Futurum RESEARCH

Red Hat



# Sovereignty Is Control, Not Location

The four capabilities that decide whether digital sovereignty is real, and how to build them.

OCTOBER 2026

# Table of Contents


Executive Summary ..... 02

1. Funded, Governed and Not Yet Built ..... 03

2. The Workloads Are Not Leaving ..... 12

3. Worry Without the Ability to Act ..... 16

4. What Buyers Will Actually Pay For ..... 23

5. AI Turns Sovereignty Into a Purchase Decision ..... 28

6. What To Do Next ..... 31

Methodology ..... 33



# Executive Summary

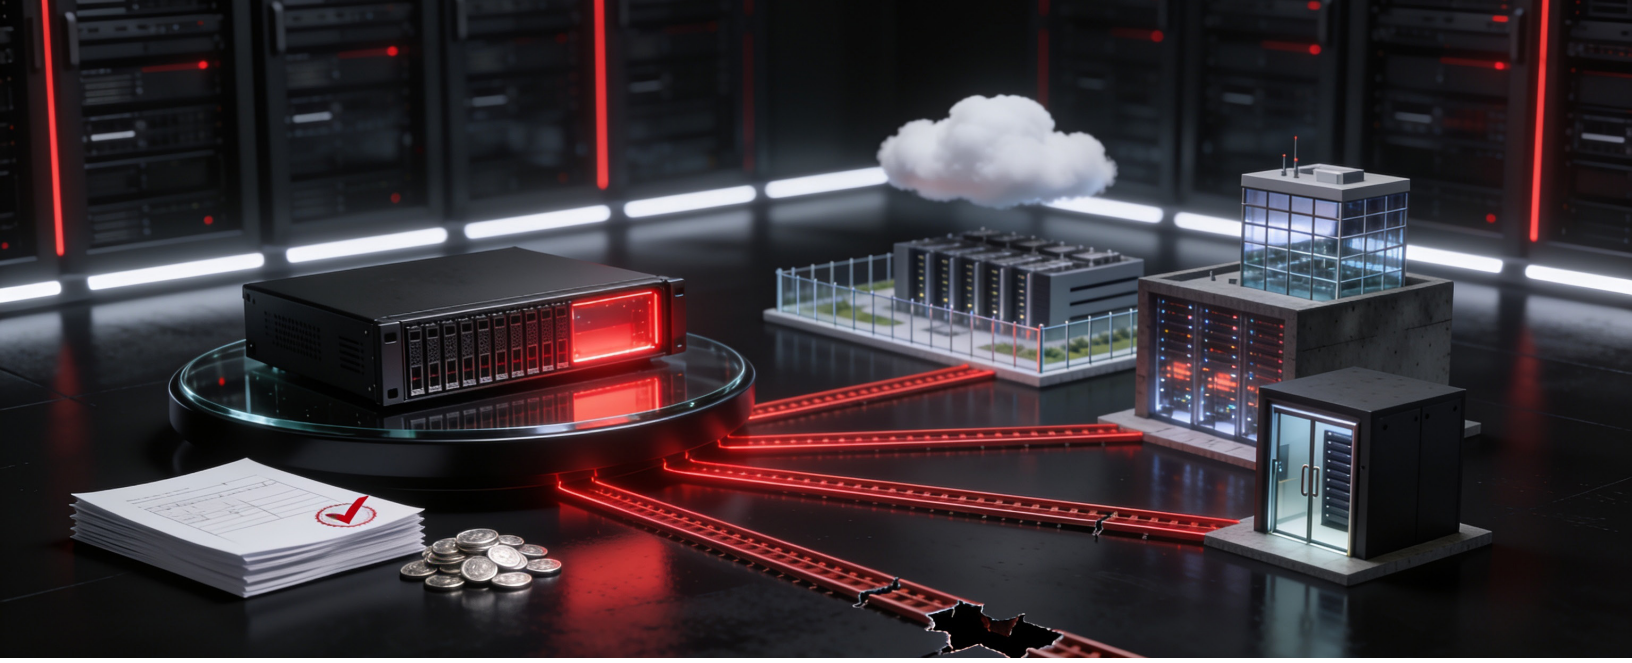
Digital sovereignty is no longer about good intentions; it is about whether organizations can actually put it into practice – and right now, almost nobody is prepared. Organizations are putting money toward sovereign technology, assigning managers to lead the effort, and setting up rules. However, very few can prove they could keep key services online if a cloud vendor, software supplier, or political crisis cut off their access. This gap between what companies plan to do and what they can actually achieve is now the main challenge in digital sovereignty.

**The concept spans several core priorities at once:** data residency and its governing laws; operational continuity during vendor or government disruptions; and control over where AI models originate and run. Organizations weigh these three differently by region and by industry. Most have already worked out where their data sits; what they are working on now is keeping the business running under their own control, and controlling the AI they use.

At the same time, production workloads continue to leverage US-headquartered hyperscalers. In fact, hyperscalers are the only environment expected to gain share over the next decade. That concentration does not make sovereignty less relevant. It makes operational readiness more important.

**The sovereignty readiness gap is the divide between strategic intent and operational reality**. Closing it requires shifting focus from where data resides to whether the enterprise can maintain independent control. Four capabilities decide that: a tested contingency plan for losing a critical vendor, a plan for losing support for an open-source dependency, a sovereignty risk framework applied to every supplier, and the ability to move a critical workload off a single cloud provider within 30 days. Only 2.9% of organizations hold all four. What buyers are buying to close the gap is narrower, i.e., a second source of support, commercially supported open source, and cross-environment portability.

Digital sovereignty now commands dedicated budget and executive governance, yet only one in seven organizations has progressed past planning into daily operations. Risk mitigation drives the investment case: operational resilience, regulatory compliance, and data control consistently win internal funding over growth-oriented outcomes. Strategy ownership is split almost evenly between infrastructure and security leadership. The CISO or chief risk officer owns sovereignty strategy in 28.6% of organizations. A message built only for platform engineering teams therefore never reaches the person who owns the decision in close to three in ten cases.



**The workloads are not leaving**. Measured as a share of production workloads, sovereign and in-country offerings sit in roughly the same place in ten years as they do today. While private cloud, on-premises infrastructure, regional providers, and managed hosting all lose ground, US hyperscalers are the only environment that gains. That does not mean sovereignty has failed. It means organizations are pursuing sovereignty inside the estate they already run, rather than replacing that estate wholesale. They are seeking more control over risk, support, governance, and workload movement while continuing to use hyperscale cloud.

The paradox is that consolidation can make the readiness gap more consequential. Organizations most concerned about dependence on US-headquartered providers still expect US hyperscaler use to grow, and those that say they are bringing work back in-house project the steepest increase in US hyperscaler share. In practice, many buyers are treating contingency as insurance around a concentrated operating model rather than building a genuinely independent alternative.

The ten-year projection reflects the options buyers can see today. Futurum's three- to five-year market outlook is less linear. Regulation is becoming more region-specific and enforceable; sovereign and in-country cloud offerings are becoming more operationally distinct; and local providers and telecom operators are expanding their role in cloud and AI supply. The likely result is a more fragmented set of deployment and supplier options than respondents' long-term forecasts imply. Leaders should therefore treat a ten-year allocation to any provider model as a strategic assumption to test regularly, rather than as a plan that can be set once and left unchanged.

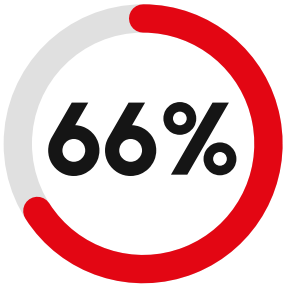
Worry does not translate into the ability to act. Nearly nine in ten organizations fear being cut off by a supplier or a government, but the risk named most is being unable to leave, not foreign government access to data. Few hold the full set of capabilities needed to respond, as current strategies prioritize compliance and planning over operational readiness. Moving a critical workload takes most organizations months, and the largest, most exposed enterprises are the least able to do it quickly. That reframes what counts as protection: if the risk buyers name first is being unable to leave, then the relevant property of any technology is whether somebody other than its original supplier can keep it running.

**Readiness is bought as support and portability**. Asked what they need to be ready, buyers rate a second source of support and maintenance above every other requirement, including data location. Open source underpins most sovereignty strategies and has become the preferred way to avoid lock-in at the software layer, but it does not remove the need for vendor support: organizations are highly concerned about running unsupported packages and place a premium on a credible alternative source of maintenance. Notably, organizations that already hold a support contract are the ones looking hardest for a second maintenance source – support is being treated as a contingency requirement, not a procurement formality.

Portability remains unmet. Sovereignty has become decisive in a large share of infrastructure renewals and is now considered in most. But buyers who have made it decisive are no more able to move a workload than anyone else, and reliability – not data location – is what they say makes a vendor sovereignty-ready. Portability is the requirement buyers name and have not yet bought.

**Recommendations:** Closing the gap starts with testing the exit against a clock, and writing that exit into the contract at the next renewal, with a second source of support identified in advance. From there, organizations should prefer open, commercially supported software wherever being trapped is the risk; build an inventory that reaches the encryption keys and open-source components underneath; assess every supplier, and make the assessment include a way out; settle who owns sovereignty and report it to the board as a risk; revisit the decision to consolidate every year; and, especially for large organizations, don't delay while risks are growing.

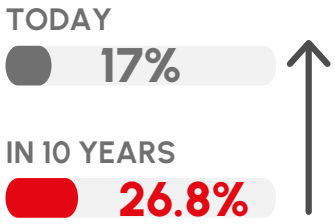
## KEY METRICS




call digital sovereignty a critical priority; 0 of 1,940 called it irrelevant

**2.9%**

hold all four capabilities sovereignty requires




share of production workloads on US hyperscalers, today (17%) vs. in ten years




call an alternative source of support critical if their primary vendor could no longer serve them





# 1. Funded, Governed, and Not Yet Built

**Nearly every technology leader calls sovereignty important and most have money set aside for it; far fewer have built anything yet**

Whether digital sovereignty matters is no longer in question. Two-thirds of technology leaders call it a critical priority, and none of the 1,940 surveyed called it irrelevant. The money has followed. Most expect a major investment within a year, and spending is rising in every region. Delivery, however, is not growing: nearly half of organizations have not moved past planning, and only one in seven actually have sovereignty embedded in operations.

**+8.4%**

average expected change in sovereignty spend over the next 24 months

**↑ RISING IN EVERY REGION**



Latin America
**+11.1%**



North America
**+8.1%**



Asia-Pacific
**+8.0%**



EMEA
**+6.6%**

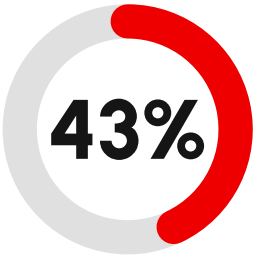
Source: Q4 (expected change in sovereignty spend over 24 months), mean of category midpoints; regions weighted equally.

## Regulation puts sovereignty on the agenda and other pressures fund it

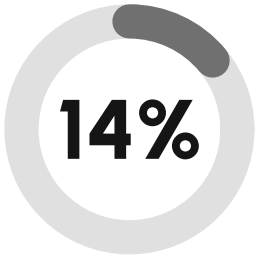
Regulation is a major factor for putting sovereignty on the agenda. Nearly half say the regulatory environment in their region drives their roadmap to a great extent and, similarly, nearly half describe the cultural pressure around it as high or very high, usually a board-level concern beyond IT. However, regulation does not, on its own, get the work funded.

What organizations expect back is protection: resilience, compliance, and control over their data lead the list, and every growth outcome trails every protective one (Figure 1). Sovereignty is primarily being funded the way General Data Protection Regulation (GDPR) compliance was funded, i.e., as a cost of doing business. The trend is visible across industries, with the manufacturing and energy sectors leading in budget increases (43% planning to boost spending by more than 10%), ahead of the public sector and defense at 34%. Ultimately, the most successful funding justifications focus on maintaining continuity and control rather than driving expansion.

### BUYERS STILL SEE UPSIDE




call sovereignty a significant competitive differentiator




call it a compliance cost only

Source: Q22 (sovereignty as competitive differentiator vs. compliance cost; 3-point scale).

*Figure 1. Organizations Expect Protection From Sovereignty, and Growth Trails*




Note. Percentage selecting each outcome among their top two. Global figures weight the four regions equally (n = 1,940).
Source: Q21 (outcomes most expected from sovereignty investment).

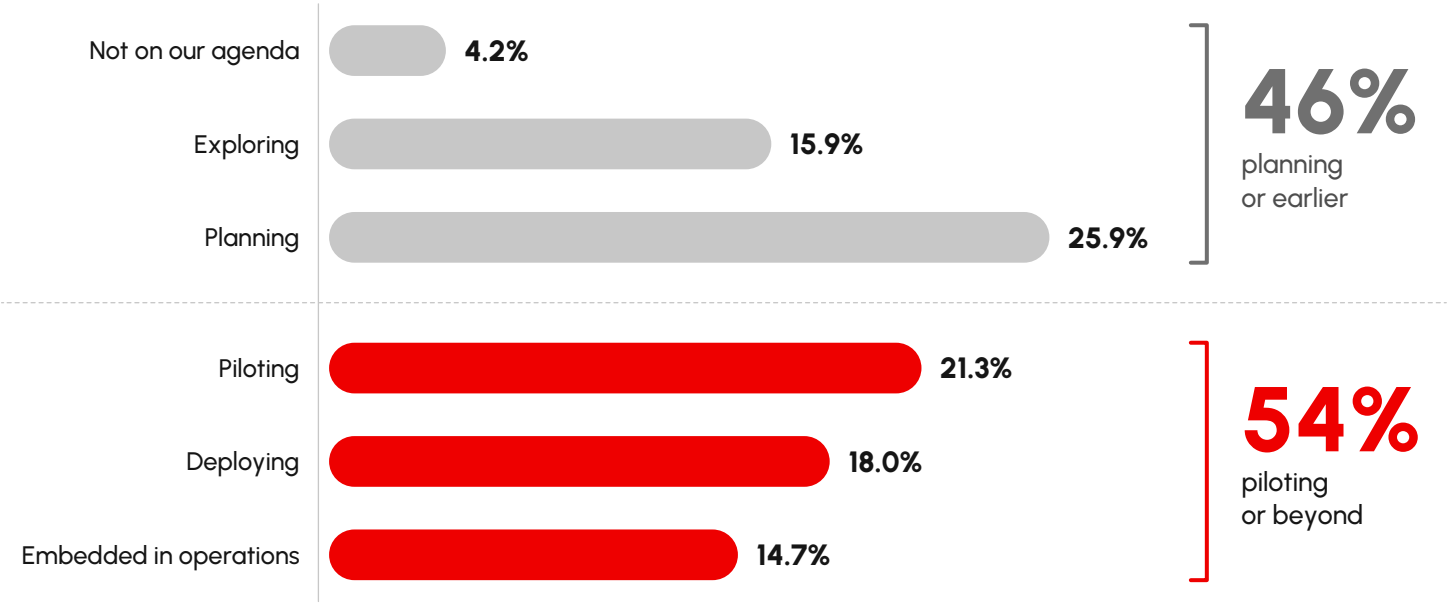
## Every additional jurisdiction adds complexity

Sovereignty rarely comes down to one country's rules. Most large organizations operate in several countries, often in several regions, and every jurisdiction they add brings its own set of rules, its own conditions for market access, and often its own infrastructure and vendor landscape. Trends run by region, but the countries inside a region can still need very different things; Canada and the United States are the clear example. Local jurisdictions inside a country can add rules of their own. The more distinct sets of rules and requirements an organization has to work around, the more complex sovereignty becomes, and regulation is only part of it. Which markets an organization can serve, what infrastructure exists there, and which vendors can operate there all change from one border to the next. Each added jurisdiction is one more place where vendors must be evaluated, business must be adapted, and readiness must be proven.

## Planning is further along than delivery

Nearly half of organizations have a sovereignty mandate and a budget, but have not yet run anything (Figure 2). Organizations with sovereignty embedded in how systems are chosen, contracted, and run are the target state the rest of the market is working toward, but the gap is significant, and progress is slow.

*Figure 2. The Sovereignty Maturity Ladder: Where Organizations Stand on Sovereignty*




Note. Percentage of respondents at each stage. Global figures weight the four regions equally (n = 1,940).
Source: Q1 (current stance on digital sovereignty).

Program maturity brings governance and money with it: organizations deploying sovereignty or with sovereignty embedded are more likely to expect a major investment and to increase spend. For a leader, the useful question is which program maturity stage an organization is in. Organizations still planning should know that their current planning will dictate how fast they overcome this gap. Organizations piloting or deploying should check whether their pilots include a tested exit and a second source of support. Section 3 demonstrates how few pilots include an exit plan.

## MATURE ORGANIZATIONS ARE ACCELERATING

● Mature organizations ● All others



**81% vs. 60%**

expect a major sovereignty investment within 12 months: mature vs. all other organizations



**72% vs. 62%**

are increasing sovereignty spend over 24 months: mature vs. all others

Source: Q25 and Q4 by Q1; mature = deploying or embedded in operations; unweighted.

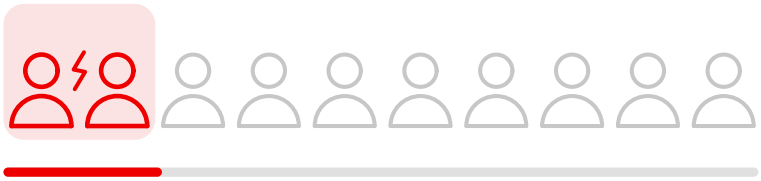
### Ownership is split and security holds it nearly as often as infrastructure

Sovereignty is a board-level decision about which vendor dependencies an organization accepts. Some dependence is unavoidable, but a dependency with no tested way to migrate is different. For critical systems and customer-facing services it is one the board should approve knowingly. Those systems cannot degrade because a supplier or government changes terms, and identifying them is leadership's job, beyond IT alone.

Ownership in practice is less settled than that. While nearly every respondent indicated a primary owner exists for digital sovereignty, one in five report active internal disagreement about direction or ownership. Where it is settled, the CIO or CTO holds it most often and the CISO or chief risk officer nearly as often, so security owns sovereignty almost as often as infrastructure does (Figure 3). Interestingly, C-level survey respondents were significantly more likely to attribute sovereignty ownership to the CIO or CTO (45%) than any other respondent levels from managers and architects through SVPs (27%–37%), indicating that senior executives disagree with the rest of their business on who actually owns sovereignty responsibility.

**1 in 5**

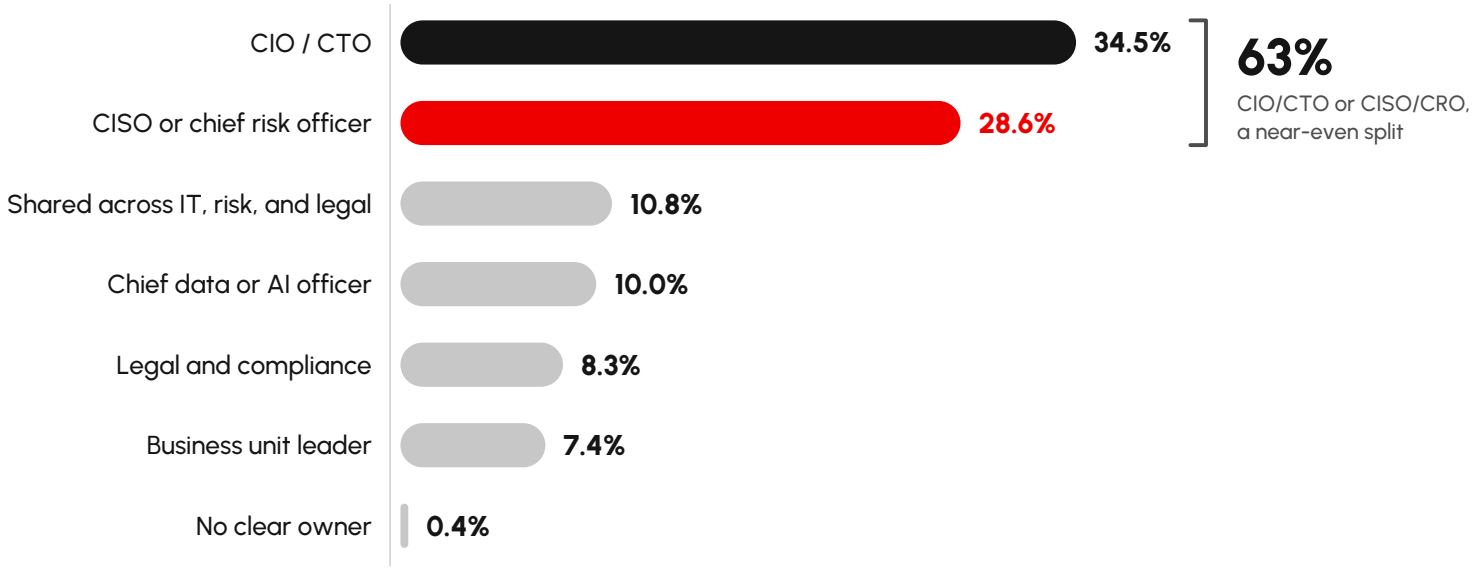
(21%) report active internal disagreement on direction or ownership



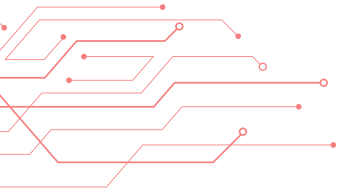
**21%** of respondents

Source: Q27a (internal disagreement or political tension about sovereignty priorities).

*Figure 3. Ownership of Sovereignty Is Split Between Infrastructure and Security Leaders*




Note. Percentage naming each role as the owner of digital sovereignty strategy. Global figures weight the four regions equally (n = 1,940). Source: Q3.



*"Senior Information Risk Owner responsibility for location of data and control of assets from a tech point of view is generally in one person... it's very clear that's my responsibility to advise the board on."*

– CDIO, Healthcare (Europe)

## Maturity splits sharply by region, but not industry or size

North America is furthest along, building sovereignty inside a familiar jurisdiction, with only a third of organizations (33%) not yet piloting sovereignty programs (Figure 4). Asia-Pacific and Latin America both sit in the middle (42%–47% not yet started with pilots).

North America's lead stems from industry regulations and supplier concentration rather than region-wide rules. US technology leaders view heavy reliance on US cloud providers as a concentration risk rather than a legal issue. Consequently, they focus on building resilience within their home region rather than having any plans to leave it. While North America leads in program stage (67.1% piloting or beyond) and demand for secondary support (72%), only 21.5% can relocate a workload within 30 days, reflecting progress in planning rather than execution capability.

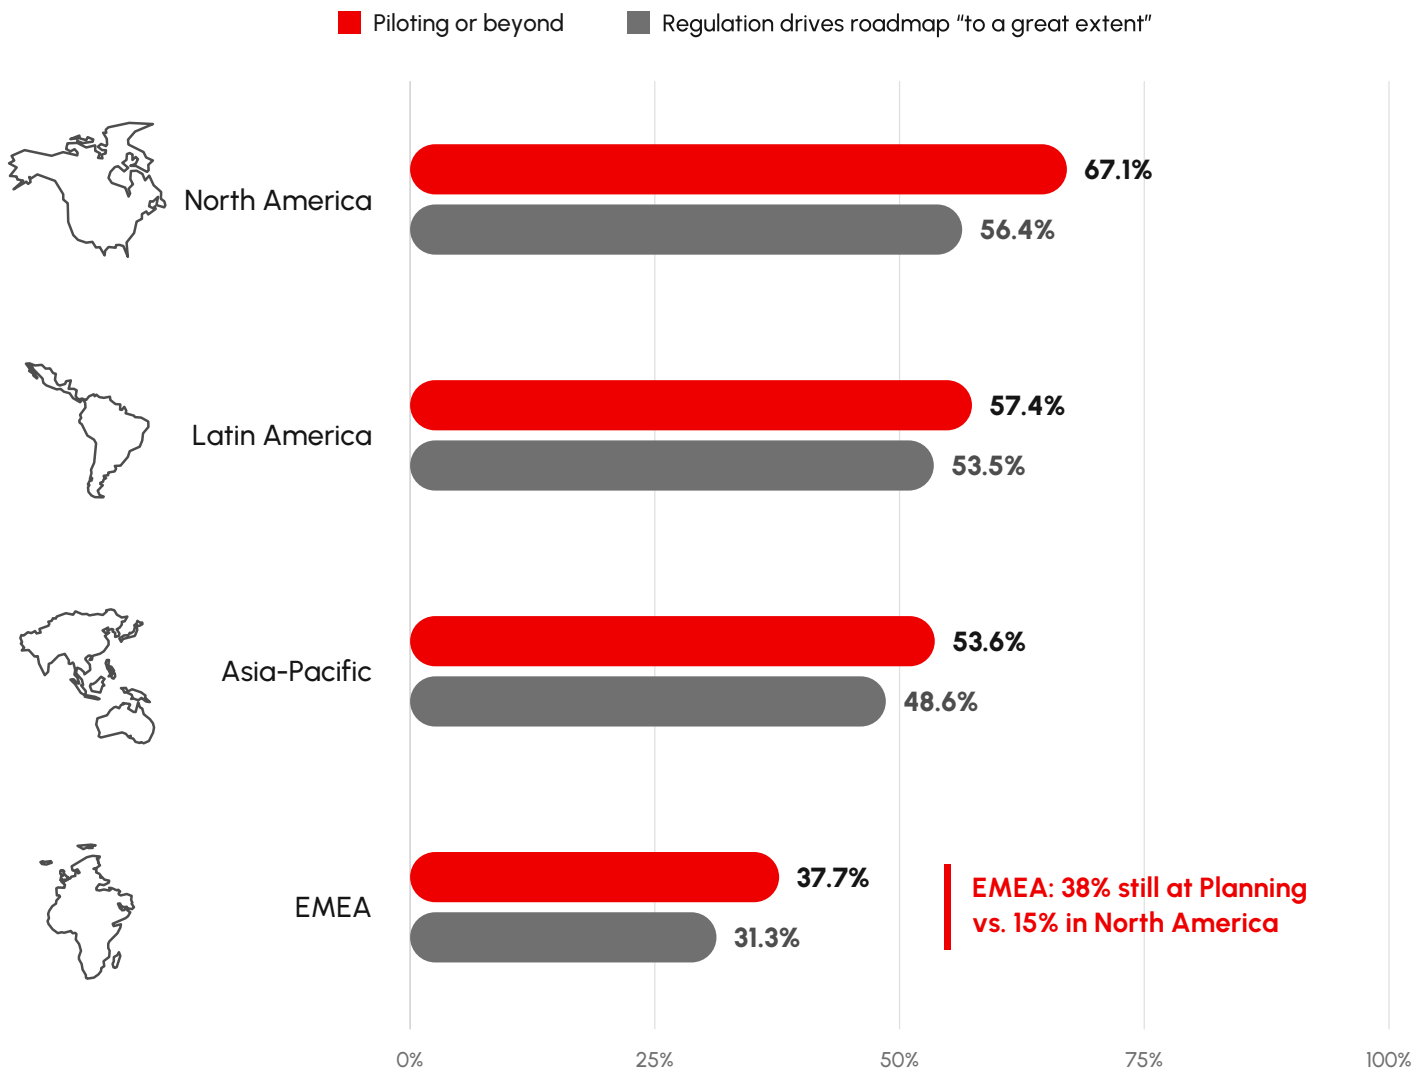
EMEA's position reflects a different sequence of priorities. The region tackled data residency early through GDPR, treating location as a resolved issue. However, operational continuity remains largely unaddressed. As a result, EMEA scores lower on making sovereignty decisive at renewal (31.7%) and requiring secondary support (52.1%), even while sharing high concern levels. Broad local definitions also make pilots larger undertakings. EMEA is not lagging in priority, but in solving operational continuity.

Within EMEA, Mainland Europe is furthest behind: only a third have started piloting sovereignty programs, against more than four in ten in the United Kingdom and nearly half in the Middle East. Across markets, maturity ranges from 22% in Italy and 29% in the Netherlands, Belgium, and Luxembourg to 70% in Hong Kong and 69% in Singapore (Figure 5).

These figures reflect self-reported answers, but local definitions of a "pilot" vary. For instance, a US firm running on-premises systems alongside a domestic hyperscaler may consider that a sovereignty pilot, whereas a Spanish or Italian firm facing stricter residency and supply-chain rules might not. These variations highlight how local expectations differ and demonstrate that digital sovereignty must be defined locally to be effective.

Interestingly, neither the respondents' industry, role, nor organization size have any consistent, meaningful effect on sovereignty maturity as reported in the survey.

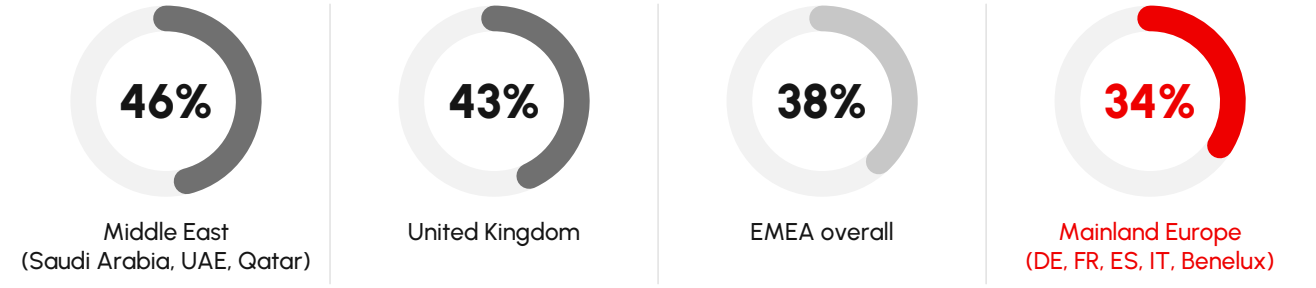
*Figure 4. Regional Maturity: North America Leads, EMEA Lags*



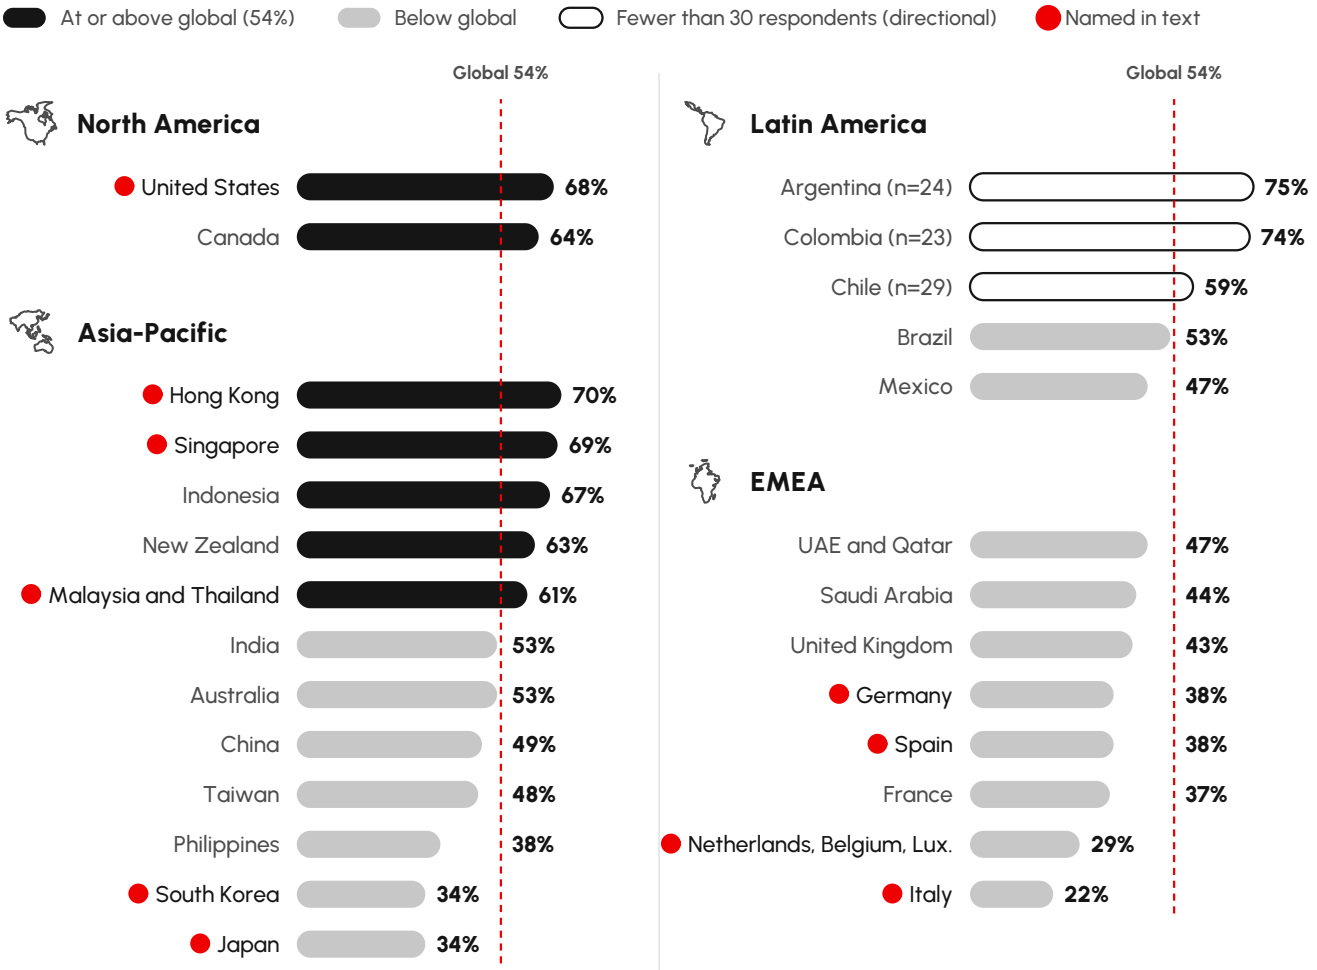

Note. Percentage piloting or beyond (Q1) and percentage saying regulation drives their roadmap to a great extent (Q9), by region. Regional figures are unweighted: North America n = 289, EMEA n = 470, Asia-Pacific n = 979, Latin America n = 202. Source: Q1, Q9.

*Figure 5. Sovereignty Maturity Varies Sharply by Subregion and Market*

## A Inside EMEA: Mainland Europe is furthest behind




## B Across markets: from about one in five in Italy to seven in ten in Hong Kong




Note. Percentage piloting or beyond, by EMEA subregion (Panel A) and by market (Panel B). Mainland Europe = Germany, France, Spain, Italy, and the Netherlands, Belgium, and Luxembourg (n = 299); United Kingdom n = 68; Middle East = Saudi Arabia, UAE, and Qatar (n = 101). Several markets were fielded as combined groups (Netherlands/Belgium/Luxembourg, UAE/Qatar, Malaysia/Thailand) and cannot be separated. Market figures are unweighted; markets with fewer than 30 respondents (Argentina, Colombia, Chile) are directional only. The dashed line marks the global figure (54%). Source: Q1 by S1 (country).

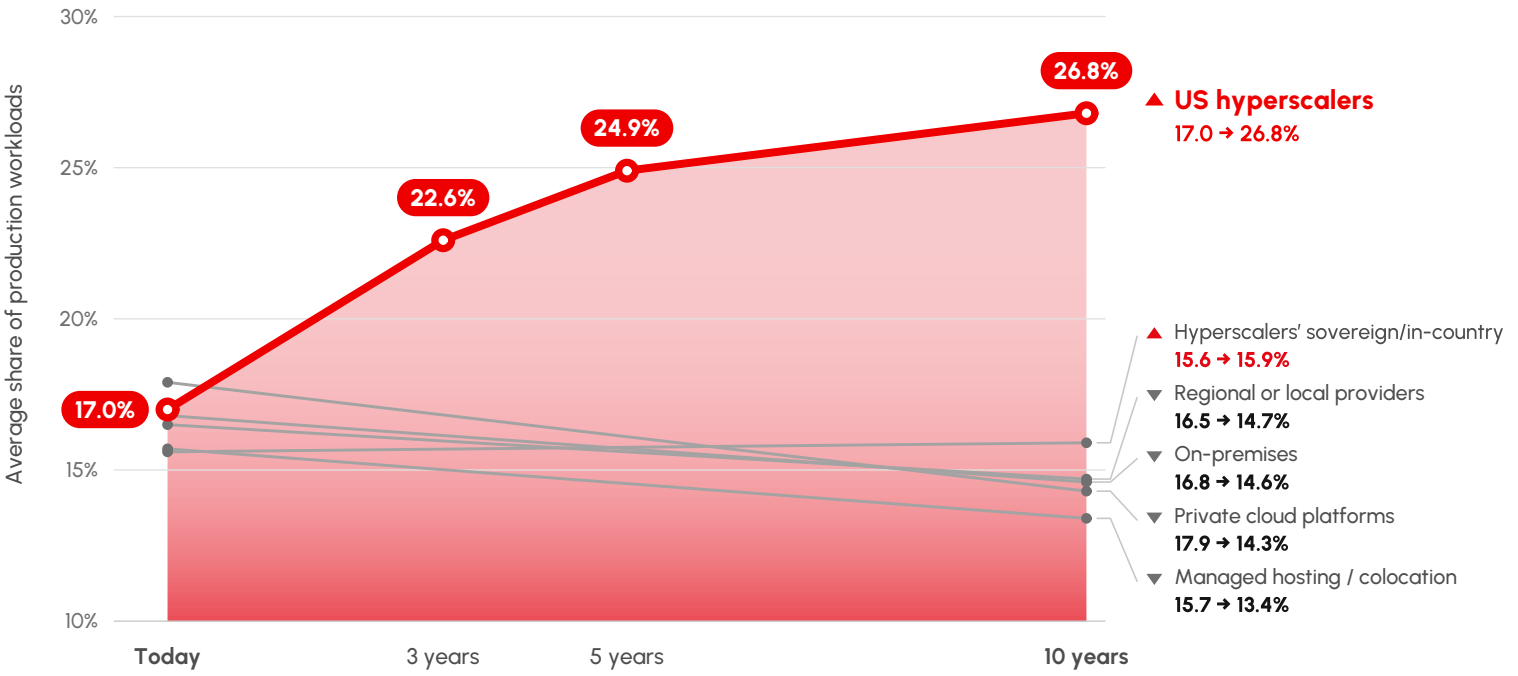
**What to Do About It: Find out which stage you are actually at, not which you report. If your program is funded but has never run anything, you are in the 46% at planning or earlier, and the next decision is which single workload to test first.**

# 2. The Workloads Are Not Leaving

### Where production workloads run today and in ten years, and what buyers are changing about them

If sovereignty were moving workloads to sovereign alternatives, the ten-year picture would show sovereign and regional providers gaining share. The largest US cloud providers are the only environment that grows across the decade; every other environment (sovereign and in-country hyperscaler offerings, regional and local providers, on-premises infrastructure, private cloud platforms, and managed hosting or colocation) loses share (Figure 6). The pattern is pervasively consistent across respondent roles, regions, industries, sizes, and sovereignty maturity. For an organization building a sovereignty plan, this settles where the workloads will sit in ten years. What remains open is the terms attached to that hosting: the contract and the architecture around it.

*Figure 6. Where Production Workloads Run, Today and in Three, Five, and Ten Years*




Note. Mean share of production workloads by environment, as reported today and projected at three, five, and ten years. Global figures weight the four regions equally (n = 1,940). Source: Q18 (share today), Q18a (expected share at three, five, and ten years).

*"[We aren't moving] on-prem, but we are diversifying…we're working with another hosting company based in France for some of our workloads…it became a lot more important criteria [in RFIs and RFPs]."*

– CIO, retail organization (France)

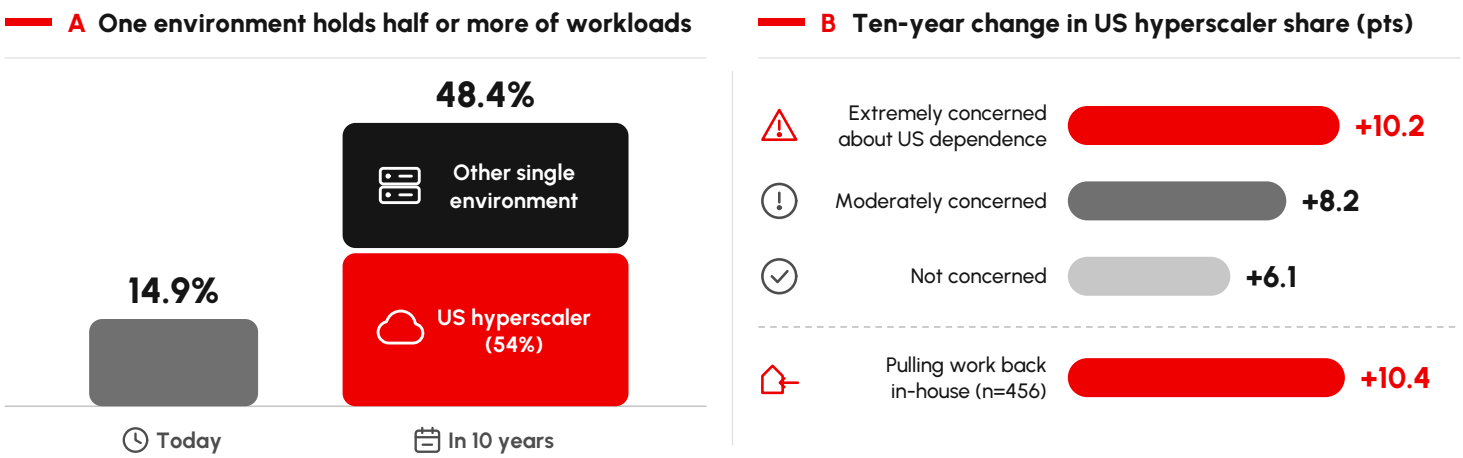
## Organizations worried about depending on US providers are concentrating the fastest

Concentration is rising alongside that shift: within ten years, many organizations expect a single environment to hold most of their production workloads, and for a little over half of them, that environment will be a US hyperscaler (Figure 7). Concentration compounds the readiness problem: the more an organization's workloads sit in one place, the larger any move would be and the fewer places it could go. Organizations most concerned about depending on US-headquartered providers, and organizations pulling work back in-house, expect the steepest rise in US hyperscaler share of anyone in the study. These buyers are staying with the largest providers because that is where new capabilities arrive first, treating a fallback plan as insurance while governing from inside an already concentrated position. Doing that safely requires a tested exit and a second source of support beforehand.

A US hyperscaler with a sovereign or in-country offering is the joint most-cited provider model at 30.4%, level with private cloud platforms and well ahead of regional, local, or sovereign providers at 23.4%. Buyers are not planning to leave the largest providers in order to become sovereign; they are planning to buy sovereignty from them, or to place a portable platform layer on top. In interviews, several executives described continuing with hyperscalers as a reluctant conclusion rather than a preference; the capability, the pace of new services, and the absence of a like-for-like alternative made the decision for them. That is acceptance, rather than enthusiasm, and it drives a different set of planning assumptions.

Notably, even holding data in-country in a facility owned by a company headquartered elsewhere does not put it outside the reach of that country's law, and encryption at rest does not change that as long as the provider holds the keys. The defensible test is what is contractually enforceable: who can be compelled to hand over the data, and what migration path exists before anyone needs it.

*Figure 7. Workloads Are Concentrating, and the Most Worried Organizations Expect to Concentrate Fastest*




Note. Panel A: percentage of organizations with one environment holding half or more of production workloads; the ten-year bar is split by which environment that is (US hyperscaler = 54% of that group). Panel B: mean change in US hyperscaler share over ten years, in percentage points, by concern about dependence on US providers (Q7) and among the 456 organizations pulling work back in-house (Q24); cohorts unweighted. Panel A is weighted; Global figures weight the four regions equally (n = 1,940). Source: derived from Q18 and Q18a; Q18a by Q7 and Q24.

*"We have no concern using US cloud. Our group has bought AI companies in the US But regulations in India mandate that our data reside in our own data centers."*

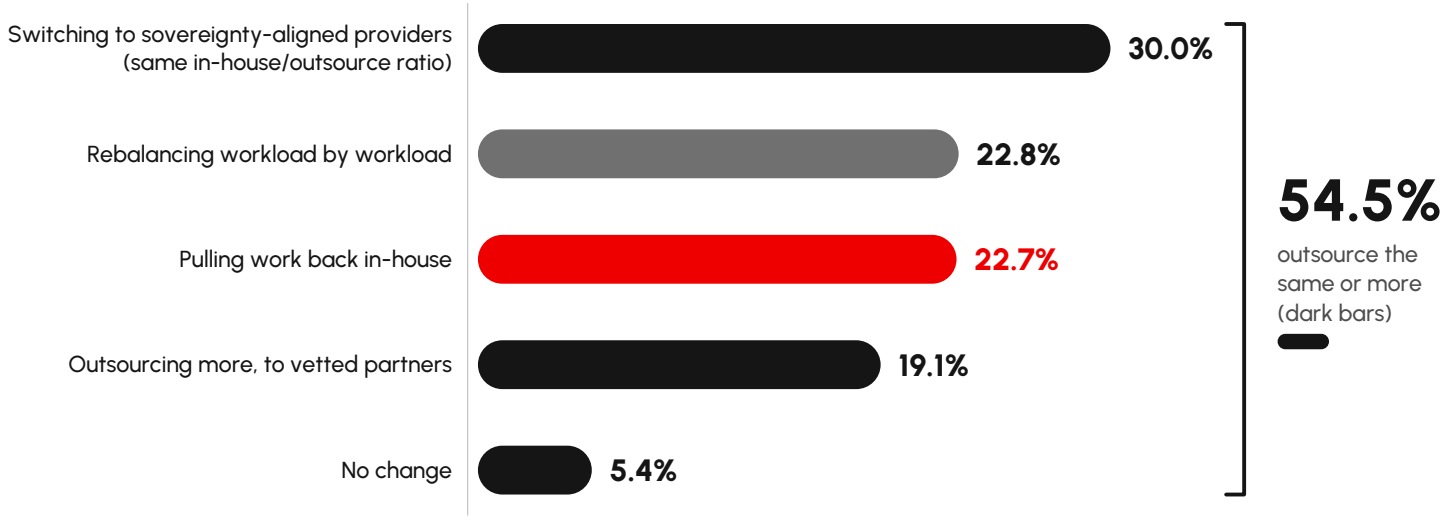
– CTO, financial services group based in India, headquartered in South Korea (Asia-Pacific)

## The money is moving between vendors, not back in-house

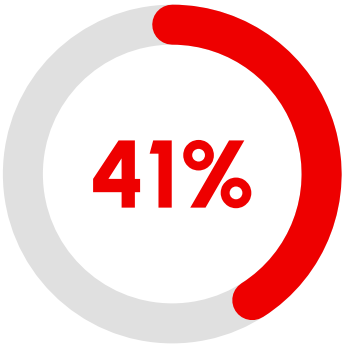
Most organizations are switching providers, rebalancing workload by workload, or outsourcing more; only a minority are pulling work back in-house (Figure 8). Most buyers also want that sovereignty delivered by the vendors that already run their cloud and infrastructure (see the callout). For a leader, the practical lever is the renewal conversation with an incumbent supplier, and the terms that matter are a documented migration path and a second source of support.

The provider models buyers expect to engage with confirm it: as noted above, a US hyperscaler with a sovereign or in-country offering and a private cloud platform provider sit jointly at the top, ahead of regional, local, and sovereign providers. No single class of workload dominates what is moving: analytics, backup and disaster recovery, AI training and inference, and core transactional systems are all being moved by roughly a third of organizations.

*Figure 8. The Money Moves Between Vendors, Not Back In-House*



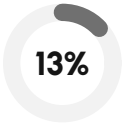

Note. Percentage choosing each description of how sovereignty is changing what they keep in-house versus outsource. Global figures weight the four regions equally (n = 1,940). Source: Q24.




say it is extremely important that sovereignty solutions come from the vendors that already run their cloud and infrastructure




say it is somewhat important




say it does not matter

Source: Q20b (importance of sovereignty solutions coming from current cloud and infrastructure vendors); asked only outside Asia-Pacific (n = 961); North America, EMEA, and Latin America weighted equally.

## Respondents expect consolidation and Futurum expects fragmentation

Figure 6 captures what respondents expect from the estate they know today: fewer environments and a larger share of production workloads on US-headquartered hyperscalers. That is a useful record of buyer intent. It is not necessarily a forecast of the market those buyers will face.

Futurum's three-to five-year outlook points in a different direction. Regulation is becoming more region-specific and more enforceable. Sovereign and in-country cloud offerings are becoming operationally distinct rather than merely positioning. Local and regional providers are gaining relevance, and telecommunications operators are moving from buyers of infrastructure to suppliers of sovereign cloud and AI services.

The result is likely to be a more fragmented supply landscape, not a simple continuation of today's concentration. Public-sector and regulated buyers will have more reasons – and more options – to separate sensitive data, AI, and critical workloads from dependencies they judge difficult to govern, support, or exit. The question for leaders is not whether US hyperscalers will remain central. They will. It is whether a strategy built around today's provider hierarchy remains appropriate as regulation, local capacity, and sovereignty requirements change around it.

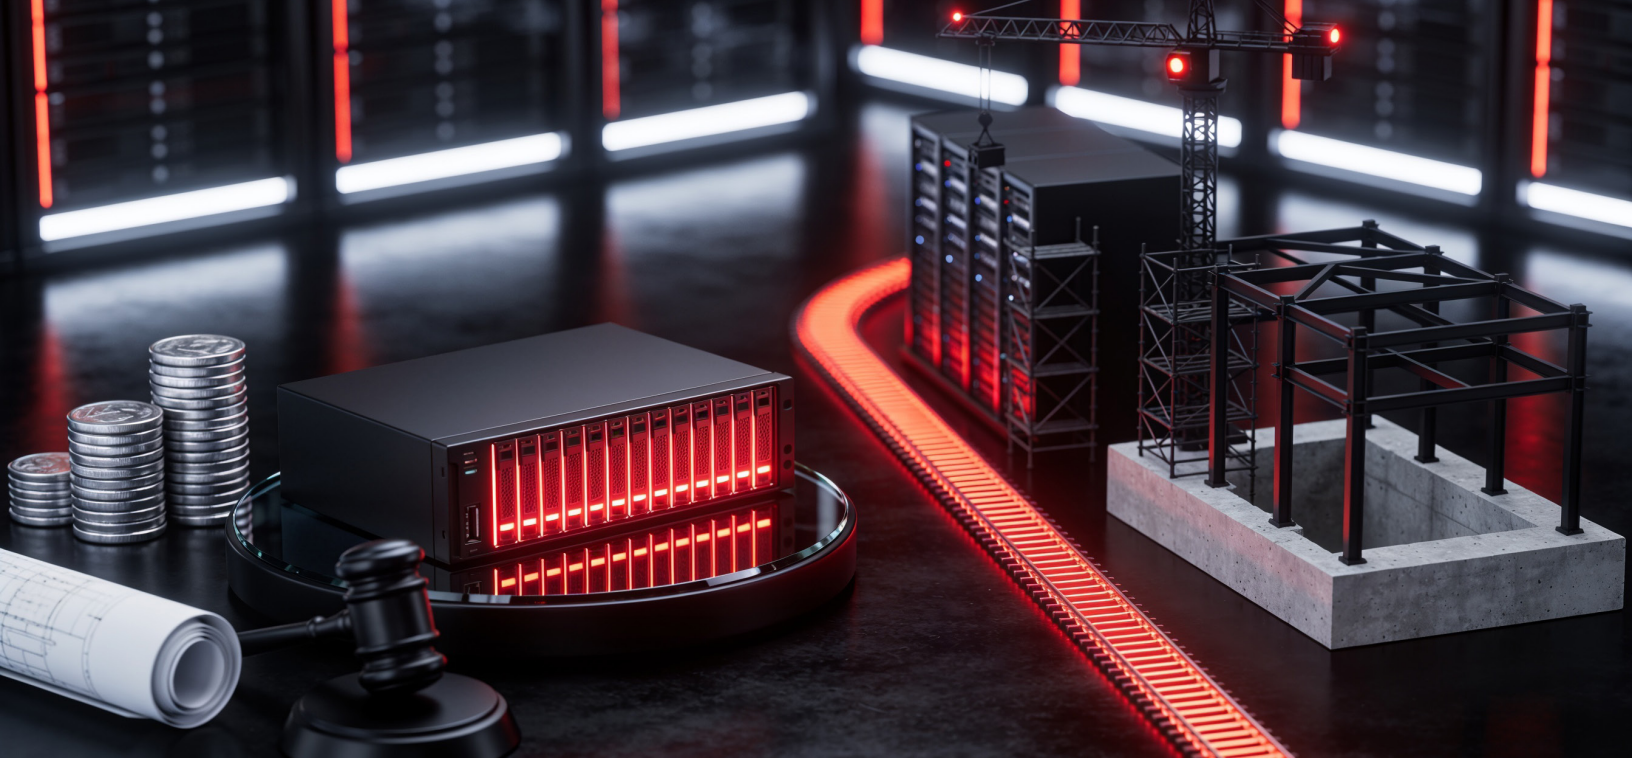
Concentration is not a hyperscaler-specific problem and treating it as one would be a mistake. Of the organizations expected to hold half or more of their production workloads in a single environment within ten years, a little over half are concentrated on a US hyperscaler; the rest are concentrated on a regional provider, a private cloud platform, or their own infrastructure. The regulatory position reflects this: operational resilience rules such as DORA address systemic overreliance on any single provider, not on foreign ones. An organization that moves from one hyperscaler to one local cloud provider has changed its jurisdiction and kept its concentration risk. What reduces the risk is the ability to move, including a documented exit, a second source of support, and a platform layer whose tooling is not tied to the provider underneath it.

The direction of travel is clearer than its speed and that uncertainty requires active management. The survey asked respondents to project their workload mix ten years ahead, yet the technology, regulatory, geopolitical, and supplier landscape has changed materially in only the past two or three years.

Leaders should test that assumption regularly. They should identify which workloads require the newest cloud capabilities today, which could run on regional, sovereign, private, or on-premises platforms as those alternatives mature, and what events would change the preferred placement. The relevant decision may also change as AI and cloud capabilities become more widely available.

The practical conclusion is not to avoid consolidation. It is to avoid making an irreversible bet on it.

**What to Do About It: Write down the assumptions behind your provider concentration and put a review date on them. A ten-year allocation is a strategic assumption, not an operational plan.**



# 3. Worry Without the Ability to Act

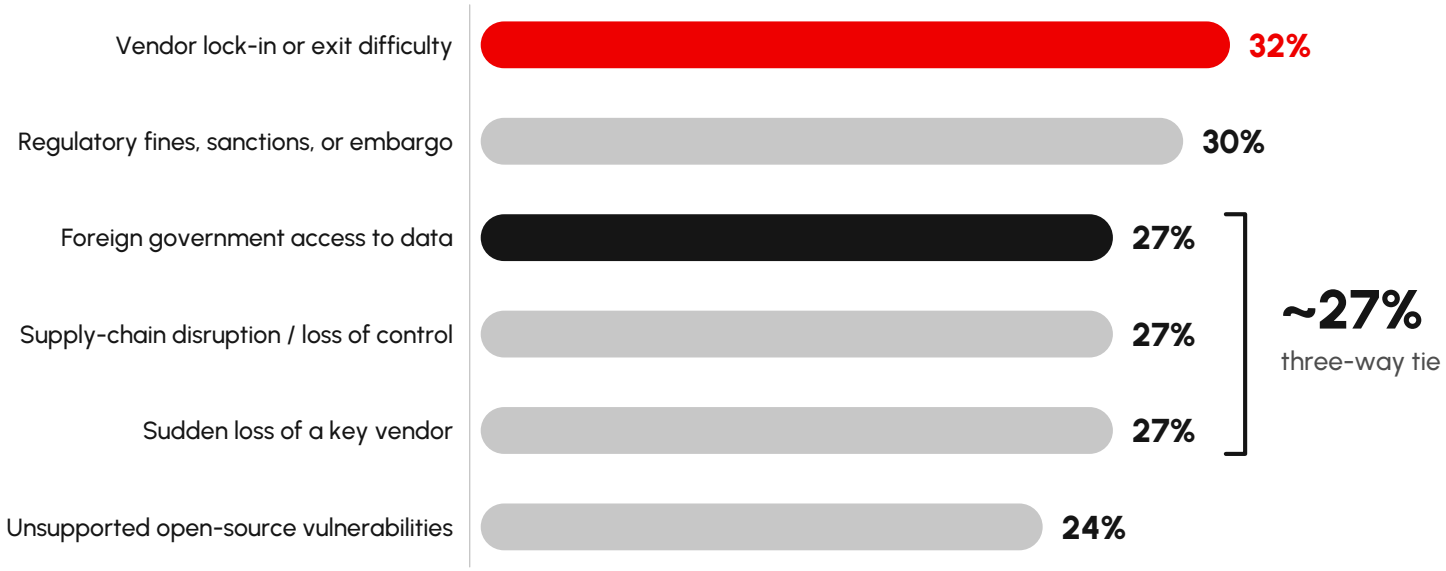
### Organizations fear being cut off, but struggle to move critical workloads quickly, and nothing has made a dent

Most organizations are worried about a provider cutting them off, and only one in four could move quickly if one did. As governments place greater strategic weight on control of critical technologies, dependence on providers governed by a single jurisdiction is becoming a more material operational risk. The US decision in June 2026 to suspend access to a leading US AI provider's most capable models turned that scenario into a precedent. Yet, even organizations with mature sovereignty programs say that they cannot act quickly if they were cut off.

### The risks buyers fear most are the ones that trap them

Nearly nine in ten organizations are at least moderately concerned about a supply-chain kill switch (a vendor or government shutting off hardware or software they depend on), and eight in ten about depending on US-headquartered providers for critical workloads. Asked which risks concern them most, buyers put vendor lock-in and exit difficulty first (Figure 9); foreign government access to data, the risk sovereignty usually brings to mind, ties for third. What buyers want is the ability to leave, which is a procurement and architecture problem more than a data-residency one. That points to a specific answer. If the risk buyers name first is being unable to leave a supplier, then the relevant property of a technology is whether anyone other than its original vendor can keep it running. Open source is the only category where that is structurally possible, and buyers have already chosen it: it has a role in 87% of sovereignty strategies and is foundational to 37%. What they have not necessarily secured is the support underneath it: 56% are extremely concerned about running unsupported packages, the highest level of concern recorded on any item in this study, and 52% call commercial support for the open source they run critical.

*Figure 9. The Risks Buyers Fear Most Are the Ones That Trap Them*




Note. Percentage naming each risk among their top two sovereignty concerns. Global figures weight the four regions equally (n = 1,940). "Reduced access to innovation" (14%, asked only outside Asia-Pacific) is not shown. Source: Q5.

*"The ability for [a country or vendor] to actually…to invoke a kill switch and lock us out of our data is very real."*

– CDIO, public-sector social security agency (Europe)

## Only 2.9% of organizations hold all four capabilities needed to act

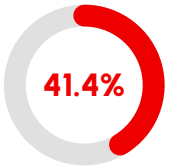
Four capabilities enable an organization to respond if a supplier cuts it off: a formal, documented, and tested contingency plan for vendor disruption; a plan for open-source software that can no longer be supported; a sovereignty risk framework applied to every supplier; and the ability to move a critical workload off a single cloud provider within 30 days (Figure 10).

Each capability is insufficient on its own. A contingency plan does not ensure workload portability. Portable workloads do not resolve an unsupported software dependency. A supplier-risk framework does not help if the organization cannot execute an exit. Taken together, the four provide a practical minimum test of whether a sovereignty program can operate under disruption. Only 2.9% of organizations meet that test. At the other end, 24.6% have none of the four capabilities. The gap is not simply the one of policy maturity; it is a gap in operational readiness.

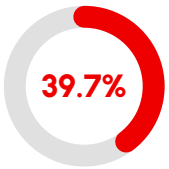
For leaders, the self-assessment is straightforward. Organizations that hold three of the four capabilities have identified part of the problem but still have a critical exposure. Those that hold one or none should begin with a full inventory of critical workloads, cloud providers, software dependencies, support arrangements, administrative access, and contractual exit rights. The objective is not to document every risk. It is to establish the minimum conditions for keeping critical services running when a dependency fails.

*Figure 10. Only 1 Organization in 35 Holds All Four Capabilities That Sovereignty Requires*

**A Share holding each capability**



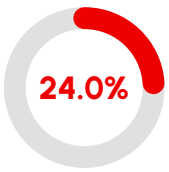

Formal, documented, tested contingency plan for a critical vendor being cut off




Formal plan for open source that can no longer be supported

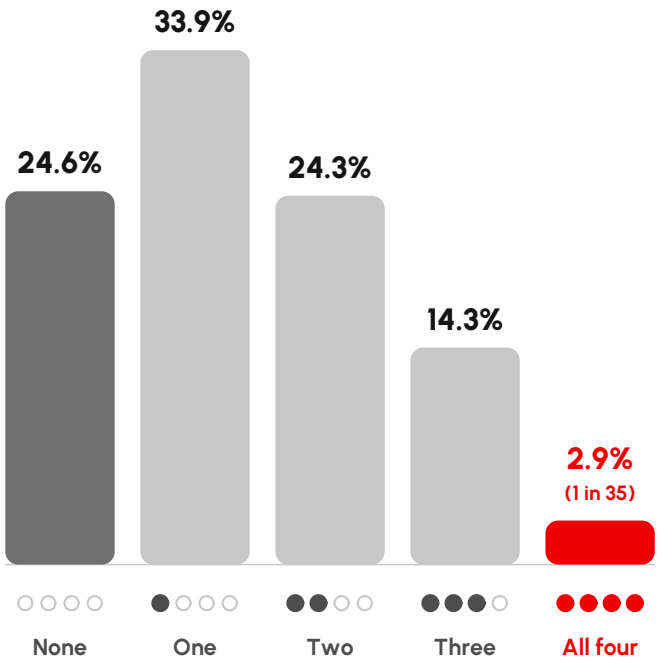



Sovereignty risk framework applied to every supplier




Can move a critical workload off a single cloud provider within 30 days

**B How many of the four do organizations hold?**



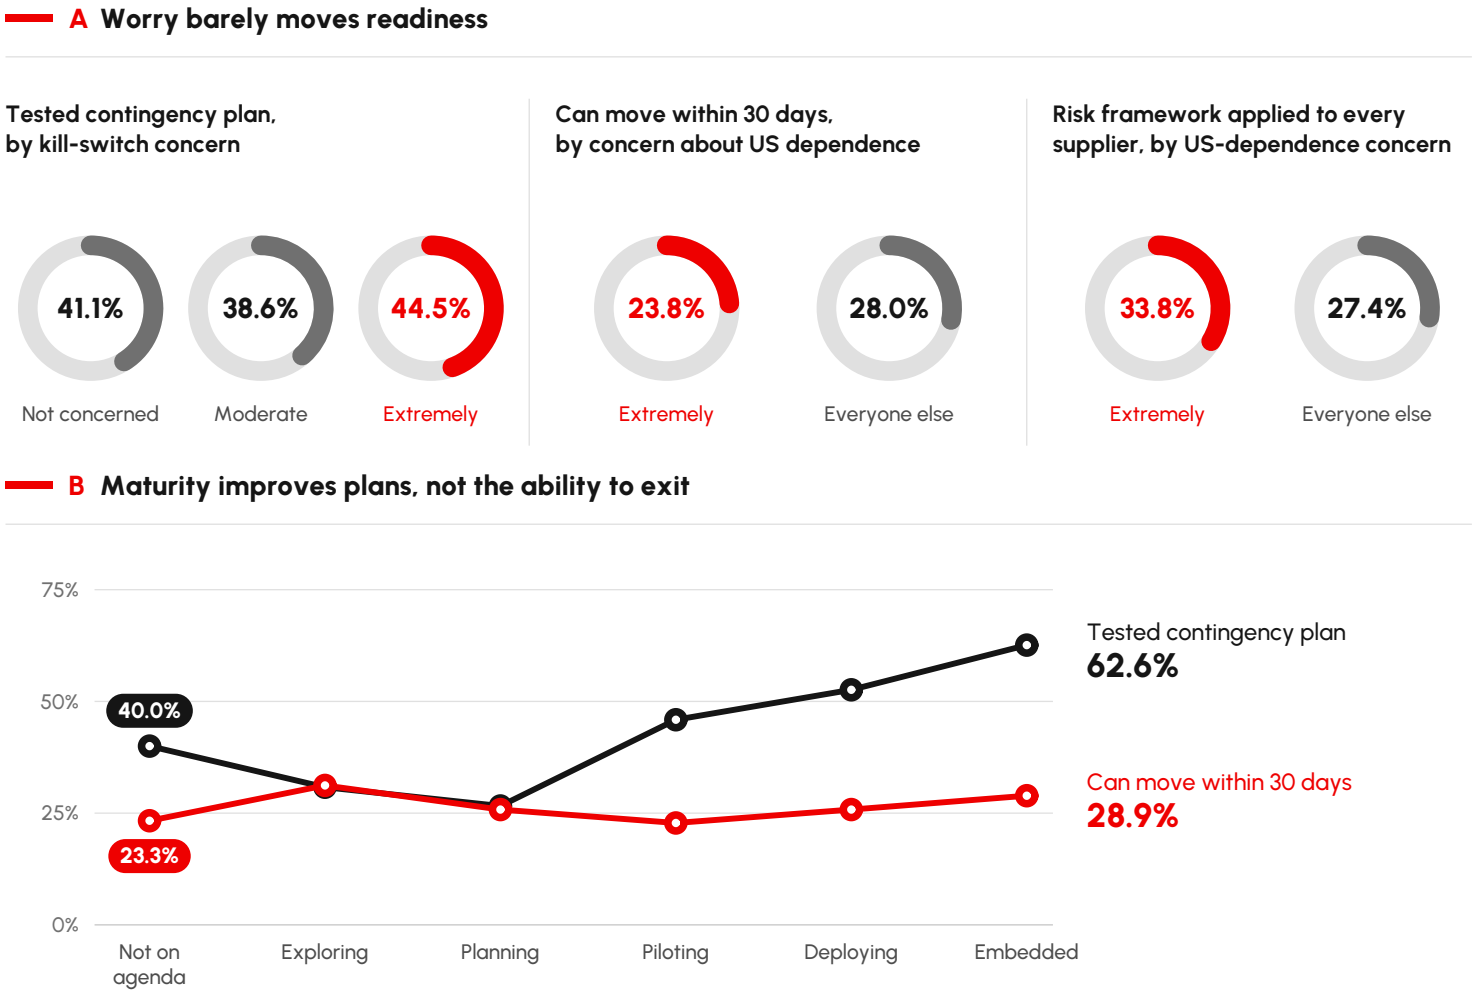

Note. Panel A: percentage holding each capability. Panel B: distribution by number of the four capabilities held. Capabilities: formal, documented, and tested contingency plan (Q8a); formal plan for unsupported open source (Q13); sovereignty framework applied to every supplier (Q6); able to move a critical workload within 30 days (Q19). Global figures weight the four regions equally (n = 1,940).



## Neither worry nor a mature program produces the ability to leave

Worry does not predict capability. A tested contingency plan is barely more common among organizations extremely concerned about a kill switch than among organizations not concerned, and even they are no more able to move a workload within 30 days. Program maturity shows the same pattern: organizations with sovereignty embedded in operations hold a tested plan far more often than organizations that have not put sovereignty on the agenda, yet the two groups' ability to move a workload within 30 days barely differs (Figure 11). Worry and maturity both buy governance, not readiness: a program's plans improve well before its ability to exit does.

*Figure 11. Worry and Maturity Buy Plans, Not Exits*




Note. Panel A: percentage with a tested contingency plan by concern about a supply-chain kill switch (Q8a by Q8); percentage able to move a critical workload within 30 days (Q19) and percentage applying a risk framework to every supplier (Q6), extremely concerned about US dependence vs. everyone else (by Q7). Panel B: tested contingency plan and ability to move within 30 days by sovereignty stage (Q1). Cohorts and stages unweighted.

*"We do not have a tested Plan B if a key vendor's hosting fails. Migrating from [a US hyperscaler] to [another US hyperscaler] remains a future ambition."*

– CDIO, Healthcare (Europe)

## Many factors drive provider-switching speed: size, industry, and region

Moving a critical workload off a single cloud provider takes most organizations months rather than weeks – nearly 21 weeks on average – and the ability to do it within 30 days falls as organizations get larger (Figure 12), long enough that a disruption reaches far beyond the IT department. The largest enterprises are also the most likely to call sovereignty a critical priority, and the reasons they move slowest compound: more systems and legacy, slower change, multiple jurisdictions, and a larger, more visible target. As a result, the organizations with the most at stake are also the ones least able to respond quickly.

**21 weeks**

average time to move a critical workload off a single cloud provider

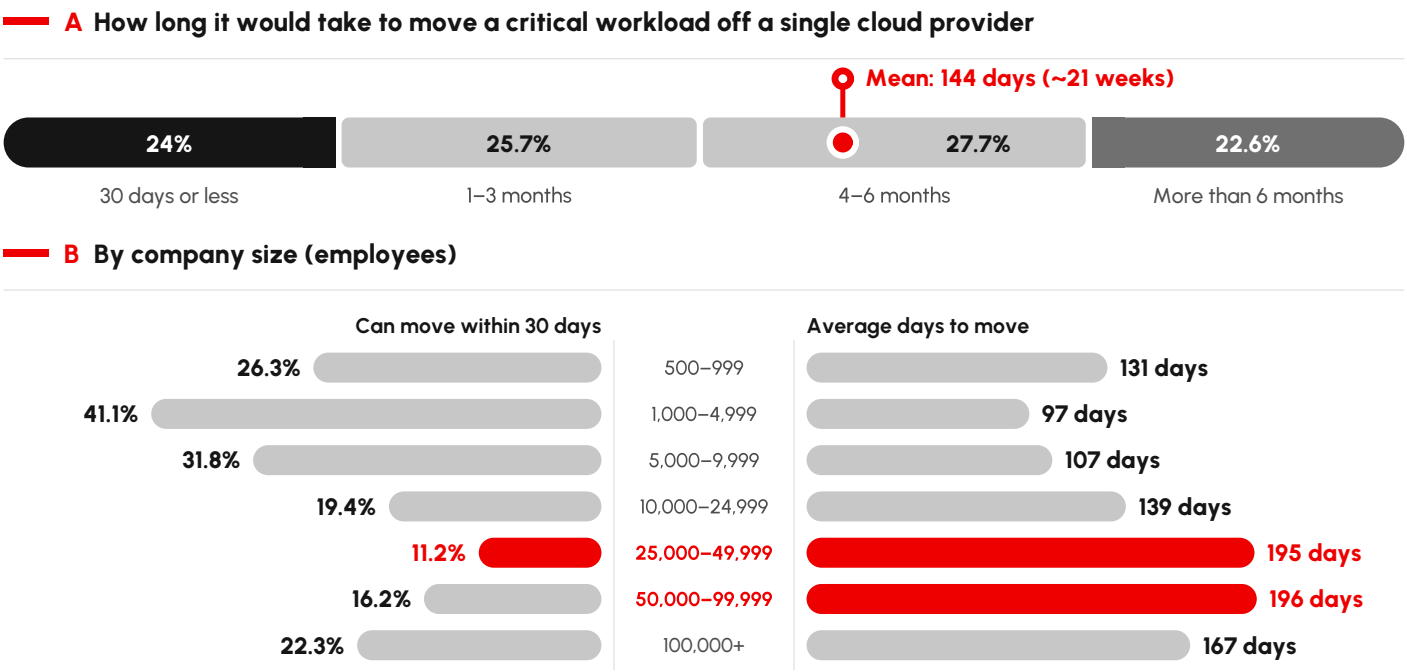



Source: Q19, mean of category midpoints (144 days); regions weighted equally.

The organization's maturity in its sovereignty rollout has surprisingly little impact, with embedded-sovereignty organizations actually performing the slowest of any cohort in the sample. This doesn't indicate the sovereignty work slowed them; it likely reflects that these organizations had particularly strong need for such programs and still have many blockers, but the result is surprising nonetheless.

Other factors also make a difference. Asia-Pacific organizations are able to migrate much faster: a median of 10 weeks, versus 17 to 19 weeks for the other three regions. Public sector organizations are uniquely slow: 17 weeks versus 10–13 weeks for other industries.

*Figure 12. Moving a Critical Workload Takes About Five Months, and the Largest Organizations Take the Longest*



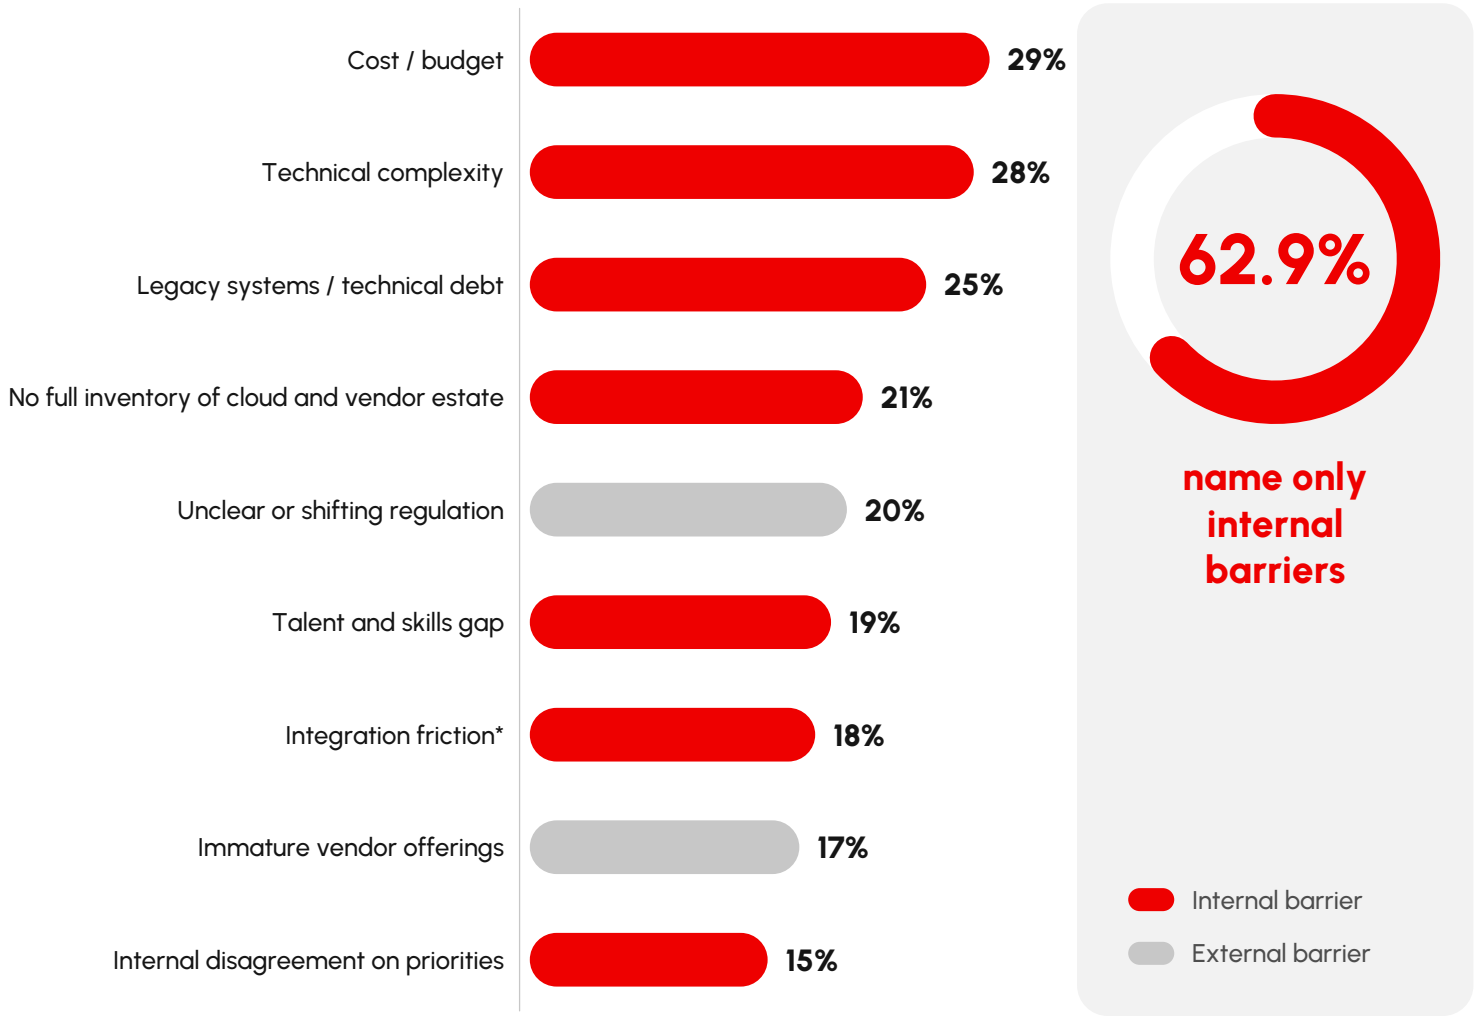

Note. Panel A: time to move a critical workload off a single cloud provider; mean of 144 days (about 21 weeks) based on category midpoints; n = 1,938 excluding two who said they cannot migrate. Global figures weight the four regions equally (n = 1,940). Panel B: percentage able to move within 30 days and mean days to move, by company size (S3); bands unweighted. Source: Q19 by S3.

## Facing significant barriers, organizations procrastinate on sovereignty, waiting until external shocks force workload moves

Every leading event that would trigger moving workloads off a cloud provider is an external shock: new or stricter regulation, a security or compliance incident, a geopolitical event, and a vendor kill switch. Organizations are waiting for something to go wrong instead of building the capability in advance.

Barriers to sovereignty programs are primarily internal (in order): cost, technical complexity, legacy systems, a missing cloud and vendor inventory, and a skills gap (Figure 13). For organizations without sovereignty on their agenda, cost and budget is the runaway driver – noted as a top barrier 11 points higher than average. Small-to-midsize businesses and nonprofits identify both cost and technical complexity as major barriers at notably higher rates than enterprises; while those barriers still top the list for enterprises, they do report more significant barriers with legacy systems and technical debt, lack of inventory, and talent and skill gaps at higher rates.

*Figure 13. The Biggest Barriers to Sovereignty Are Inside Organizations*




Note. Percentage naming each barrier among their top two. Global figures weight the four regions equally (n = 1,940). *Integration friction was asked only outside Asia-Pacific (n = 961). 62.9% named only internal barriers. Source: Q23.

Although buyers may fear a government cutting them off, the disruptions practitioners actually described were commercial: a vendor acquired and then shut down, an open-source component whose new owner began charging for it, cost changes buried in a contract. Price pressure ranks alongside geopolitical events among the triggers.

The barrier is rarely conviction. It is that the work competes with everything else on the roadmap and wins only when something forces it – as one retail CIO put it:

*"It's hard to make time for a project that brings no direct value and is purely about controlling risk."*

– CIO, retail organization (EMEA)

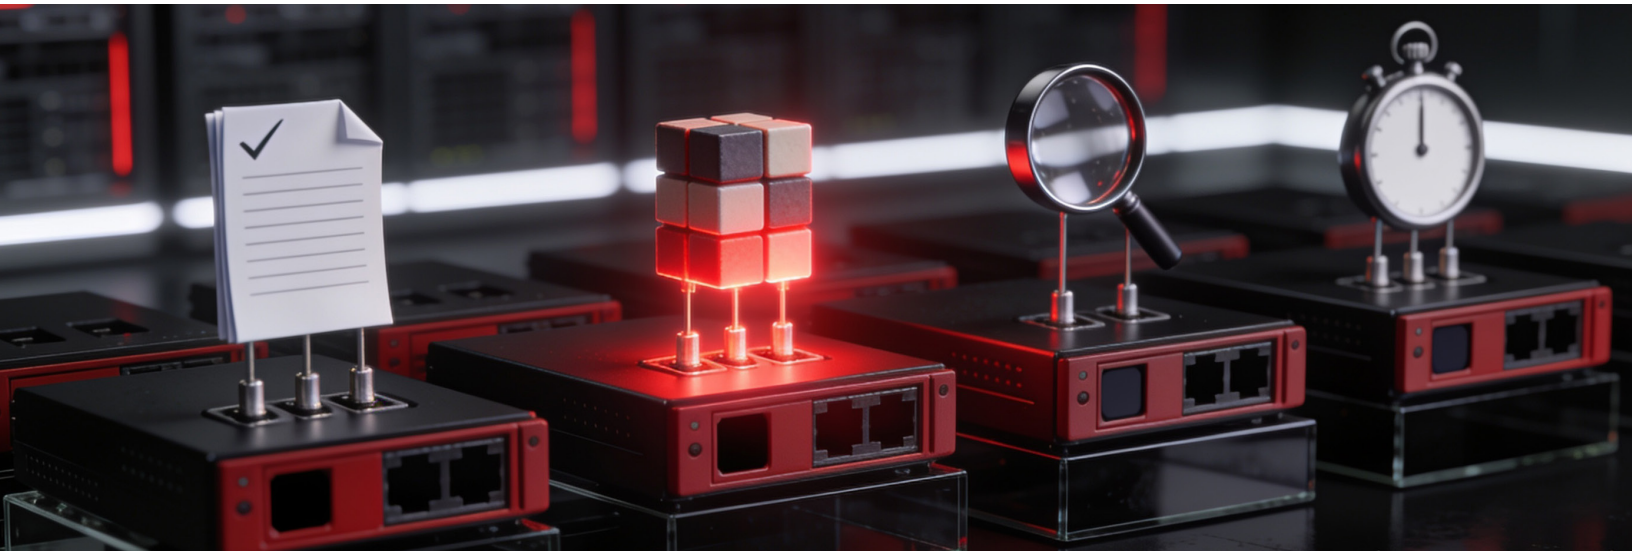
**What to Do About It: Score yourself against the four capabilities. Holding three still leaves a critical exposure; holding one or none means starting with an inventory, not a policy.**

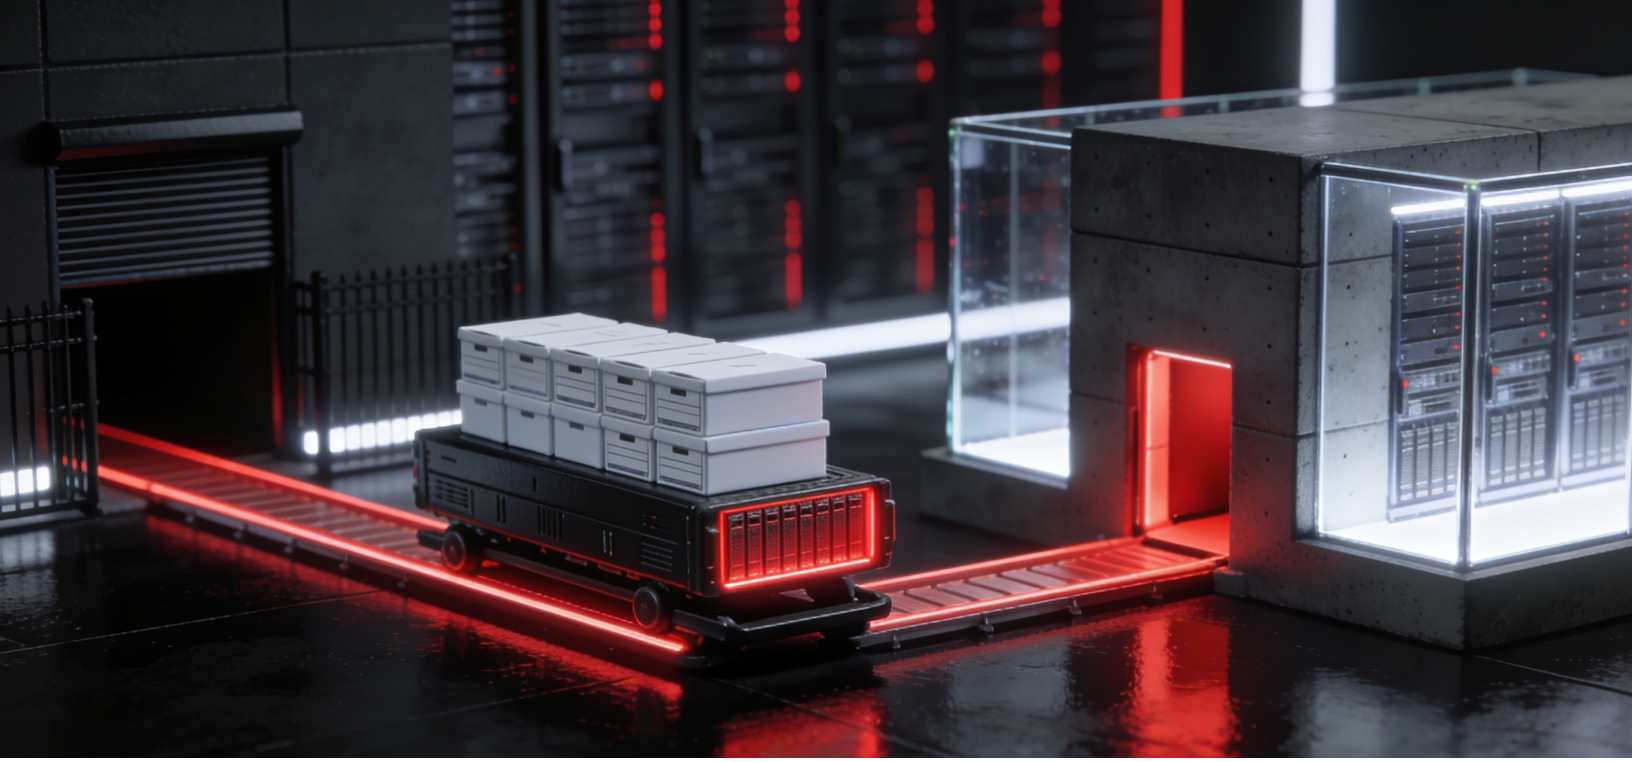
*"A vendor we used was acquired, then shut down with 2 years' notice. The real risk was migrating data we're required by regulation to retain for 9 years."*

– CTO, financial services group headquartered in South Korea (Asia-Pacific)

*"My main problem is I don't have enough people on my team, and there's no appetite to let me hire. I'm dependent on third parties, and that worries me."*

– CIO/CISO, electricity distribution company (Asia-Pacific)





# 4. What Buyers Will Actually Pay For

### The top-rated requirement is somebody else able to keep the software running

If workloads are staying with the largest cloud providers, readiness has to come from somewhere other than where data sits. Asked what they need for sovereignty, buyers rate an alternative source of support and maintenance above everything else measured, well ahead of the next requirement (Figure 14). The requirement is strongest in the Americas and weakest in the Middle East, but even there 98% call it at least moderately important.

Buyers rated support as a set of separate requirements (Figure 14): a second source of support and maintenance if the primary vendor fails; access to the vendor's engineers and community; commercial support for the open-source code they run, with service-level guarantees, patching, and legal indemnification; a formal plan for open-source software that loses its support; and a say in who staffs the support. A vendor's support offer has to answer all five.

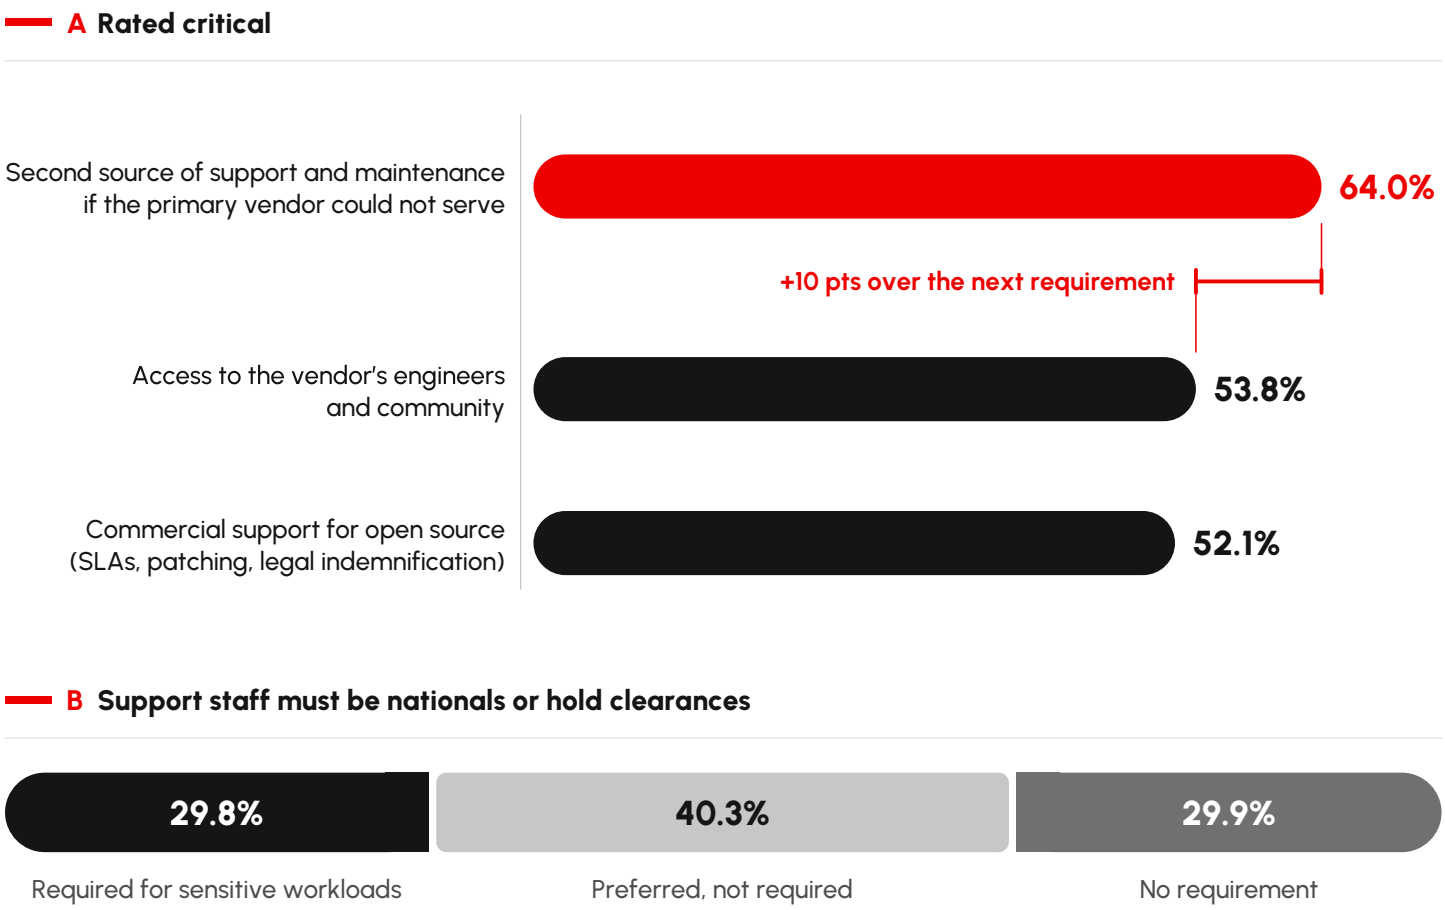
*"We require fail-safe vendors we can flip to if the primary fails. We identify a competitive vendor, pilot both, and choose the best, always keeping a second option in reserve."*

– CISO & Cloud Architect, regional health insurer (North America)

*"Enterprise support is non-negotiable. Even a millisecond of downtime is unaffordable, because price swings in the stock market mean our customers can't absorb losses caused by us."*

– CTO, financial services (Asia-Pacific)

*Figure 14. The Top Requirement Is a Second Source of Support, and Support Has Several Parts*



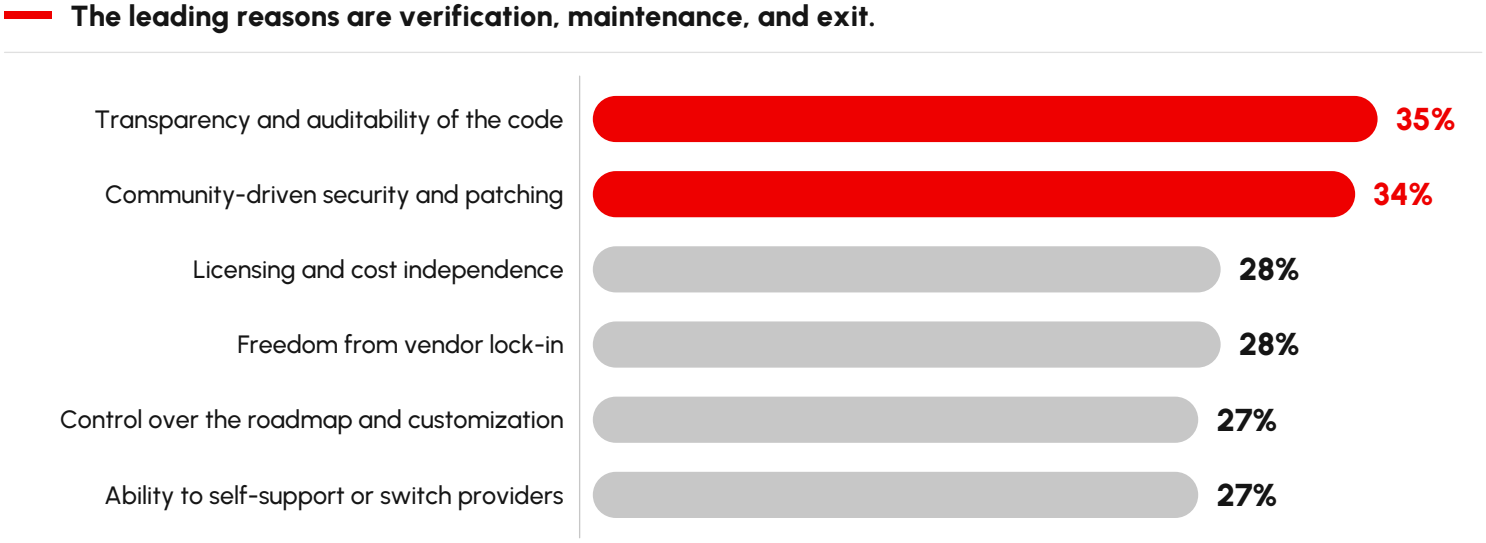

Note. Panel A: percentage rating each requirement critical. Panel B: whether support personnel must be citizens or nationals or hold clearances. Global figures weight the four regions equally (n = 1,940). Source: Q26a (alternative source of support), Q26 (vendor talent and community), Q12 (commercial support for open source), Q26b (support personnel nationality or clearance).



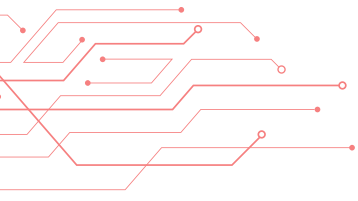
## Open source underpins 87% of sovereignty strategy, but leaders worry most about running it unsupported

Open source has a role in 87% of every sovereignty strategy, and is foundational for 37% of them. It is chosen for the rights the code carries: to inspect it, patch it, and move it without asking permission (Figure 15). Confidence in open source carries real anxiety with it: concern about running unsupported open-source packages is the highest level of concern on any item measured (56% of respondents), and it is even more concerning among the organizations that call open source foundational than anywhere else. Sovereignty is creating demand for commercially supported open source (important or critical for over 90% of respondents): the code grants the right to inspect, patch, and move it, and the support contract makes that right usable on a system the business cannot afford to lose. In Futurum's interviews, no C-suite officer was ambivalent about this: unsupported open source was a non-starter for critical systems. Where open source was used, they wanted it commercially supported and able to run consistently across the cloud, private, and local environments their estates require.

*Figure 15. Why Open Source Suits Sovereignty*




Note. Percentage selecting each reason, among the 1,711 respondents who give open source a role in their sovereignty strategy; regions weighted equally. Source: Q10a.



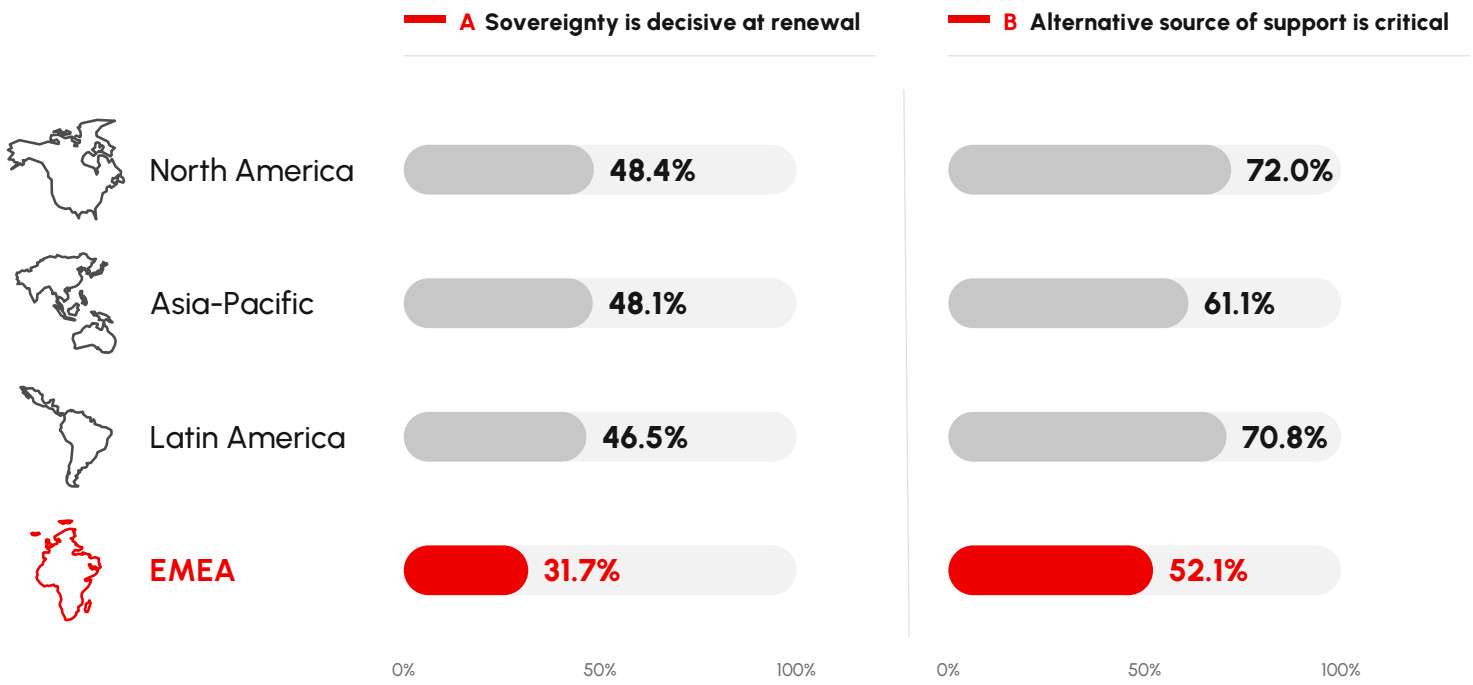
*"Open source is probably more secure than closed source, but support is the issue. I wouldn't switch to an open-source ERP or point-of-sale system without support for critical failures."*

– CIO, retail organization (EMEA)

## Sovereignty has entered the renewal cycle, ahead of the ability to act on it

Sovereignty has become a deciding factor in infrastructure renewals, considered in 86% of reviews and renewals, and critical in 44% (Figure 16). Buyers who treat sovereignty as decisive at renewal already know what they want: a second source of support and commercially supported open source, at far higher rates than everyone else. They are no more able to move a critical workload within 30 days than anyone else. A renewal is the one point where a buyer holds real leverage. Most are using it to add sovereignty language to contracts, ahead of a documented exit path or an identified second source of support.

*Figure 16. EMEA Lags on Making Sovereignty Count at Renewal*



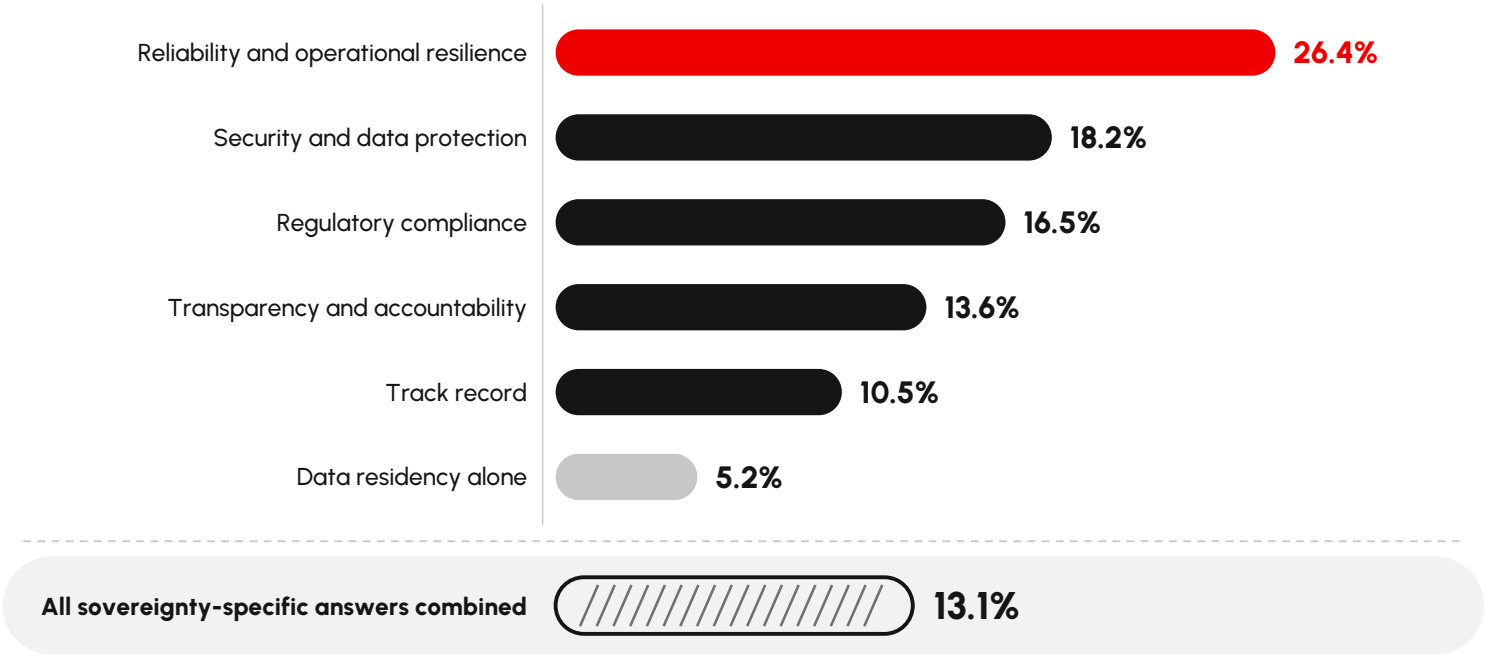

Note. Panel A: percentage saying sovereignty is a decisive factor when reviewing or renewing infrastructure subscriptions (Q12a). Panel B: percentage rating an alternative source of support and maintenance critical (Q26a). Regional figures are unweighted: North America n = 289, EMEA n = 470, Asia-Pacific n = 979, Latin America n = 202. Open-source adoption is the same in North America and EMEA (87%; Q10).



## Buyers judge a sovereignty-ready vendor on reliability, not on where the data sits

Asked in their own words what makes a vendor ready for sovereignty work, buyers name reliability and operational resilience first, well ahead of every other reason offered, including data location (Figure 17). Security and data protection place second, regulatory compliance third; every sovereignty-specific reason combined still does not outrank reliability alone. Practitioners interviewed at length gave the same answer unprompted: each one, asked directly, named a capability, a reputation, or a relationship. The reliability buyers mean is reliability under disruption: a second source of support ranks first among requirements, and being unable to leave ranks first among feared risks, because a single provider, however dependable, cannot keep a customer running through its own failure. For most buyers, sovereignty opened the conversation, and reliability and operational outcomes led them to invest.

*Figure 17. Buyers Name Reliability First and Data Location Last*




**What to Do About It: Identify where support would come from if your primary vendor could not provide it, and raise it at the next renewal rather than the next incident.**





# 5. AI Turns Sovereignty Into a Purchase Decision

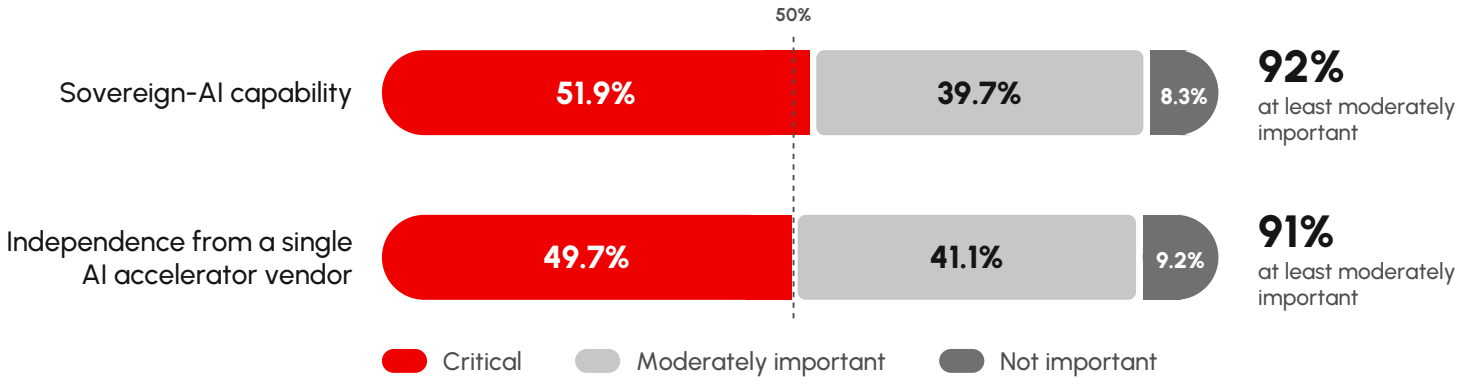
### AI turns the readiness question into an urgent purchase

AI is where sovereignty becomes an immediate operating decision. Organizations must decide which models they can use, where sensitive data is processed, where inference and training run, who controls the underlying infrastructure, and what happens if a model, cloud, or hardware supplier becomes unavailable.

Just over half of organizations call sovereign AI capability critical to their AI strategy (Figure 18). Almost as many call independence from a single AI accelerator vendor critical. That is an aspiration more than an immediate available choice. NVIDIA remains the dominant supplier of AI data center accelerators and its CUDA software ecosystem makes moving to alternatives more difficult than simply replacing one chip with another. The finding does not mean buyers can substitute suppliers tomorrow. It means they recognize the exposure and want the ability to reduce it over time.

For AI, the operational question is straightforward: can the organization continue to run a critical AI service if a model provider, cloud provider, accelerator supplier, or critical software dependency changes terms or becomes unavailable? AI commitments are made faster than the platform, data governance, and infrastructure choices needed to support them. Without those foundations, organizations recreate the cloud sovereignty readiness gap in AI, only sooner.

*Figure 18. Half Call Sovereign AI Critical, and Half Say the Same of Chip Independence*




Note. Importance of sovereign-AI capability (Q14) and of independence from a single AI accelerator vendor (Q17). Global figures weight the four regions equally (n = 1,940).

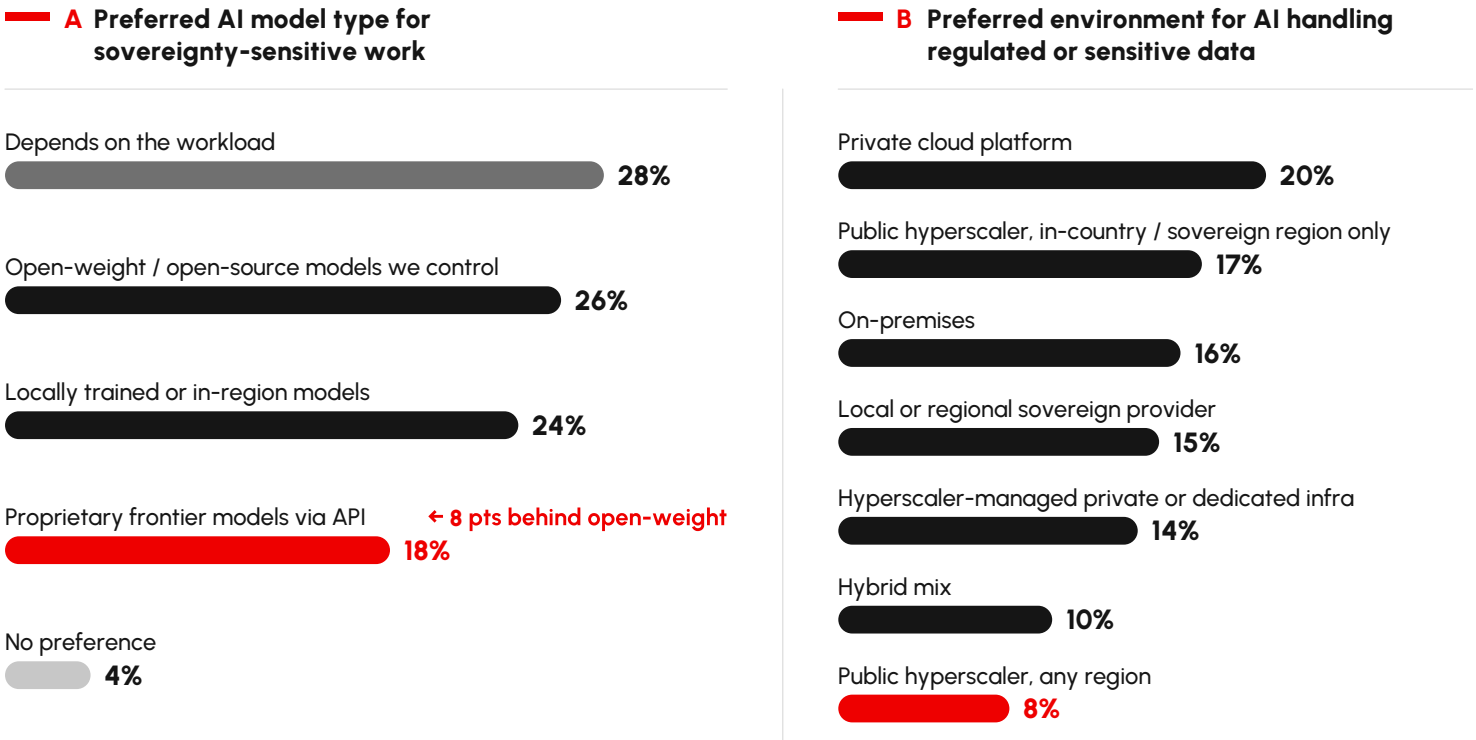
## Buyers choose models by workload and place sensitive AI where control matters

Buyers choose AI models by workload, not by type. Where they commit, open-weight models, meaning models with published parameters an organization can download and run itself, outrank proprietary frontier models reached through a vendor's interface. Preferred environments for sensitive AI follow the same logic: a private cloud platform leads, then a public hyperscaler limited to a sovereign region, a jurisdiction that keeps the workload in place, with an unrestricted public hyperscaler last (Figure 19). Most organizations choose the strongest frontier model where capability decides the outcome, and run open-weight models where control decides it, so a technology decision-maker needs a commercially supported platform that moves workloads across public cloud, sovereign region, private cloud, and on-premises.

As governments and government-backed national AI companies fund models trained locally and in-region, open-source AI models will appear from many more countries, widening the set an organization can run under its own control. And as other classes of models catch up to present-day frontier capabilities, even despite the lag, leaders may find themselves much more able to embrace open or local models even for workloads requiring advanced, dependable capabilities, opening the door to more prioritization of sovereign, control, or open-source options.

AI workload movement is part of the near-term sovereignty agenda: 30% of organizations say they are actively repatriating or migrating AI workloads to private or local trusted infrastructure within the next 24 months, rising to 41% in Asia-Pacific against 26% elsewhere. Agentic AI, AI that acts and decides on its own instead of only answering a prompt, raises the stakes. As AI systems move from answering prompts to taking actions in regulated processes, organizations need clearer controls over where inference runs, what data and tools the system can access, how decisions are observed, and how the system can be stopped or moved if a dependency changes. A leader planning agentic AI for a regulated process should decide the required level of control over inference, data access, tools, and observability before choosing a model.

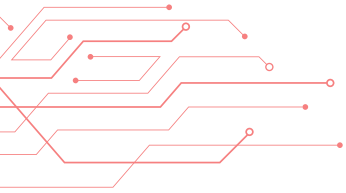
*Figure 19. Buyers Choose Models by Workload and Place Sensitive AI Where Control Matters*




Note. Panel A: preferred AI model type for sovereignty-sensitive work (Q16). Panel B: preferred environment for AI workloads handling regulated or sensitive data (Q15); all seven response options shown. Global figures weight the four regions equally (n = 1,940).

## Practitioners still pick the best model, for now

What companies say they prefer and what they actually use are often different. While buyers say they favor models they control, convenience and performance usually win out in practice. For control to matter, leaders must make it a firm requirement for sensitive workloads rather than just a general preference. The actual decisions vary substantially by the organization's needs; region may matter little, but industry may make a significant difference. In the case of these two European executives, the retailer optimizes for capability while the public agency accountable for citizen data focuses on control:



*"We don't factor sovereignty into AI. We use [a leading US AI assistant] and [another leading US AI assistant] because they're better than [a French AI vendor]. We want the best model."*

– CIO, retail organization (EMEA)

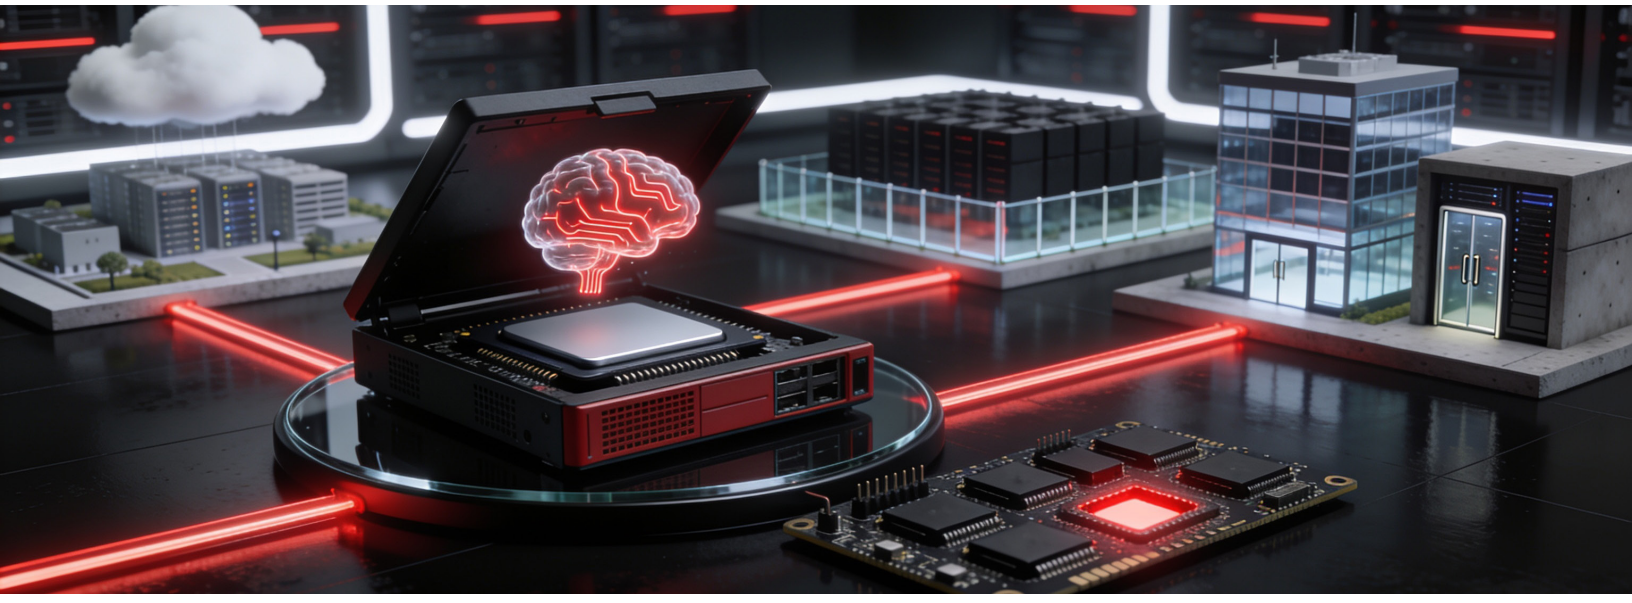


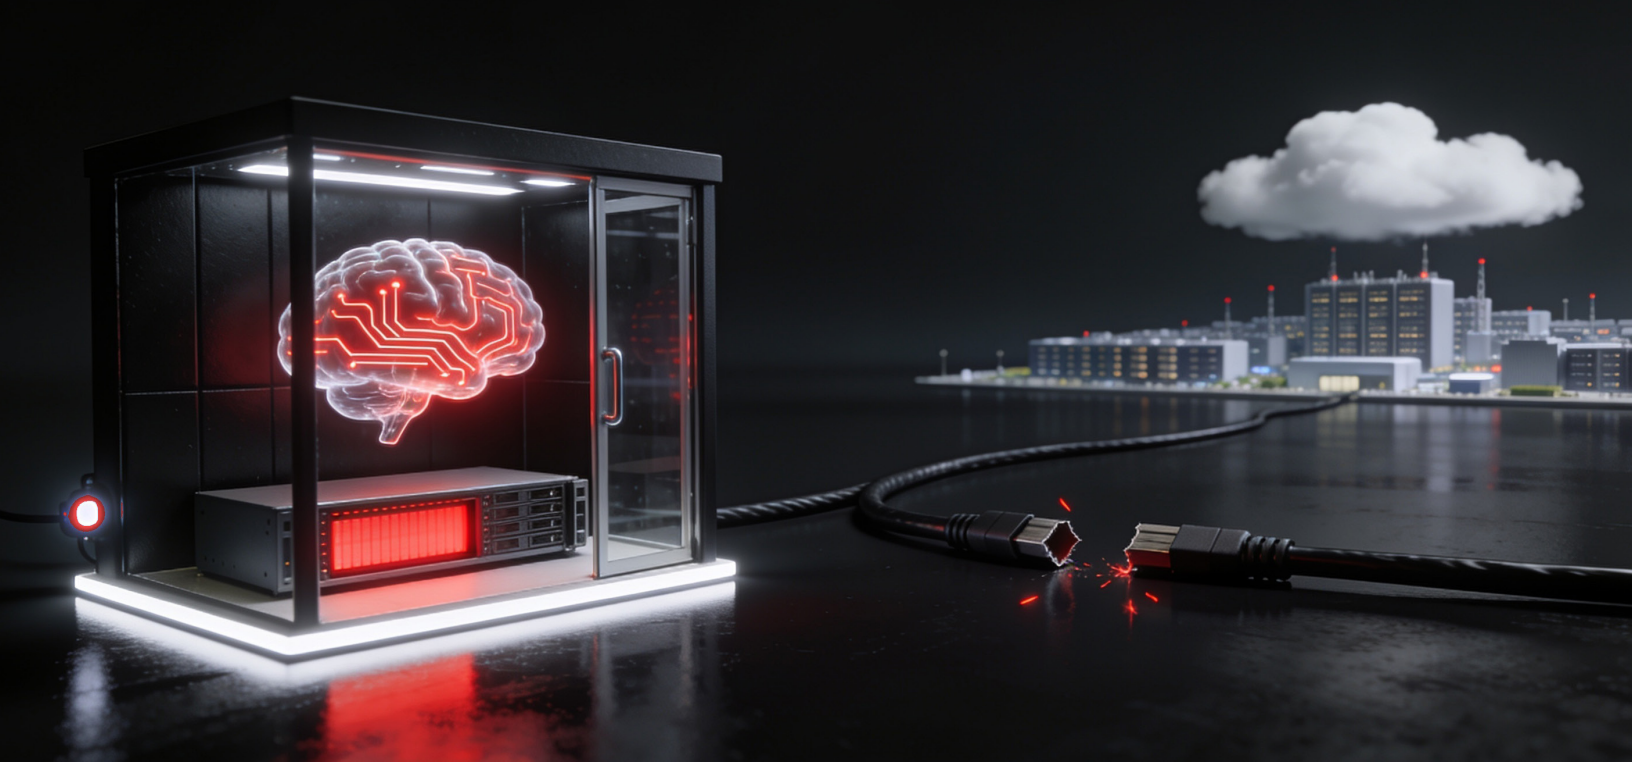
*"We'll run our own LLMs in our own environment, get them from a vendor but own the entire IPR, control it, and make sure it doesn't talk back to a mothership elsewhere."*

– CDIO, public-sector social security agency (EMEA)

A technology decision-maker should know whether their own organization has already made that choice. Organizations that hold to a sovereign-AI policy under pressure are the ones that chose a controlled environment for their sensitive workloads before the choice became urgent.

**What to Do About It: Decide the required level of control over inference, data access and tooling before you choose a model, not after.**





# 6. What To Do Next

### What a technology leader should know, and do next

Most large organizations have a sovereignty budget and a governance program. Most are not yet ready to keep critical services running if a supplier, cloud service, software component, or support path becomes unavailable. The recommendations below close that readiness gap by turning intent into tested operational capability.

- **Test the exit from an environment, against a clock**. Choose one critical production workload, move it or a full copy of it to a second environment, and time the move, at least once a year and after any major change. A plan that has never been run describes what you hope will happen: four in ten organizations hold a formal, tested contingency plan, and only one in four could move a critical workload within 30 days.

- **Write exit plans into the contract at renewal**. At the next infrastructure renewal, add a documented migration path with timelines and data-export commitments, and identify in advance where support would come from if the primary vendor could no longer provide it, which can mean a third party, or your own staff certified to support the components you depend on. Sovereignty is already decisive at renewal for more than four in ten buyers, and a few have either term in writing.

- **Prefer open, commercially supported software where being trapped is the risk**. Choose software you can inspect, patch, and move, backed by a commercial support contract with service-level guarantees, patching, and legal indemnification, and run it the same way across a public cloud, a sovereign region, a private cloud, and on-premises. Unsupported open source is the wrong choice for anything critical.

- **Build an inventory you can act on**. List every cloud, vendor, and hosted service the business depends on, and for each record where the data sits, who holds the encryption keys, which open-source components run underneath, and what would replace it. Most inventories stop at the application layer, where exits actually fail.

- **Assess every supplier, and make the assessment include a way out**. Extend your sovereignty framework beyond critical vendors to every supplier: every one should answer whose law reaches the data held for you, who staffs your support and where, and how you leave. Disruptions practitioners described came from acquired vendors and small suppliers as often as from the largest.

- **Settle who owns sovereignty, and report it to the board as a risk**. Give one owner, whether the CIO, the CTO, the CISO, or the chief risk officer, authority to accept or reject a dependency on the organization's behalf, and report accepted dependencies to the board like any other operating risk, with the exit tested and the second source named. One organization in five reports active internal disagreement about direction or ownership.

- **Revisit the decision to consolidate every year**. If the plan is to place most workloads with a single provider, write down the assumptions behind it and review them on a fixed schedule, keeping the exit tested and the second source live. A data center's location does not grant legal immunity, and encryption does not settle the question while the provider holds the keys.

- **If you run a large enterprise, go first**. Size, legacy systems, and a footprint across several jurisdictions make the largest organizations the most exposed and the slowest to move, so start with one workload and the exit test this year. What that test shows decides the rest of the list.

*"If we cannot administer their cases or make their payments, it disrupts their ability to put food on the table. In the worst case, it could lead to civil unrest."*

– CDIO, public-sector social security agency (EMEA)





# Methodology

Futurum Research surveyed 1,940 technology and risk decision-makers in 28 countries and interviewed 12 senior executives in July and August 2026 to evaluate the state of digital sovereignty today on a global basis. Futurum wrote and owned the questionnaires, oversaw recruitment of respondents, and independently conducted the interviews. Red Hat sponsored this research and was provided an opportunity to review the questionnaires and analysis, but Futurum retained complete final control over all research design, analysis, and content writing.

**Survey firmographics: **Respondents were surveyed across regions (Asia-Pacific: 979, Europe: 369, US and Canada: 289, Middle East: 101, Latin America: 202) and six industries (financial services 353, manufacturing and energy 334, retail and consumer 334, healthcare and life sciences 308, public sector and defense 306, telecom and digital infrastructure 305). Among organizations, 93% have 1,000 or more employees, and the average annual cloud, infrastructure, and AI spend is $173M across the sample. Most operate in multiple countries: 93% of respondents' firms generate 10% or more of their revenue from countries outside their home country, with an average of 46% of total revenue generated outside the home country.

**Survey demographics: **Every respondent confirmed decision-making or strong-influencer involvement in cloud, infrastructure, security, data governance, AI, or regulated-workload strategy in the past 12 months. Of the respondents, 73% are director level or above and the remaining are managers or senior/enterprise architects. At executive levels, 17% are c-level and 24% are VP or SVP. Respondents have long average tenures at their current organization, at 7.1 years, and they each personally influence or own an average annual budget of $124M (with 50% of respondents owning or influencing $20M to $178M of budgets).

**Survey precision and weighting**. Whole-sample margin of error is ±2.2% at the 95% confidence level before weighting. Regional margins range from ±3.1% in Asia-Pacific to ±6.9% in Latin America. Percentages are shown to one decimal where the source supports it and rounded otherwise. The achieved sample is 50.5% Asia-Pacific, so to avoid global results dominated by that region, the global all-respondent analysis weights the four regions equally: each contributes 25% of the result, and nothing from any region is discarded. Every segmented number – for a region, an industry, a company-size band, or a group defined by how it answered a question – is that segment's own and unweighted, with every respondent in it counting once.

**Analyst-led in-depth interviews with senior executives**. Twelve hour-long in-depth interviews with CIOs, CTOs, CISOs, CDIOs, and VP-level technology leaders across EMEA, Asia-Pacific, Latin America, and North America were conducted to augment the analysis with additional depth and context. Quotes and vendor names are anonymized and attributed by role, industry, and region.

# Information About This Report

## AUTHORS

**Nick Patience**
Vice President & Practice Lead,
AI Platforms | The Futurum Group

**Benjamin Brown**
Vice President, Custom Research | The Futurum Group

## PUBLISHER

**Futurum Research**

## INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

## CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, "The Futurum Group," and that the report was commissioned by Red Hat. Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

## LICENSING

This document, including any supporting materials, is owned by The Futurum Group. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of The Futurum Group.

## DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

Red Hat

## ABOUT RED HAT

Red Hat is the world's leading provider of enterprise open source software solutions, helping organizations innovate across hybrid cloud, automation, containers, and AI. With a community-powered approach, Red Hat delivers trusted, flexible technologies that enable businesses to modernize IT infrastructure and accelerate digital transformation.

Futurum®

## ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum®
**CONTACT INFORMATION:** The Futurum Group LLC | futurumgroup.com | (833) 722-5337

© 2026 The Futurum Group. All rights reserved.