor
at ShareGate



MICROSOFT MVP

Eric Overfield

CIO at Creospark



MICROSOFT MVP

Jasper Oosterveld

Data Security Consultant at Rubicon
Cloud Advisor



MICROSOFT MVP

Drew Madelung

MS Cloud Solutions Architect at
Protiviti



MICROSOFT MVP

David Drever

Data Protection Specialist at Protiviti



MICROSOFT MVP

Laura Rogers

CEO at IW Mentor



MICROSOFT MVP

Mike Maadarani

CTO at Taiga-Ai



MICROSOFT MVP

Joanne Klein

Founder of NexNovus Consulting



MICROSOFT MVP

Noorez Khamis

CTO and Co-Founder of Creospark



MICROSOFT MVP

Marc Anderson

President of Sympraxis Consulting



MICROSOFT MVP

Susan Hanley

President of Susan Hanley LLC

ABOUT THE RESEARCH

How we ran the numbers

This report draws on two independent 2026 surveys of IT professionals and leaders running Microsoft 365. Both are respondent-level and unweighted, and we never blend them into a single figure — every statistic is attributed to the survey it came from. This page is the full picture: who we surveyed, how to read the data, and how to cite it.

The two surveys at a glance

	IT operations survey	AI governance survey
Respondents	943 IT professionals	851 IT leaders
Seniority	Manager 49% · Director 39% · VP 10%	IT leadership ~73% (rest security, data-governance, workplace, compliance)
Sector mix	Broad (Tech 20%, Manufacturing, Financial Services, and more)	Skews Technology / Software (~59%)
Countries	US, Canada, UK, Germany, France, Japan, Australia	US, UK, France, Canada, Germany, Netherlands, Ireland
Questions	46	35
Data granularity	Respondent-level	Respondent-level
Fielded via	Cint	Centiment
Field dates	May & June 2026	March 2026

Who we surveyed

Both samples are respondent-level, so every figure in the report can be cut by the dimensions below. Sample sizes are read straight from the dataset that powers the charts.

IT operations survey — by country (n=943)

Country	n	% of n
United States	270	29%
United Kingdom	135	14%
Canada	135	14%
Germany	113	12%
Japan	100	11%
Australia	100	11%
France	90	10%

Also cuts by **organization size** (5 bands) and **industry** (10, including Financial Services, n=138).

IT operations survey — org size & industry breakdown

By organization size

Organization size	n	% of n
500–999	249	26%
1,000–2,999	230	24%
3,000+	220	23%
250–499	129	14%
1–249	115	12%

By industry

Industry	n	% of n
Technology / Software	190	20%
Manufacturing	164	17%
Financial Services / Insurance	138	15%
Retail / E-commerce	109	12%

Industry	n	% of n
Healthcare & Life Sciences	101	11%
Education directional	70	7%
Professional Services directional	57	6%
Energy / Utilities directional	48	5%
Public Sector / Government directional	41	4%
Other (please specify) directional	25	3%

AI governance survey — by country (n=851)

Country	n	% of n
United States	219	26%
United Kingdom	204	24%
France	108	13%
Canada	106	12%
Germany	104	12%
Netherlands directional	56	7%
Ireland directional	54	6%

Also cuts by **organization size** (7 bands), **industry** (10), and **role**. The sample composition leans toward technology-sector organizations (Technology / Software ≈ 59%) and European markets; read its results as that audience, not the whole market.

AI governance survey — size, industry & role breakdown

By organization size

Organization size	n	% of n
500-999	192	23%
1,000-2,499	160	19%
2,500-4,999	148	17%
100-249	96	11%
250-499	92	11%
5,000-9,999	91	11%

Organization size	n	% of n
Over 10,000 directional	72	8%

By industry

Industry	n	% of n
Technology / Software	500	59%
Manufacturing	105	12%
Financial Services directional	56	7%
Professional Services directional	49	6%
Retail / E-commerce directional	37	4%
Healthcare & Life Sciences directional	29	3%
Public Sector / Government directional	27	3%
Energy / Utilities directional	20	2%
Education directional	15	2%
Other (key in) directional	13	2%

By role

Role	n	% of n
IT leadership (e.g., IT director, infrastructure director, head of IT)	620	73%
Data governance or data management leadership (e.g., data governance manager, data officer)	79	9%
Security leadership (e.g., CISO, security director, risk or cyber leader)	76	9%
Digital workplace / collaboration leadership (e.g., M365 lead, head of collaboration, workplace technology leader) directional	52	6%
Compliance / privacy leadership (e.g., compliance manager, privacy officer) directional	13	2%
Other (key in) directional	11	1%

Cells below n=75 (Netherlands n=56, Ireland n=54 in the AI governance survey, and some smaller industry cells) are directional and labeled as such wherever they appear.

How to read the data

- All figures are **self-reported** and represent percentages of respondents who answered a given question.

- Figures are **unweighted**: each percentage is a straight count of respondents, not adjusted to a target population. Verified against the row data — every percentage equals $\text{count} \div n$.
- **Multi-select** questions sum to more than 100% by design.
- The two surveys have **different samples and are never combined** into a single figure. Every statistic is attributed to its survey.
- AI governance survey **leans toward technology-sector and European** organizations, so read its results as that audience, not the whole market.

Year-over-year comparisons

Where noted, results are compared to ShareGate's 2025 report, *M365 Decoded*. Eight comparisons are closely matched in question wording (labeled *clean*); two are *directional*, where wording or base differs, and are flagged as trends rather than exact deltas.

The benchmark (M365 Governance Index)

The companion benchmark scores readers on two pillars, each against its own survey. Both cohorts are **real respondent distributions** — the estate cohort is all 943 IT operations survey respondents scored on five estate dimensions; the AI-governance cohort is all 851 AI governance survey respondents scored on five AI-governance dimensions. The two are benchmarked separately and never blended. Full definitions live with the benchmark itself.

How to cite

- IT operations stats: **"ShareGate 2026 IT operations survey (n=943)."**
- AI governance stats: **"ShareGate 2026 AI governance survey (n=851)."**
- Never cite a combined figure (1794 respondents) — the two surveys are separate samples.

Index by topic

Seventeen topics, tagged section by section. A topic that runs across several chapters is listed once here with every page it appears on.

access & sharing 6, 17, 20, 71, 74, 89, 137, 139, 150
Permissions, anonymous links, guest access, and who can reach what.

AI adoption 14, 26, 34, 36, 44, 48, 130
Where Copilot and agents are actually in use, what teams expect from them, and what they cost.

AI exposure 4, 8, 17, 20, 32, 58, 89, 117
What AI can reach that it shouldn't, and the incidents that follow from it.

automation 8, 44, 71, 77, 89, 91, 148, 150, 158
Which governance work runs itself, and which still waits for a person.

compliance & audit 6, 111, 113, 124, 130
Frameworks, audit evidence, and the artefacts that unblock a roadmap.

cost & licensing 34, 36, 38, 40, 43, 44, 48, 50, 56, 95
AI spend, licence waste, storage growth, and the cost of standing still.

data residency 113, 117
Digital sovereignty, where data physically sits, and what that gates.

end of support 85, 87, 100, 130
The SharePoint deadline, the cohort that stayed, and what held them back.

estate & hybrid 50, 79, 81, 84, 85, 87, 95, 100, 113, 132
What teams run and where — cloud, on-premises, and the mix in between.

geography 10, 48, 117
Where the numbers diverge by country — confidence, spend, and sovereignty pressure.

incidents 4, 23, 58, 60, 81
Exposure events — how often, how they surface, and how teams find out.

lifecycle 23, 50, 71, 74, 77, 84, 91, 95, 100, 139, 148, 150
Provisioning, sprawl, archiving, retirement — the parts that are still manual.

maturity

..... 4, 6, 8, 10, 40, 44, 65, 68, 81, 147, 159, 160

What teams say about their own governance, and where the data disagrees.

migration

..... 50, 85, 87, 113, 124, 125, 129, 130, 132, 133, 137, 139, 143, 146

Moving content and tenants: why, how long, what breaks, and what teams would change.

monitoring

..... 8, 20, 23, 32, 38, 40, 43, 56, 60, 63, 68, 77, 111, 147, 148, 150, 158, 160

What teams can see: detection, alerting, and the difference continuous monitoring makes.

ownership

..... 26, 32, 65, 143, 149, 158, 159

Who decides, who does the work, and where governance actually sits.

tooling

..... 43, 63, 65, 111, 124, 133, 143, 146

What teams govern and migrate with, and where the built-in tools stop.

THANK YOU FOR READING

ShareGate

Index by topic · State of M365 Report

170