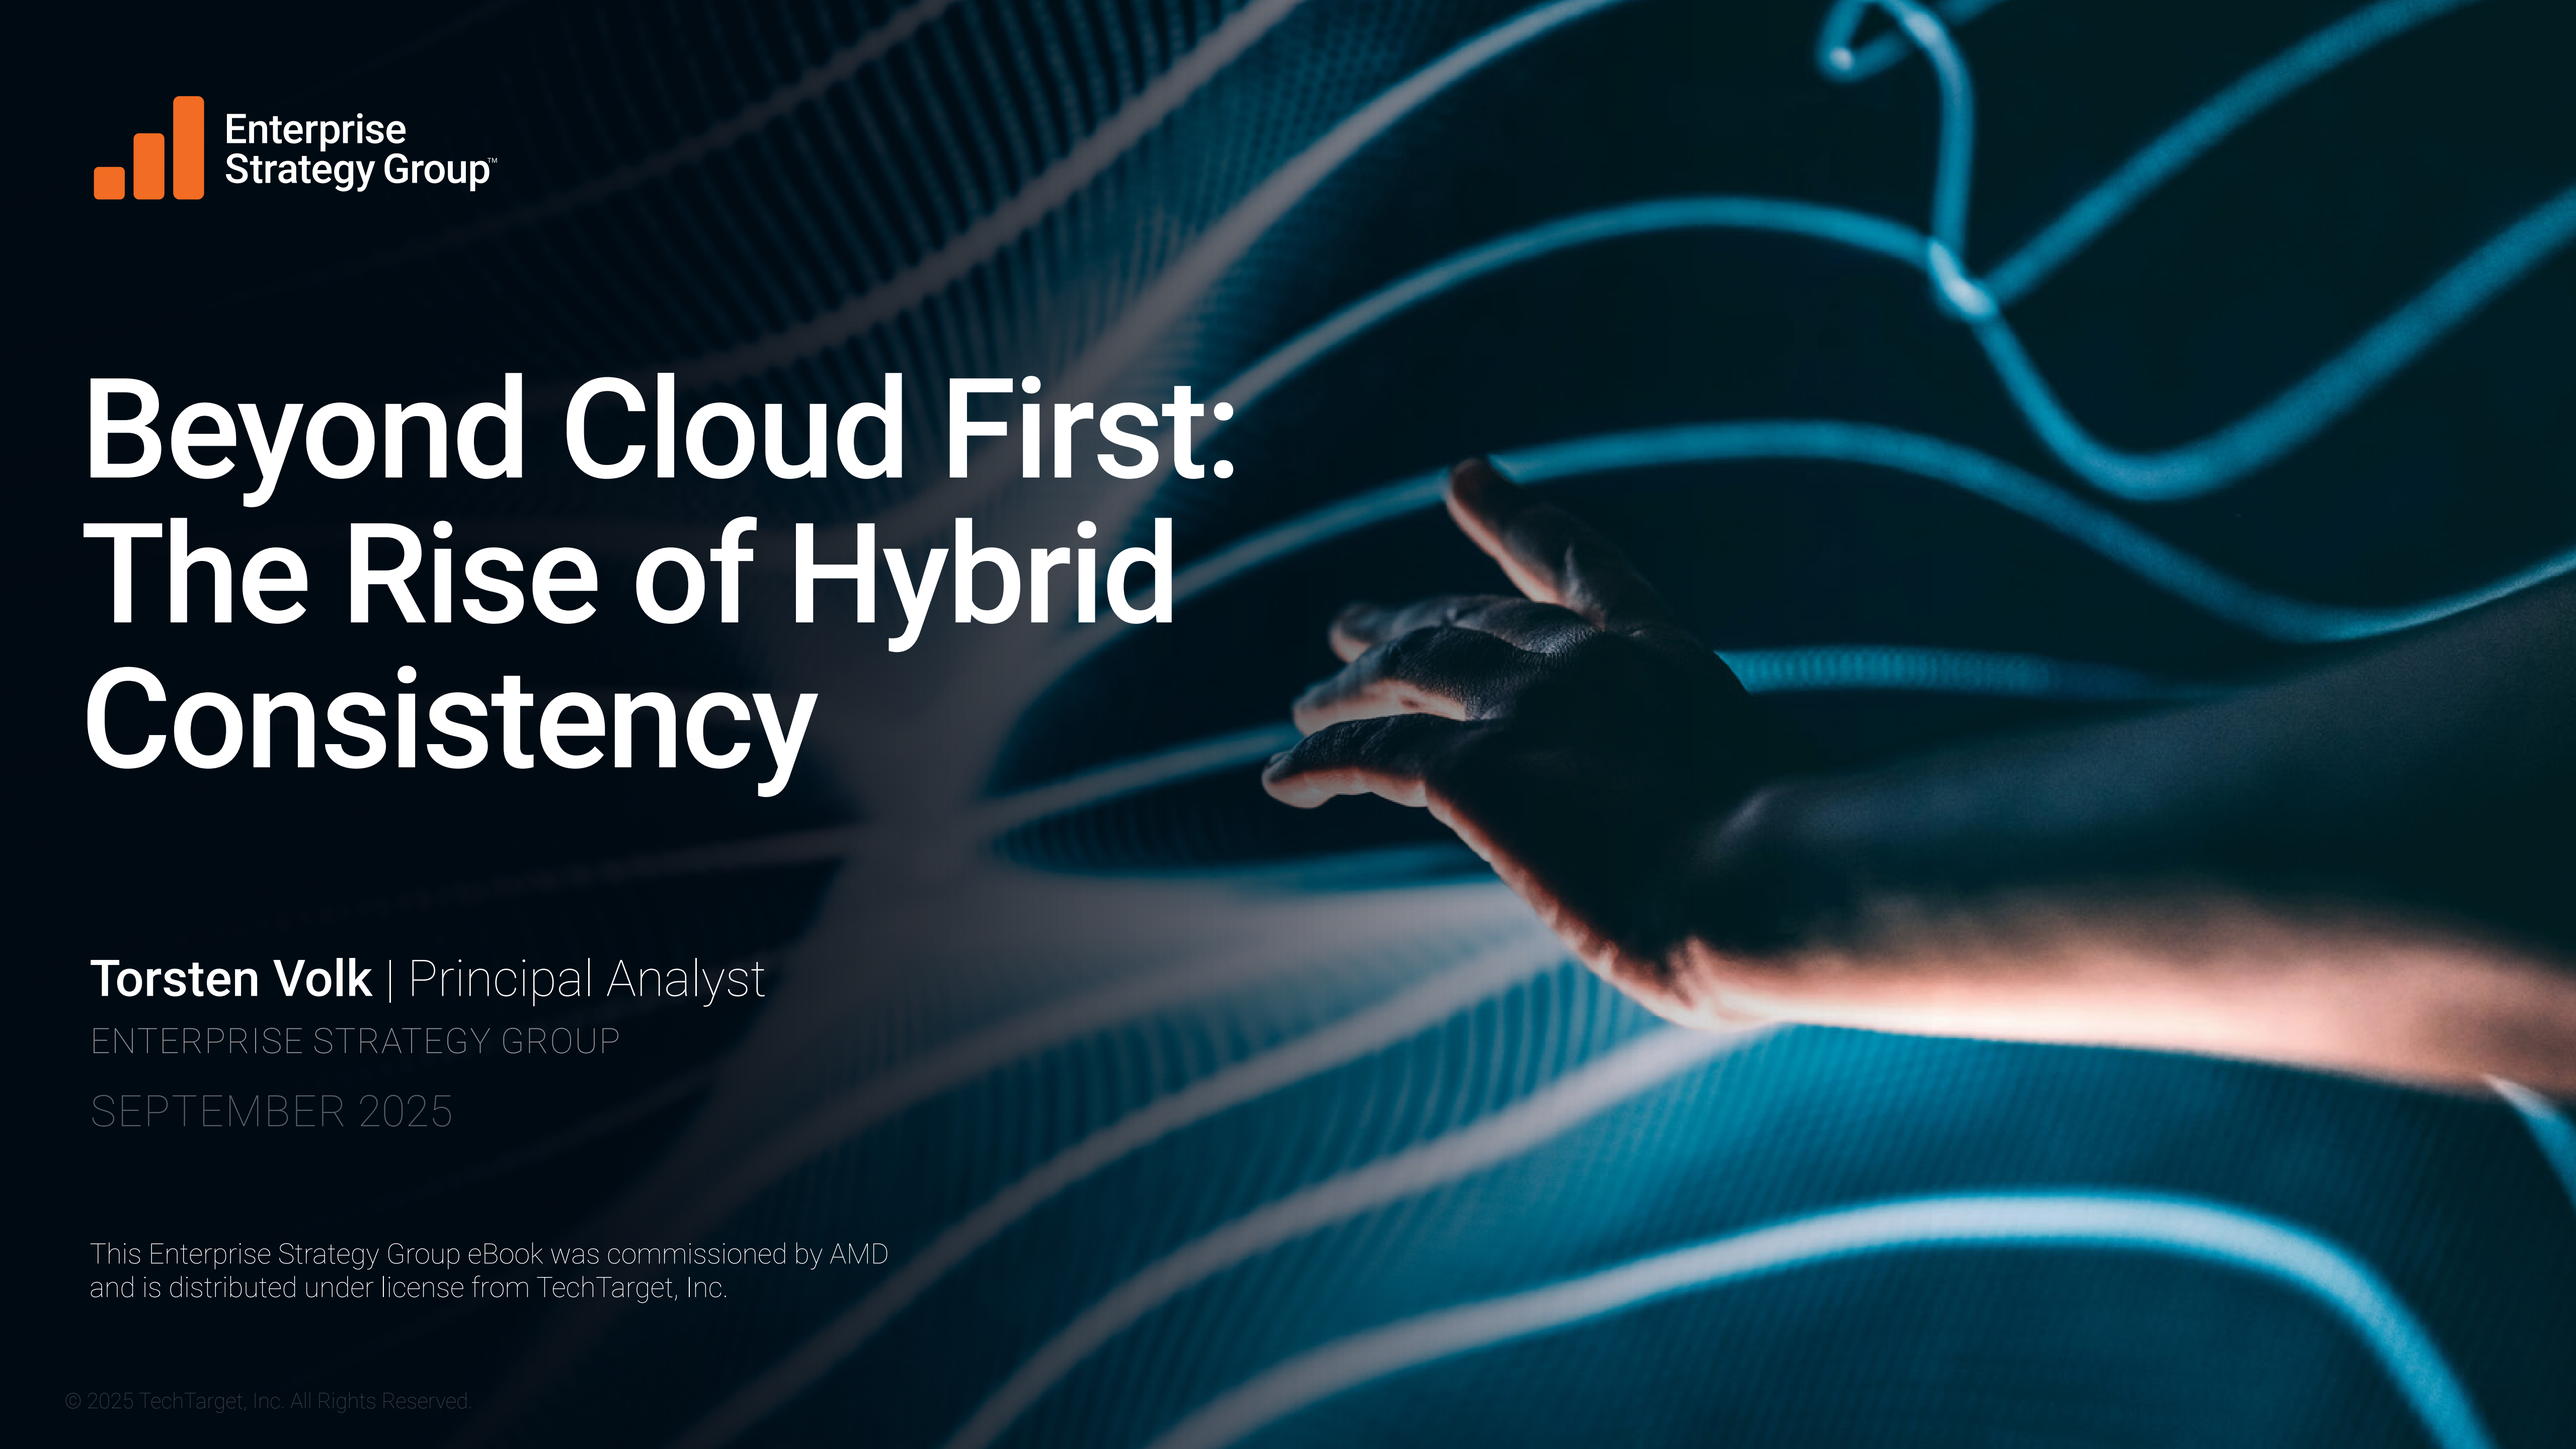

Enterprise Strategy Group™

# Beyond Cloud First: The Rise of Hybrid Consistency

**Torsten Volk** | Principal Analyst

ENTERPRISE STRATEGY GROUP

SEPTEMBER 2025

This Enterprise Strategy Group eBook was commissioned by AMD and is distributed under license from TechTarget, Inc.

# Research Objectives

This research examines how organizations navigate workload placement decisions across hybrid cloud environments, focusing on the operational realities and hidden costs of migrating between x86 and Arm CPU architectures. The study explores the technical challenges, ecosystem gaps, and unexpected expenses that impact total cost of ownership, while identifying how strategic factors, including AI support, security, compliance, and performance requirements, increasingly drive architecture decisions beyond simple cost calculations.

## THIS STUDY SOUGHT TO:

- **Assess** how organizations are placing workloads across cloud and on-premises environments and what drives those decisions.
- **Understand** the impact of hidden costs, such as monitoring gaps, staff retraining, and unexpected repatriation, on total cost of ownership (TCO).
- **Explore** the challenges and trade-offs of migrating between CPU architectures, including performance, ecosystem maturity, and operational complexity.
- **Determine** the strategic factors (e.g., AI support, security, compliance, performance) that most influence workload placement and architecture standardization.

# **Key** Findings

### Hybrid Architectures Replace One-size-fits-all Infrastructure

Organizations now balance workloads across cloud and on prem based on actual requirements, rather than defaulting to cloud vendor recommendations, leading to bi-directional movement as performance, compliance, and other business needs dictate.

PAGE 4

### Performance-critical Workloads Drive Cloud Repatriation

HPC, AI, and analytics workloads are increasingly moving back to on-premises environments to access specialized hardware and ensure reliable performance.

PAGE 8

### No Single Cloud Can Do It All

Enterprises typically use 2–3 unique public cloud infrastructure providers, mixing services for specific workloads to avoid lock-in and maximize flexibility.

PAGE 10

### Hybrid Cloud Becomes the Default

Two-thirds of organizations have repatriated some workloads, pointing to a trend that hybrid is the enduring model for optimizing performance, compliance, and cost.

PAGE 15

### Deployment Decisions Reflect the Full Business Context

Workload placement now involves security, finance, and developers, ensuring operational, financial, and technical factors outweigh simplistic “cloud-first” policies.

PAGE 20



BACK TO CONTENTS

# Hybrid Architectures Replace One-size-fits-all Infrastructure

Organizations now balance workloads across cloud and on prem based on actual requirements, rather than defaulting to cloud vendor recommendations, leading to bi-directional movement as performance, compliance, and other business needs dictate.

# Hybrid Architectures Are Replacing One-size-fits-all Infrastructure

Organizations have evolved from having simplistic cloud-first strategies to utilizing sophisticated hybrid architectures where workload placement decisions are made collaboratively across executive, technical, and financial teams based on actual workload requirements rather than cloud vendor recommendations and preferences. This flexibility is enabled in large part by x86's ecosystem breadth; its consistent tooling and monitoring across all major cloud providers enables organizations to move workloads freely without retooling their operations, making true multi-cloud optimization achievable rather than aspirational.

## No Single Cloud Can Do It All

Organizations are increasingly navigating a complex multi-cloud landscape, averaging two to three providers simultaneously[1] and employing diverse migration strategies from re-platforming existing applications to completely replacing legacy systems with cloud-native alternatives. This multi-vendor approach reflects a mature understanding that no single cloud provider excels at everything, driving enterprises to mix and match services based on specific workload requirements, while maintaining the flexibility to shift between platforms as their needs evolve.

**Common Approaches to Migrating Workloads to the Public Cloud.**

67%
OF ORGANIZATIONS

Re-platform (lift, tinker, and shift): Migrate existing workloads to the cloud with selective optimizations and/or architectural changes

62%
OF ORGANIZATIONS

Replace in cloud: Replace on-premises workloads with SaaS or off-premises hosted versions or substitutes

57%
OF ORGANIZATIONS

Rehost (lift and shift): Migrate existing workloads to the cloud with few or no code changes

54%
OF ORGANIZATIONS

Refactor and shift: Reengineer existing workloads for cloud and shift to off-premises cloud environments



[1] Source: Enterprise Strategy Group Complete Survey Results, *Application Modernization and the Role of Platform Engineering*, September 2024.

BACK TO CONTENTS

# Hybrid Cloud Becomes Mainstream

Among surveyed organizations, 66% have repatriated at least one workload from the public cloud to an on-premises environment. This shows that the hybrid cloud operating model reigns supreme to make it easy for organizations to place their workloads into the location with the optimal characteristics for performance, compliance, and cost at any given time, while maintaining the flexibility to relocate those workloads as business requirements, regulatory demands, or technology capabilities evolve.

**Repatriation of a Workload From a Public Cloud to an On-premises Environment Over the Past 24 Months.**




BACK TO CONTENTS

# Deployment Decisions Reflect the Full Business Context

Deployment decisions are based on a broad evaluation of the workload's characteristics with numerous staff roles involved—from security and compliance teams assessing risk, to finance teams evaluating TCO, to developers considering technical requirements. While executive IT leadership typically makes the final call, this cross-functional collaboration ensures decisions account for operational, financial, and technical implications rather than following simplistic mandates. This illustrates that organizations are moving away from one-size-fits-all policies like "cloud first" and are instead placing workloads where they deliver the greatest overall value across all dimensions.

**Deployment Decision Involvement (Excluding Not Involved & Don't Know).**

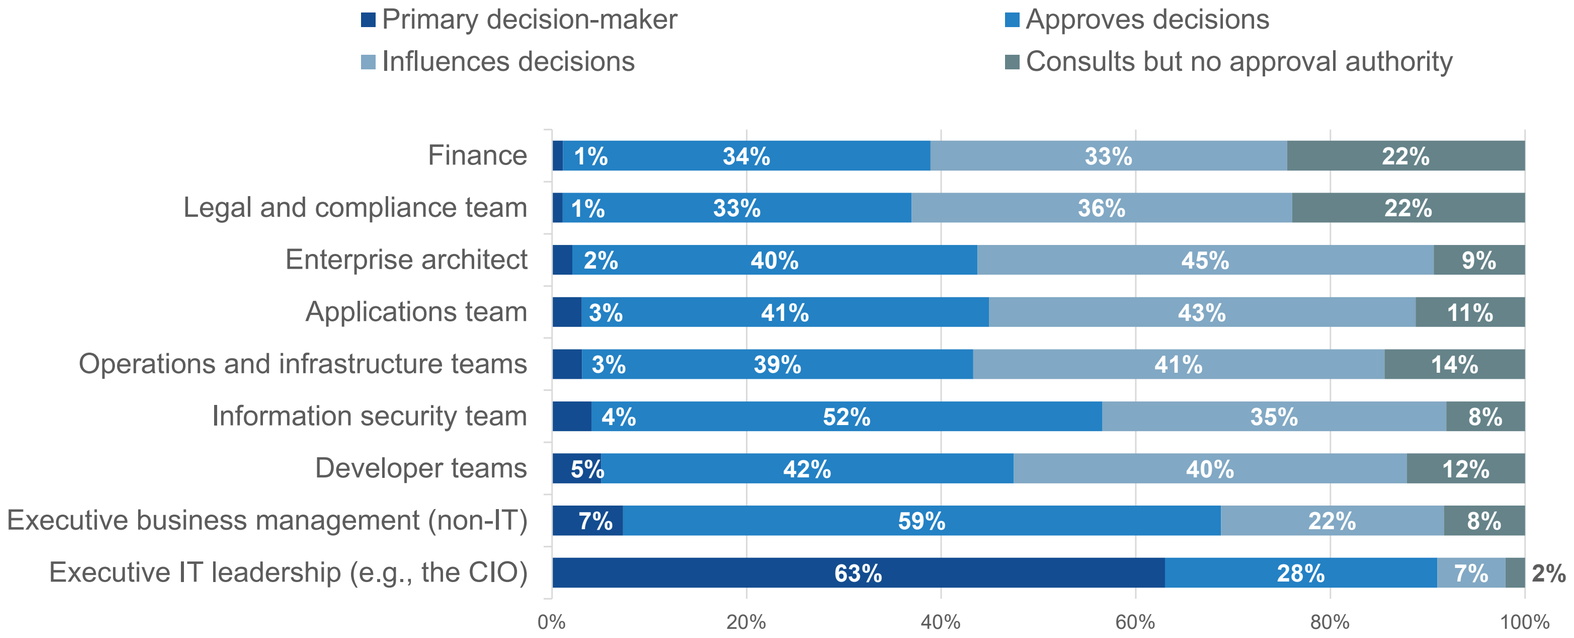




BACK TO CONTENTS

# Performance-critical Workloads Drive Cloud Repatriation

HPC, AI, and analytics workloads are increasingly moving back to on-premises environments to access specialized hardware and ensure reliable performance.

# Performance-critical Workloads Drive Cloud Repatriation

Organizations are increasingly moving high-performance workloads, such as HPC, AI training, and data analytics, back on premises, as cloud does not meet their needs.

**Types of Workloads Repatriated From the Public Cloud Back to On-premises Environments in the Past 24 Months.**

| | | | |
|---|---|---|---|
| 53% High-performance computing (HPC) | 49% AI model training | 46% Security- or compliance-critical workloads | 40% AI interference |
| 39% Database workloads | 37% Data-intensive analytics workloads | 36% Applications with steady or predictable resource usage | 32% Storage-intensive workloads |
| 31% Core business applications for strategic control | 30% Cost-prohibitive workloads (high cloud expenses) | 24% Legacy applications with licensing restrictions | 24% Latency-sensitive applications |

BACK TO CONTENTS

# No Single Cloud Can Do It All

Enterprises typically use 2–3 unique public cloud infrastructure providers, mixing services for specific workloads to avoid lock-in and maximize flexibility.

# Beyond Cheap Instances: The True Cost of Cloud Choices

Organizations prioritize AI capabilities, security, and performance over raw cost savings when choosing cloud platforms, understanding that the lowest hourly instance price often doesn't mean the best overall value. Even though cloud providers push their Arm-based options through native monitoring tools and pricing incentives, many companies stick with x86 architecture.

## The Majority of Workloads Run on x86

Seventy-one percent of organizations expect to run at least half of their workloads on x86 infrastructure three years from now. This will enable them to freely deploy, move, and scale applications within hybrid cloud environments. Additionally, they will be able to make deliberate placement decisions based on each specific workload's requirements profile and set of resource utilization characteristics.

**What proportion of your workloads do you expect to run in the public cloud on x86 architecture versus Arm architecture three years from now?**

**4%**
All our workloads will run on x86 architecture

**24%**
Most of our workloads will run on Arm architecture, while some will run on x86 architecture

**33%**
Most of our workloads will run on x86 architecture, while some will run on Arm architecture

**4%**
All our workloads will run on Arm architecture

**34%**
Our workloads will be equally split between x86 and Arm architecture



BACK TO CONTENTS

## Beyond Cost: Strategic Factors Drive Application Placement

Organizations are no longer evaluating cloud platforms primarily on price. Instead, support for AI initiatives (51%), security (50%), and performance (44%) outpace cost considerations (39%). This indicates a strategic shift: While cost remains a critical factor, enterprises increasingly prioritize capabilities that directly impact innovation, resilience, and business outcomes. In other words, cost efficiency matters, but without the right AI, security, and performance features, even a cheaper cloud option might fall short.

**Reasons for Evaluating a New Public Cloud Platform to Replace Current Platform.**

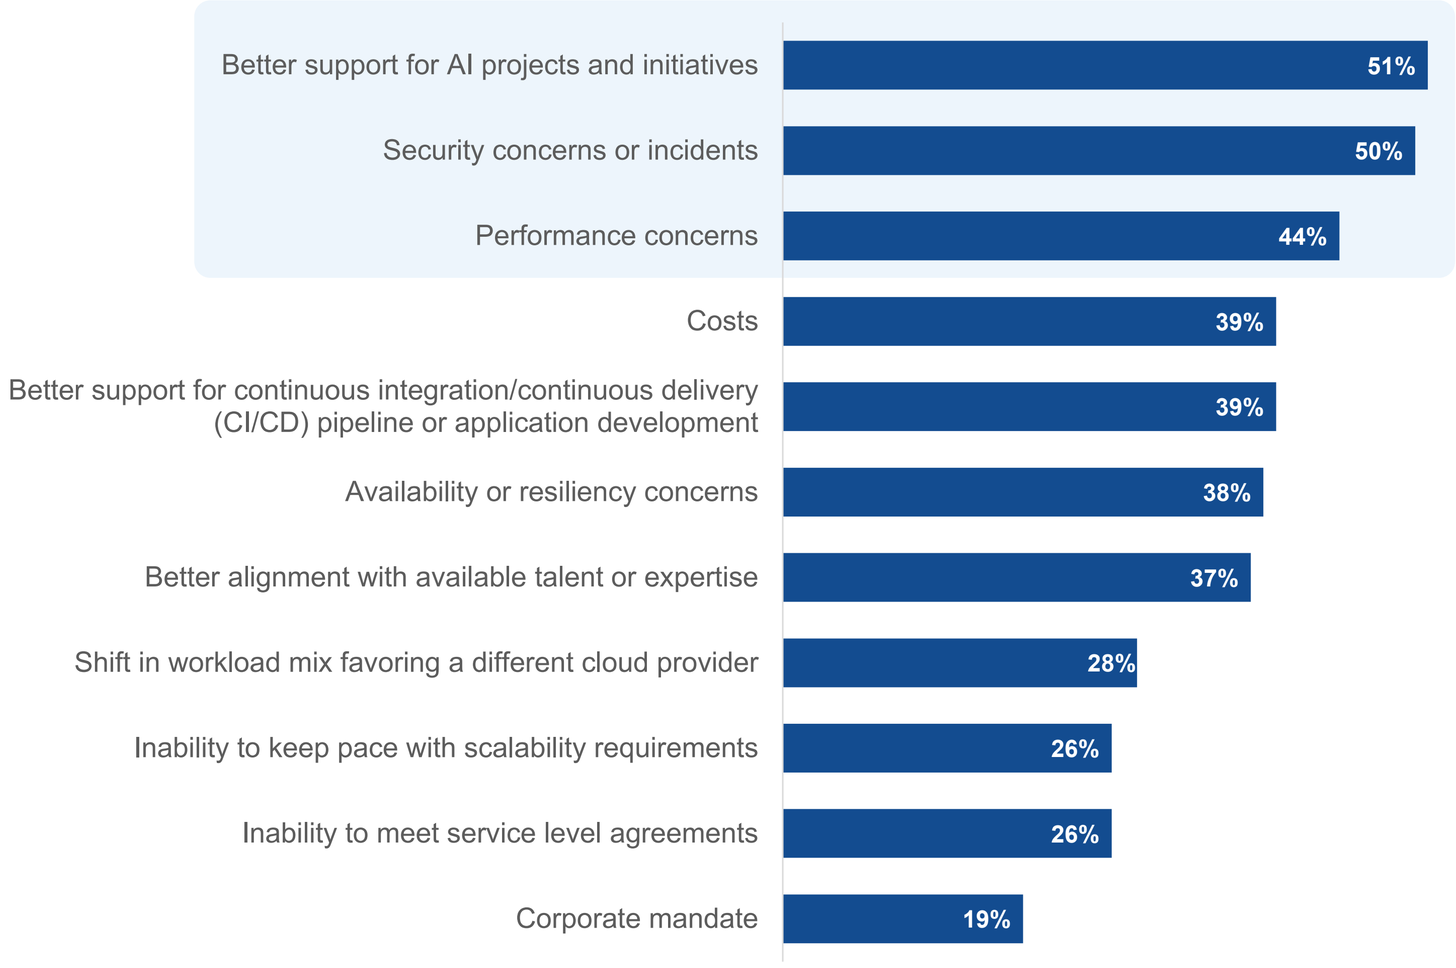




BACK TO CONTENTS

# Ecosystem Maturity Shapes Monitoring Choices

Organizations monitor Arm-based instances differently from x86, with heavier reliance on cloud provider dashboards (34% vs. 16%). This suggests that monitoring ecosystems for Arm are still maturing, while x86 workloads more often integrate with well- established enterprise APIs and tools. These differences highlight how monitoring practices evolve alongside architecture choices.

**Primary Access Point for Performance Metrics for Arm-based and x86-based Public Cloud Instances.**

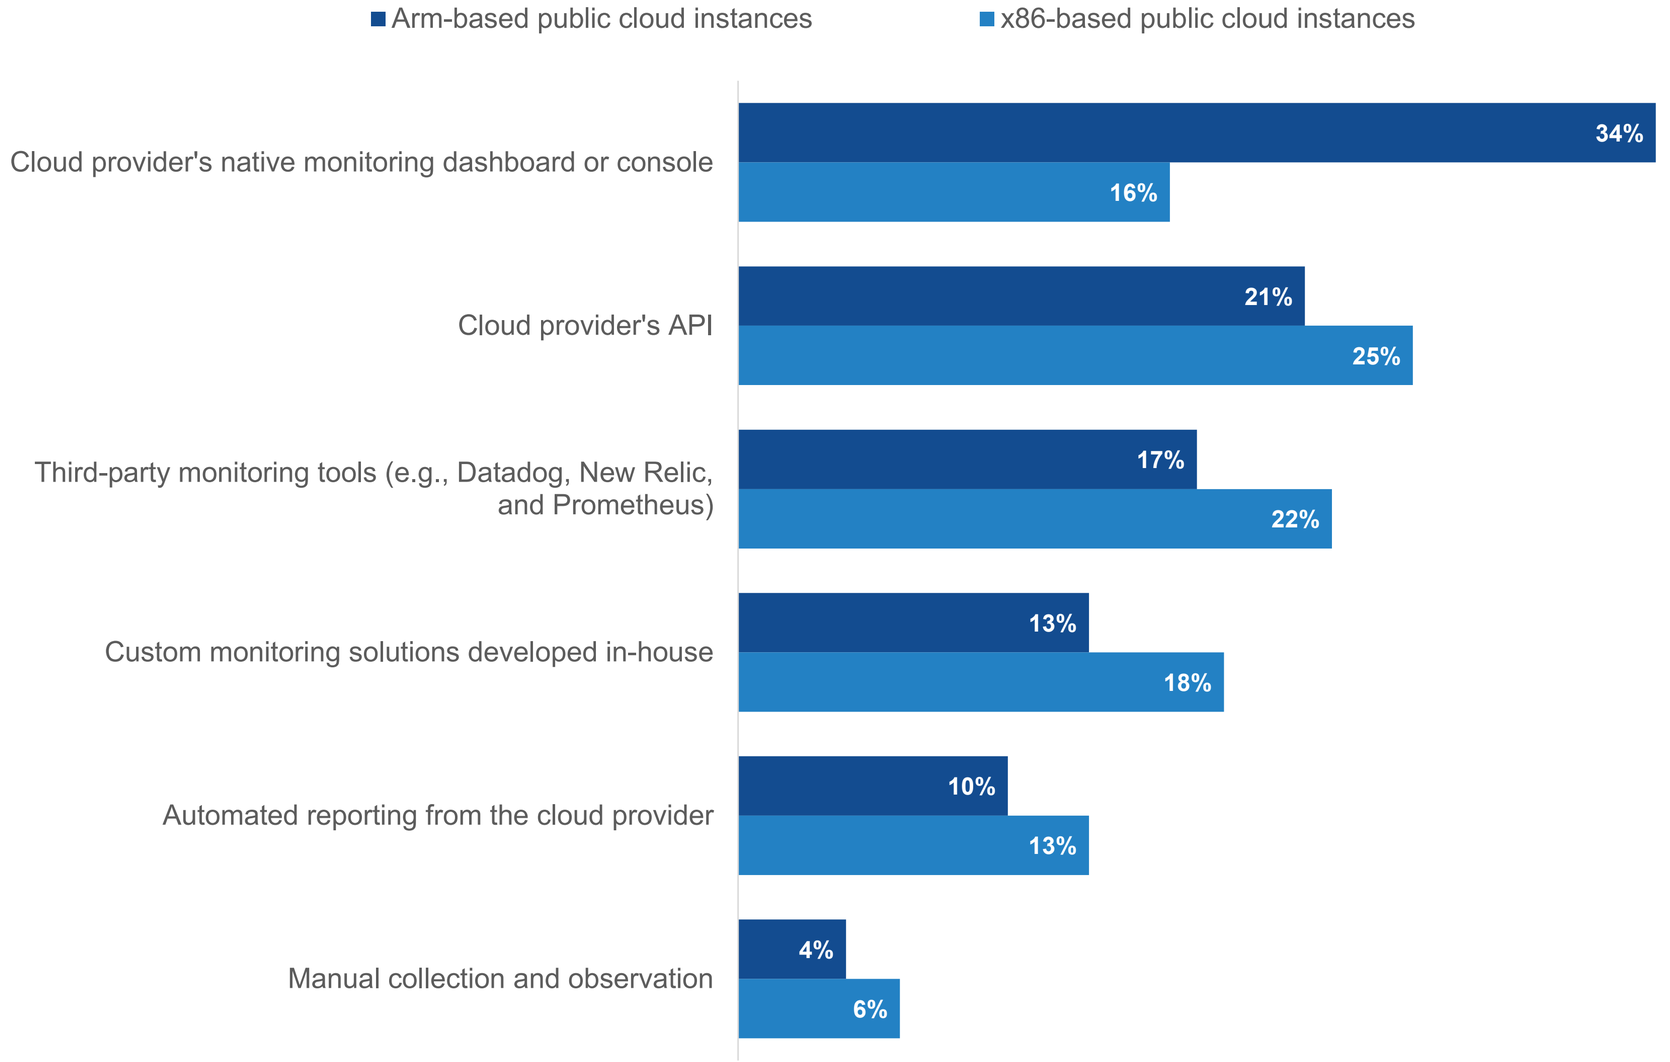




BACK TO CONTENTS

# Cloud Bills Can Be Surprising

More than half of organizations reported cloud bills coming in higher than expected, highlighting a fundamental tension: Cloud promises cost efficiency and pay-as-you-go flexibility, yet those advantages depend on mature financial operations, accurate forecasting, and disciplined workload management. Many enterprises are still climbing that learning curve, which often leads to higher-than-expected spending.

**Public Cloud Bills Are Largely Higher Than Initial Budget or Forecast.**

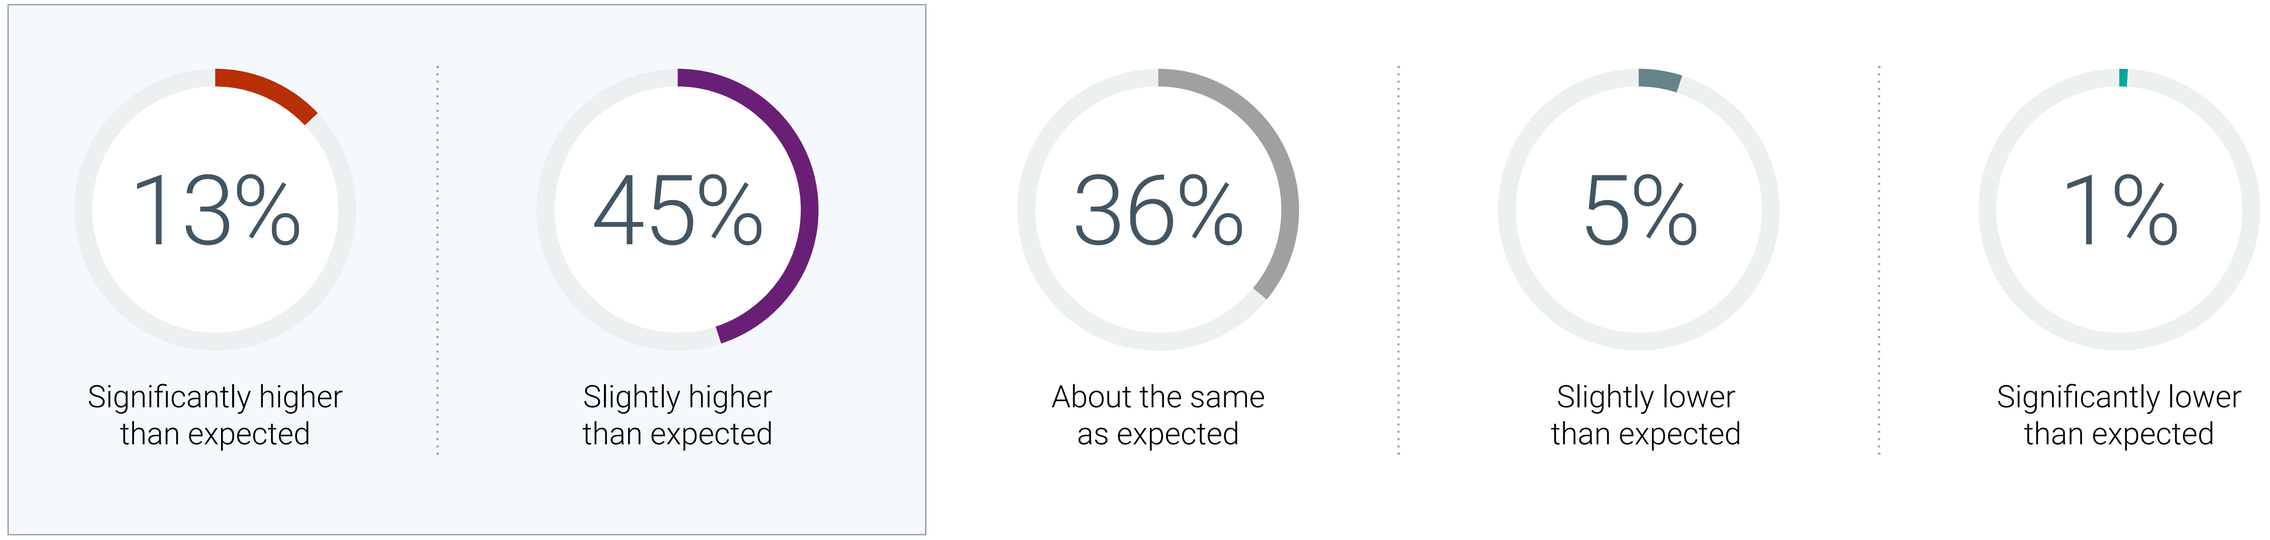




BACK TO CONTENTS

# Hybrid Cloud Becomes the Default

Two-thirds of organizations have repatriated some workloads, pointing to a trend that hybrid is the enduring model for optimizing performance, compliance, and cost.

## The Hidden Complexity and Cost of Architecture Transformation

Migrating applications between CPU architectures involves far more than recompiling code. It requires rebuilding entire continuous integration/continuous delivery (CI/CD) pipelines, replacing libraries, reconfiguring monitoring tools, and retraining staff while managing the risk of—not always subtle—performance and compatibility issues. Organizations consistently underestimate this complexity, discovering that Arm migrations prove significantly more difficult than anticipated (see Arm Migrations Are 1.5x Likelier to be Harder [vs. Easier] Than Expected). This is due to ecosystem gaps, unexpected technical challenges, and the cascading impact of architectural changes across their entire operational stack.

## Migration Between CPU Architectures Involves Significant Risk and Effort

Migrating an application to a new CPU architecture—whether from x86 to Arm or the other way around—goes far beyond a simple recompile. In fact, 77% of organizations reported that porting x86 software to Arm would require significant effort, highlighting the complexity, time, and potential risks involved. As a result, approximately one-third of organizations are hesitant to adopt new architectures.

**Current Usage of x86-based Software That Would Require Significant Effort to Port to Arm.**


**Challenges Encountered When Migrating Workloads to Arm-based Public Cloud Instances.**

37%

unexpected performance differences

32%

data-related complications

31%

Technical compabilitiy problems



BACK TO CONTENTS

# The CPU Architecture Transformation Process Between Arm and x86 Is Complex

Rebuilding an application often includes source code changes, library replacements, and adjustments to the build pipeline. In some cases, this means compiling the full application from source; in others, it might involve swapping out CPU-specific libraries, plugins, or third-party binaries.

The process often extends into the CI/CD domain, where teams might maintain separate pipelines for each architecture. Infrastructure-as-code definitions need architecture labels, and test environments must reflect CPU-specific conditions. Architecture-aware testing can uncover subtle bugs that surface only because of timing or concurrency differences between platforms.

**CPU Architecture Transformation Workflow**

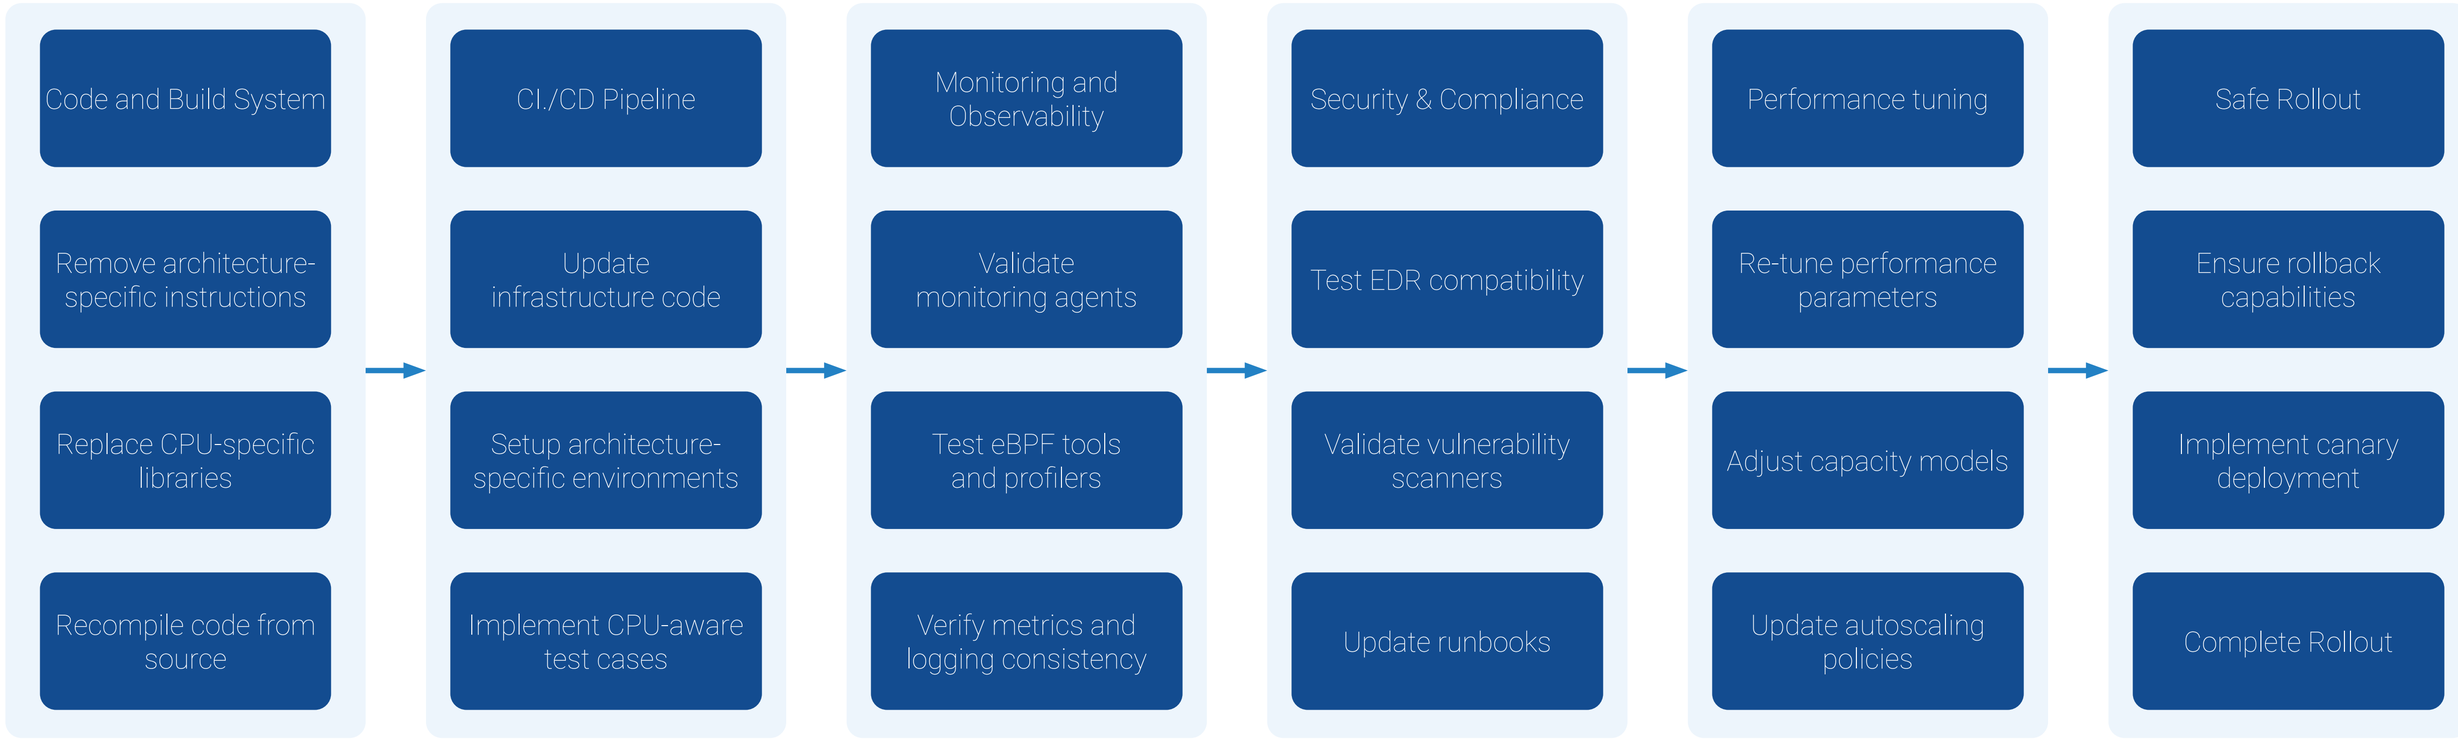




BACK TO CONTENTS

# Arm Migrations Are 1.5x Likelier to be Harder (vs. Easier) Than Expected[2]

Organizations consistently underestimate Arm migration complexity. This is not surprising, considering the foundational challenges these organizations have experienced during their migration process. From a lack of staff skills (24%) to limitations in the software ecosystem (40%), these challenges can make the migration process more difficult.

**Arm Migration Challenges.**

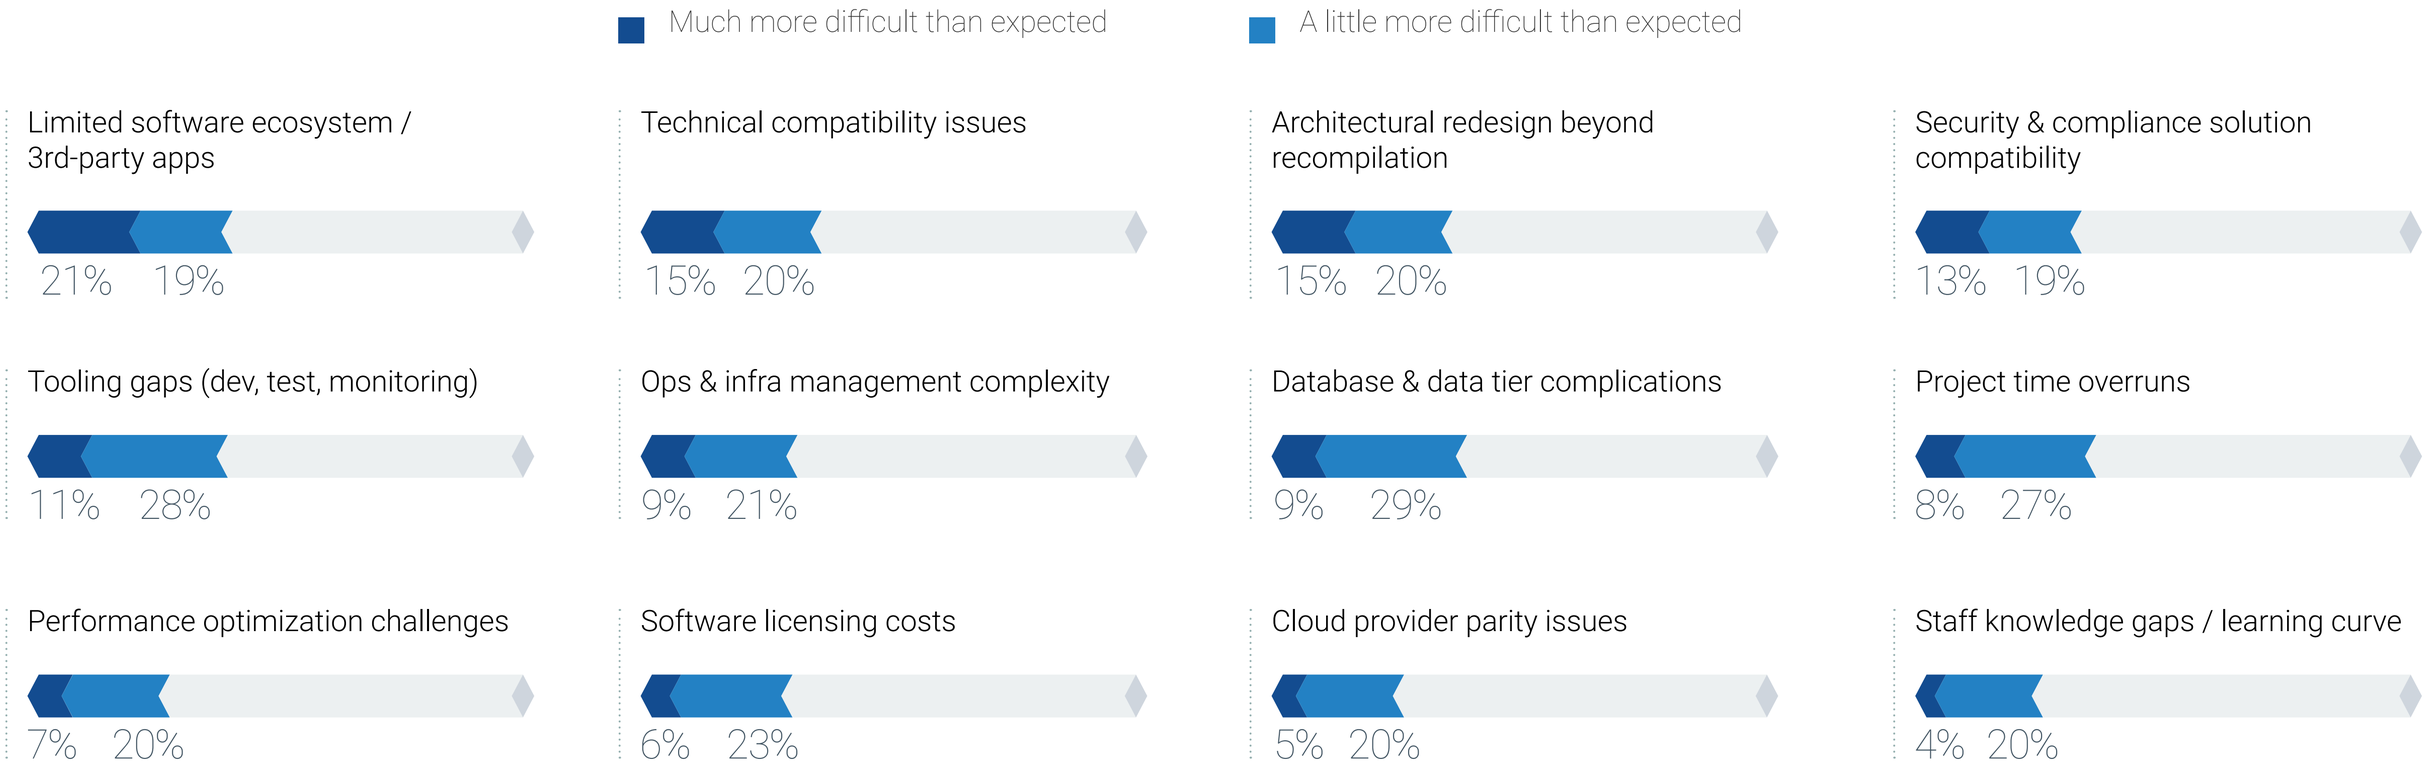


[2] The 1.5x claim was calculated as follows: Across the 12 challenge categories measured, the aggregated data on page 43 shows:

$$\frac{\text{Average responding "more difficult than expected" (~40%)}}{\text{Average responding "easier than expected" (~27%)}} \approx 1.5x$$



BACK TO CONTENTS

# Observability, Security, and Database Checks Are Non-negotiable

The observability stack, including agents, eBPF tools, and profilers, must be architecturally compatible across platforms. Security solutions like EDR and vulnerability scanners also require validation, while database and data-tier issues remain a primary migration challenge. These activities are already standard practice in on-premises environments, and enterprises expect cloud platforms to meet the same bar. To achieve this, organizations must extend compatibility testing, adjust capacity models and autoscaling policies, and fine-tune performance, while planning rollouts with safe fallback options to minimize risk.

**Challenges with Arm Migration.**

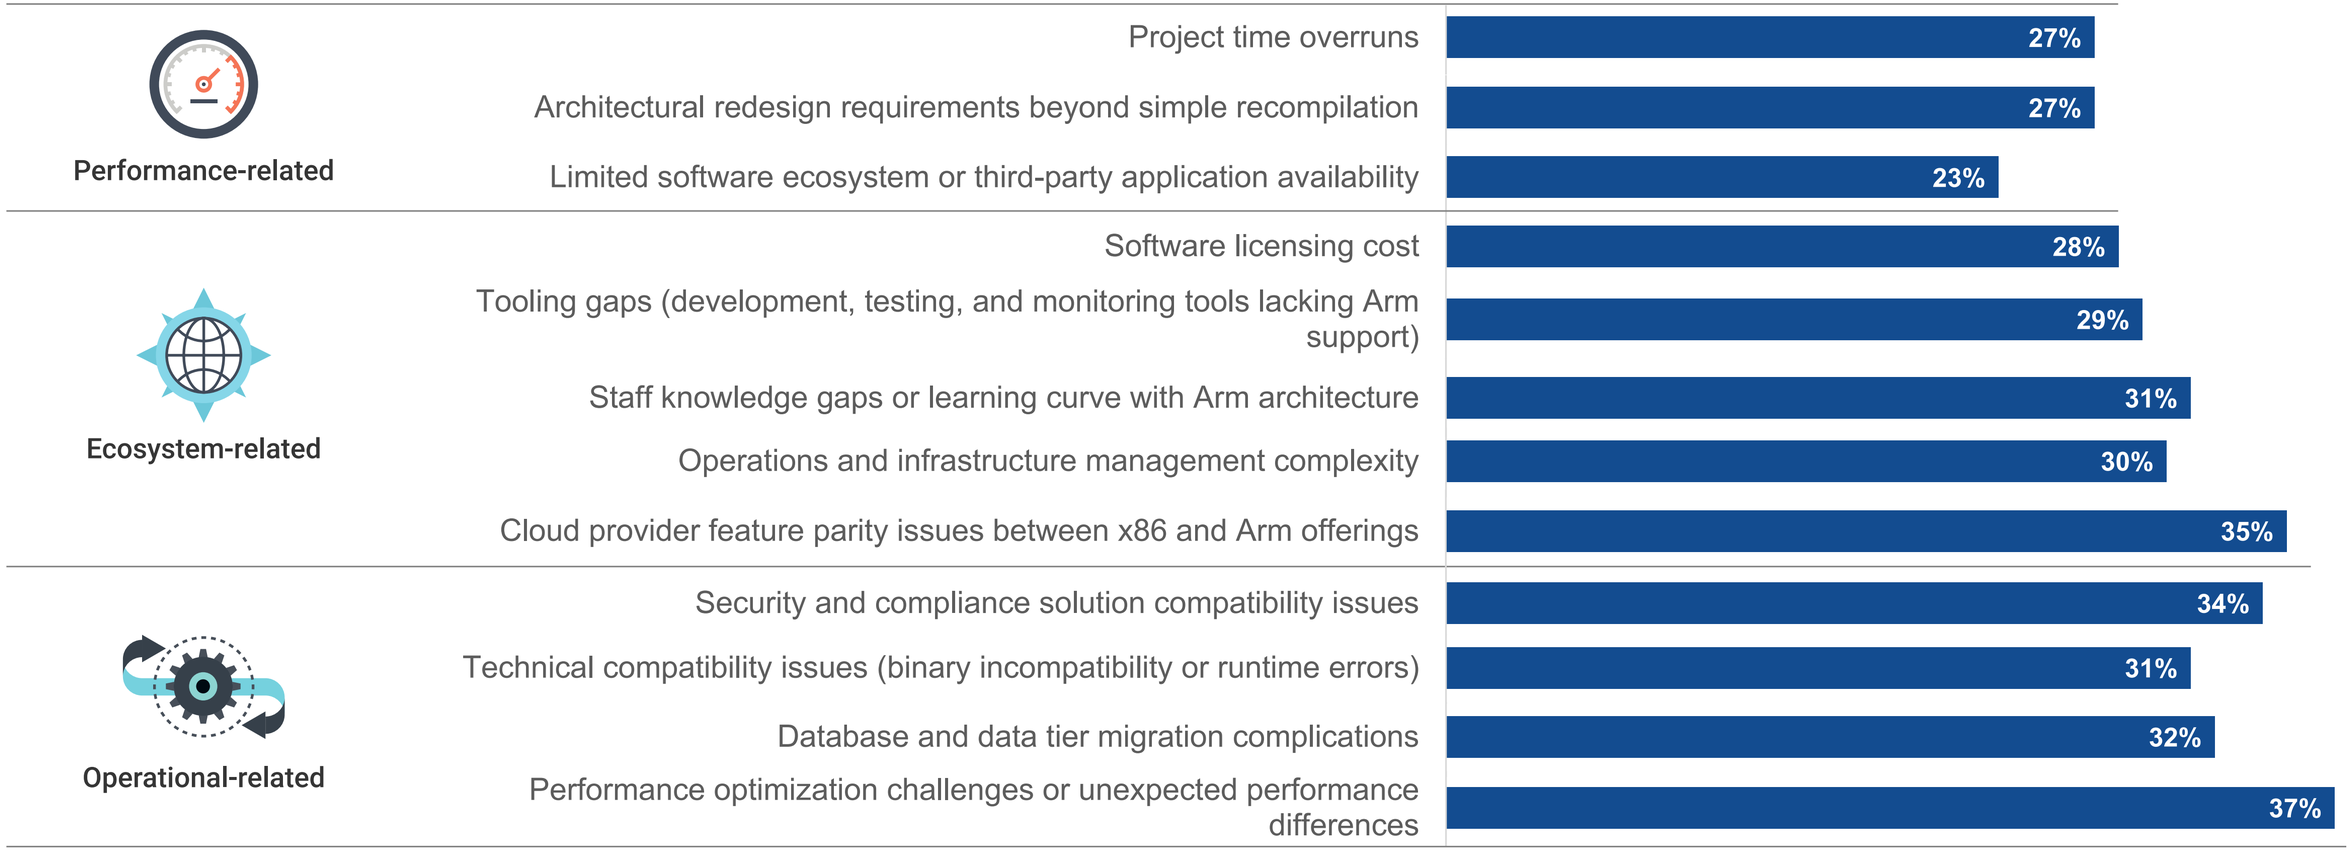




BACK TO CONTENTS

# Deployment Decisions Reflect the Full Business Context

Workload placement now involves security, finance, and developers, ensuring operational, financial, and technical factors outweigh simplistic “cloud-first” policies.

# The Great Arm Reversal

A significant number of organizations are already migrating workloads back to x86 (see Rolling Back From Arm to x86), motivated by higher-than-expected operational costs, performance issues, and integration difficulties. This pattern underscores a key lesson: Maintaining operational consistency across hybrid environments often matters more than potential, tactical cost savings. For most enterprises, architectural standardization is becoming a strategic driver of efficiency and resilience, rather than a sign of technical conservatism.

## Reevaluation Within 12 Months: Agility or Placement Risk?

Workload placement decisions are rarely permanent; 55% of organizations reevaluated placement within the first year. This reflects both sides of the cloud-value proposition: agility to adapt quickly to changing business needs and to leverage the near-constant introduction of new instance types and shapes powered by more capable CPUs and GPUs, while also highlighting the risk that initial placement could soon be suboptimal, making regular reevaluation a best practice.

**Workload Migration Reevaluation Within the First 12 Months.**

55%
Cloud-to-Cloud

45%
Cloud-to-on-premises

**Based on your experience, how long do workloads typically remain in a public cloud before being reevaluated for migration to either of the following locations?**

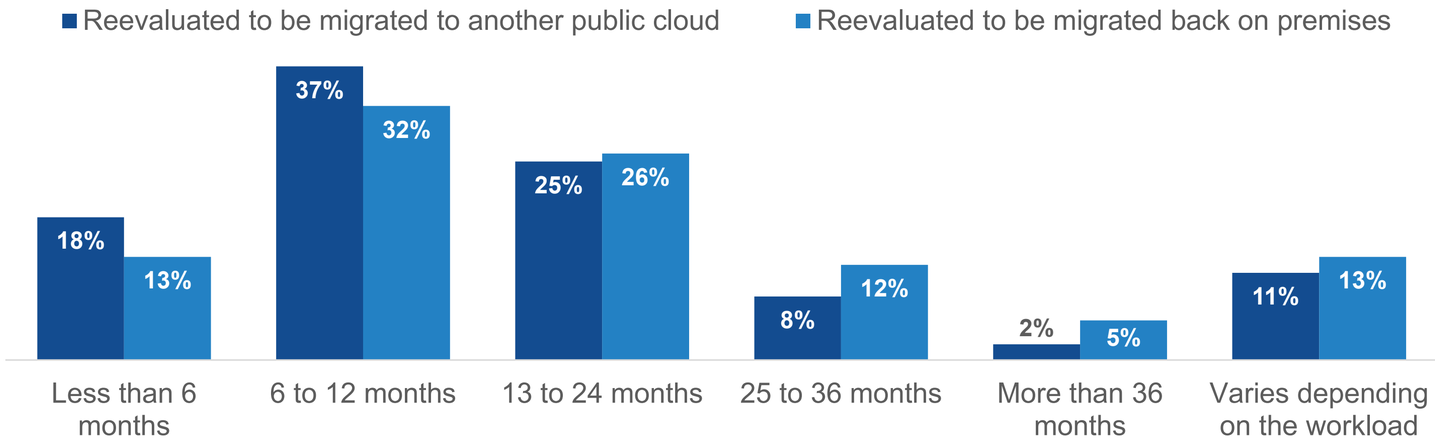




BACK TO CONTENTS

# Rolling Back From Arm to x86

Eighty-two percent of organizations have already moved workloads from Arm back to x86 cloud instances, with 78% planning more reversals in the next 24 months. On average, organizations are migrating 36% of their Arm workloads back. This can be seen as a significant course correction at scale.

**Arm Rollbacks at Scale (Past & Future)**

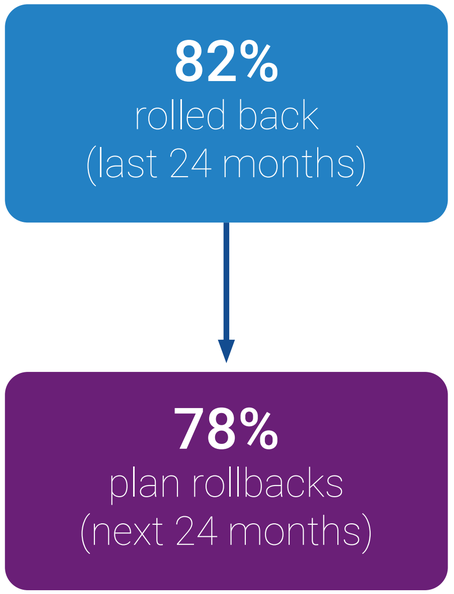


**Has your organization migrated workloads deployed on Arm-based public cloud instances to the following locations over the past 24 months?: x86-based public cloud instances.**

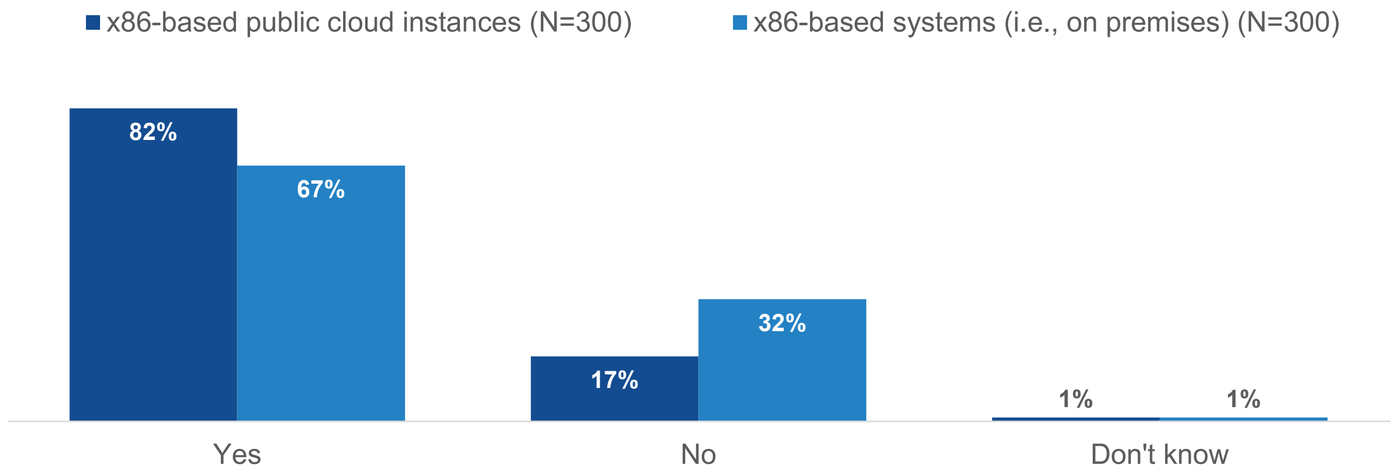


**Does your organization plan to migrate workloads deployed on Arm-based public cloud instances to the following locations over the next 24 months?**

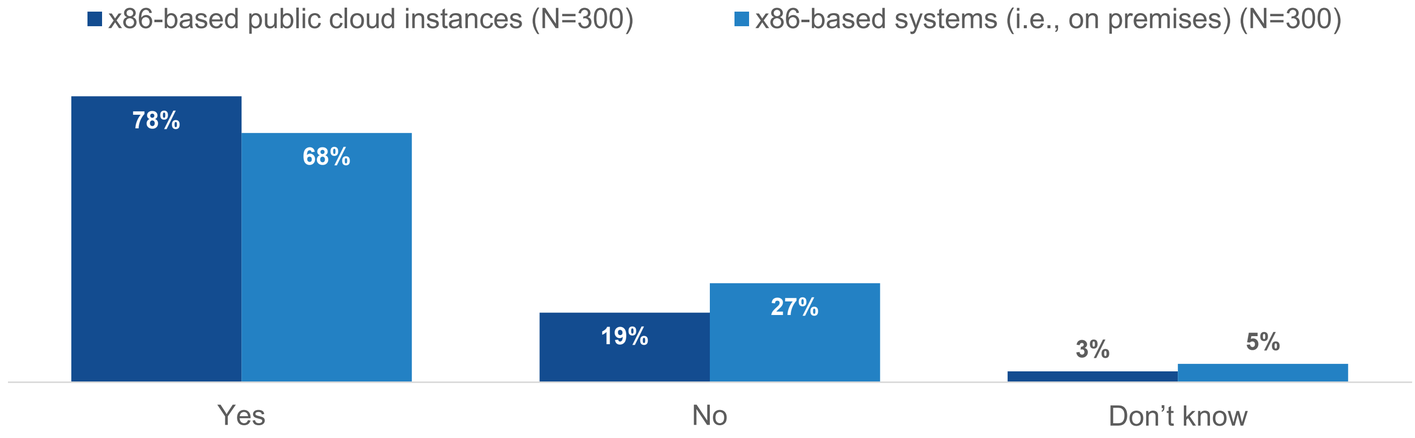




BACK TO CONTENTS

# Migration Triggers

The study revealed that the same challenges organizations encounter during Arm migrations—performance tuning, compatibility issues, and ecosystem gaps—are also the main drivers for moving workloads back to x86. Higher-than-expected operational costs (36%), networking complexities (33%), and performance limitations (31%) all directly erode Arm's cost-savings promise. These recurring challenges explain why organizations often reassess early and return to x86's mature ecosystem, where operational consistency and vendor support reduce risk and help stabilize TCO.

## Reasons for Migrating From Arm to x86.

**TCO Model Reevaluation**
Arm theoretical savings can be offset by performance limitation, and subsequent productivity inefficiencies and scaling difficulties indicate Arm's value proposition are overstated.

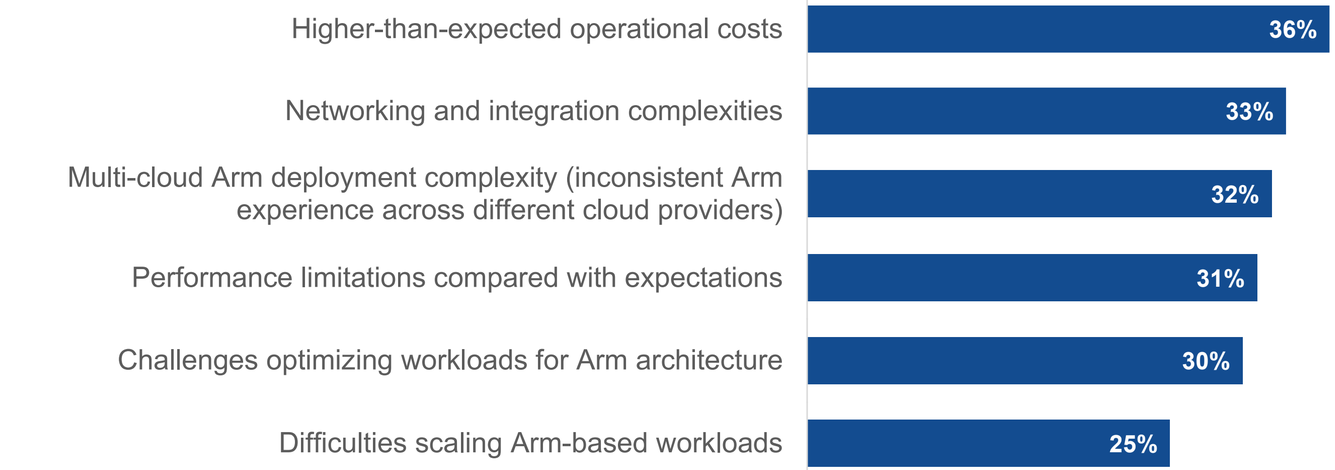


**Organizational Readiness**
Growing influence of practical outcomes leading to reevaluations of workload placement indicates organizations need realistic assessment before migration decisions.

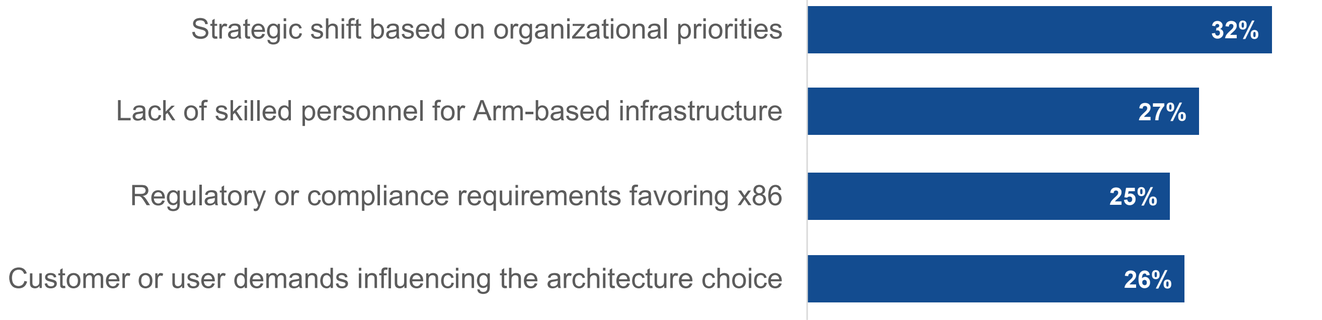


**Ecosystem Fragmentation**
Strengthening third-party relationship to re-emphasize the lagging maturity of the Arm partnership ecosystem as well as emphasizing x86 ecosystem compatibility.

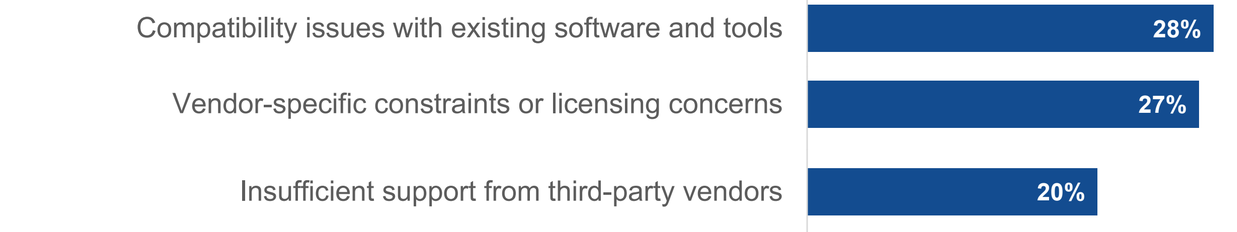




BACK TO CONTENTS

# Majority Said Consistent Architecture Simplifies IT Operations

Organizations have learned that architectural diversity has costs: 83% believed consistent architecture across on-premises and cloud simplifies operations, 70% wanted to avoid refactoring during migrations, and 68% recognized that using different CPU architectures adds operational overhead. In the real world, boring consistency beats exciting complexity.

**Architecture Sentiments.**

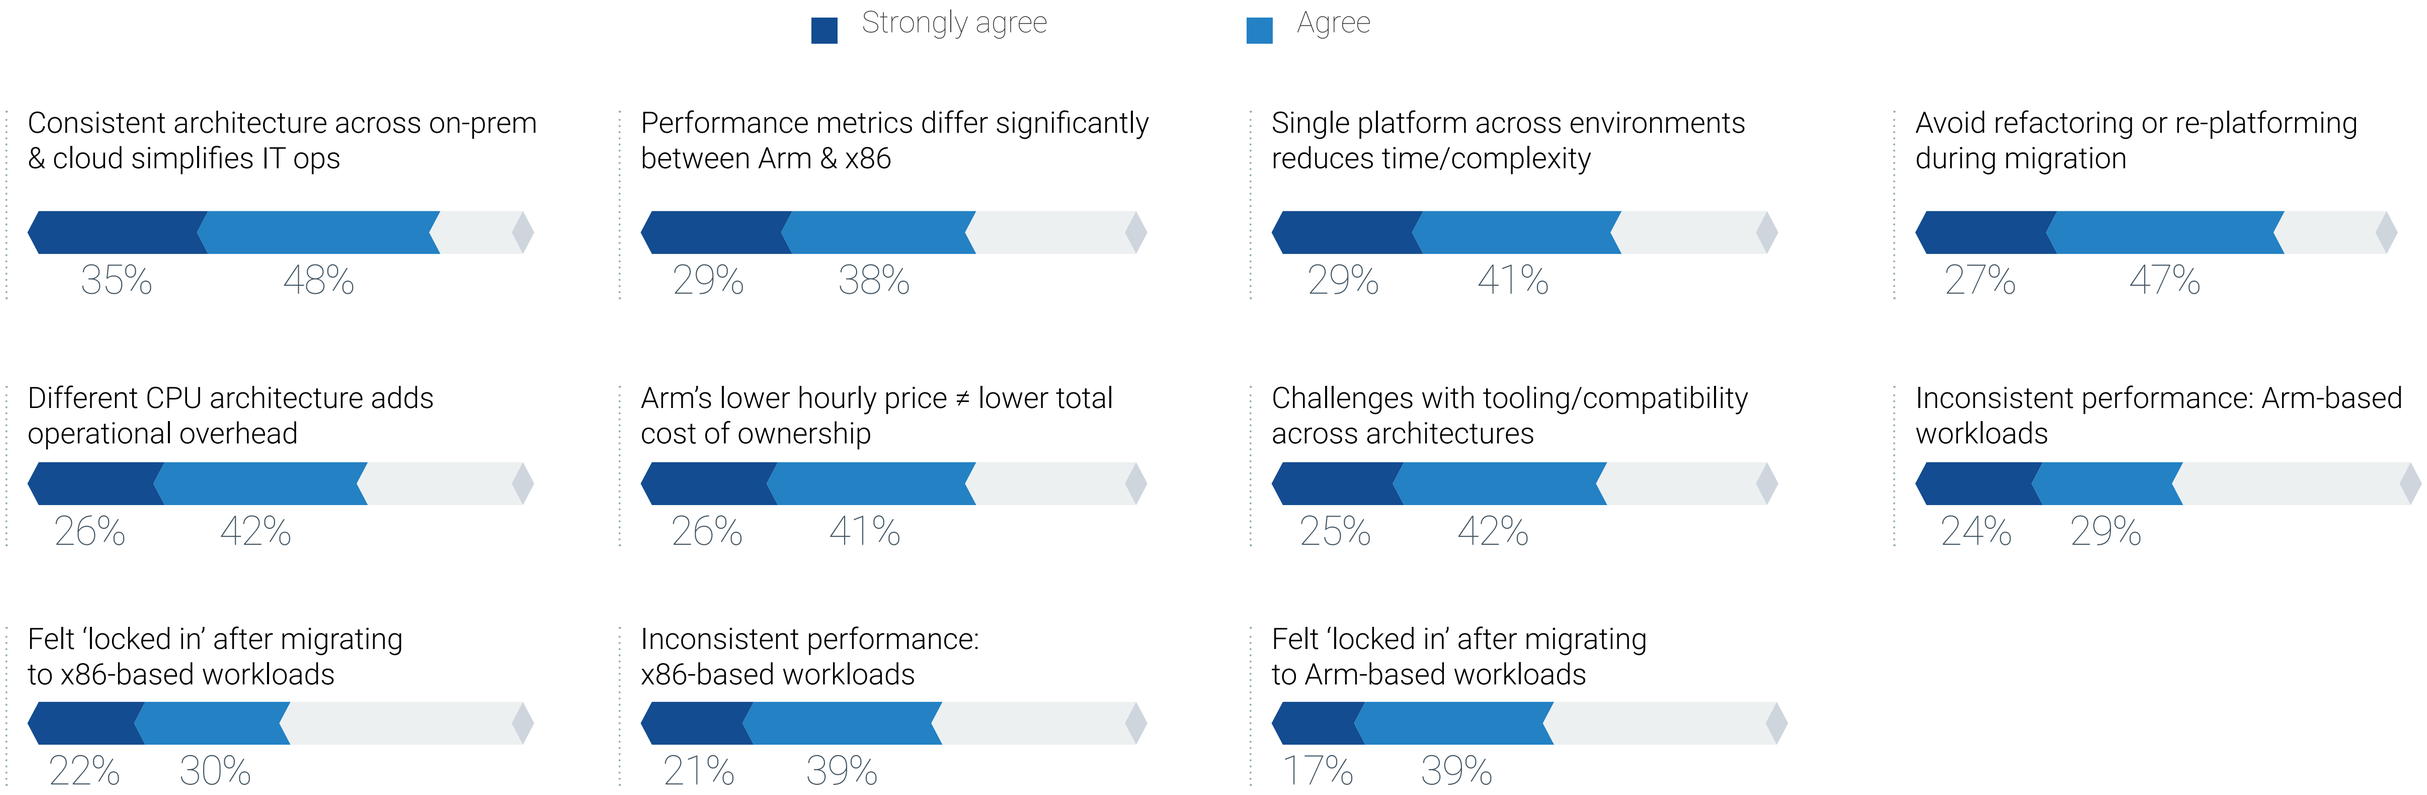




BACK TO CONTENTS

## Conclusion and Path Forward

Enterprises have learned that workload placement must be driven by performance, security, and operational consistency, not list prices or brand loyalty. The experience of Arm migrations underscores that hidden costs and ecosystem gaps can outweigh theoretical savings.

The path forward is about strategic consistency with selective diversity. Organizations should:

- Standardize where consistency reduces complexity (tooling, observability, security).
- Diversify only when specialized architectures clearly deliver differentiated value.
- Continuously validate TCO beyond hourly rates.

In summary, x86 remains the backbone of hybrid cloud flexibility, while organizations experiment pragmatically with alternatives where benefits are proven. The goal is not one architecture but a strategy that maximizes resilience, portability, and long-term business value.

AMD

ABOUT

For more than 55 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, LinkedIn, Facebook and X pages.

LEARN MORE

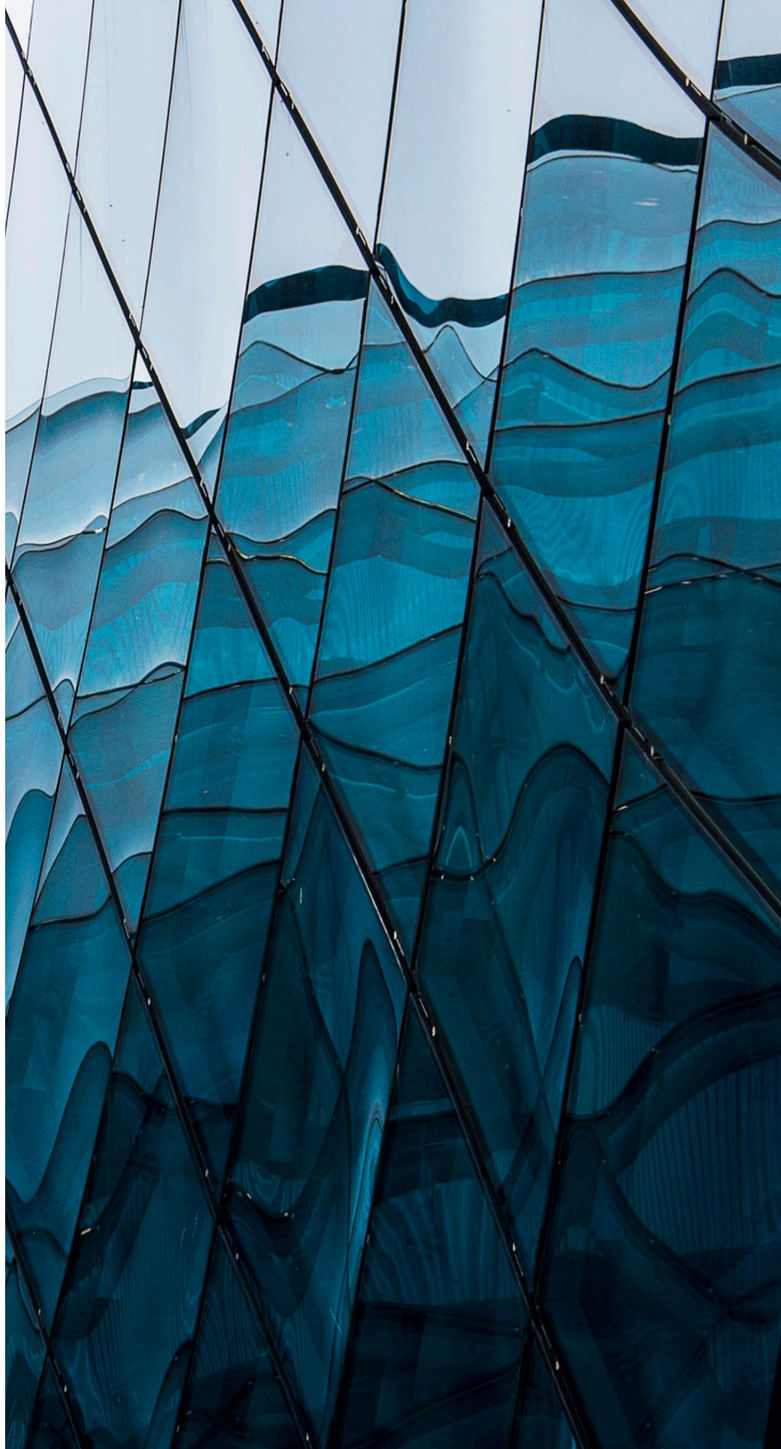

## RESEARCH METHODOLOGY AND DEMOGRAPHICS

To gather data for this report, AMD commissioned Enterprise Strategy Group to conduct a comprehensive online survey of IT professionals from private- and public-sector organizations in the United States. After filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on a number of criteria) for data integrity, we were left with a final total sample of 300 IT and data professionals.

### Respondents by Number of Employees

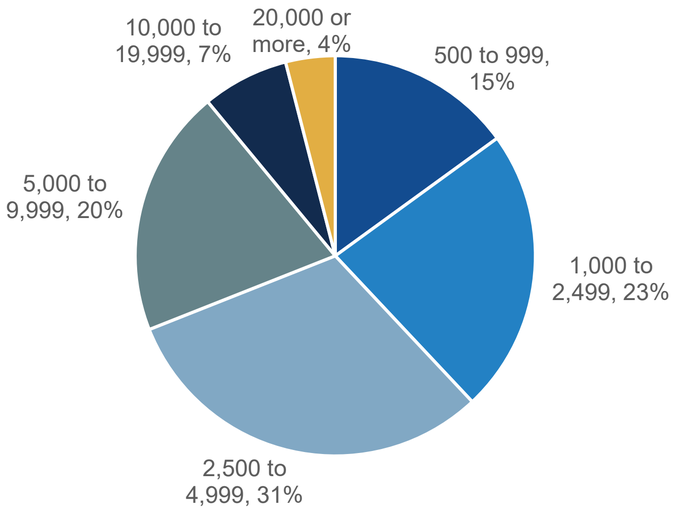


### Respondents by Job Title

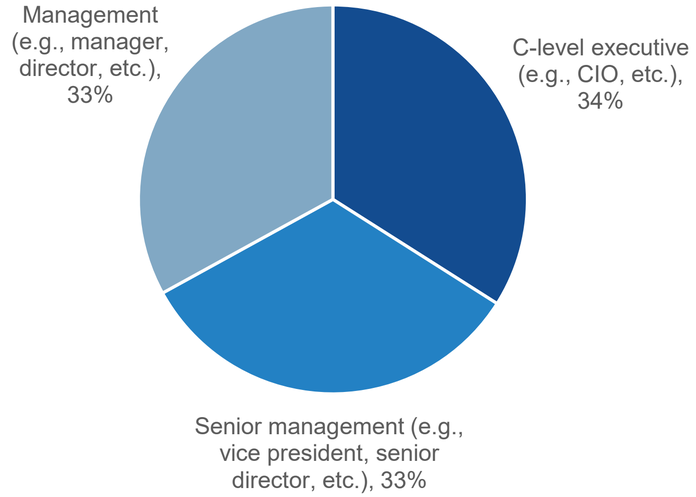


### Respondents by Industry

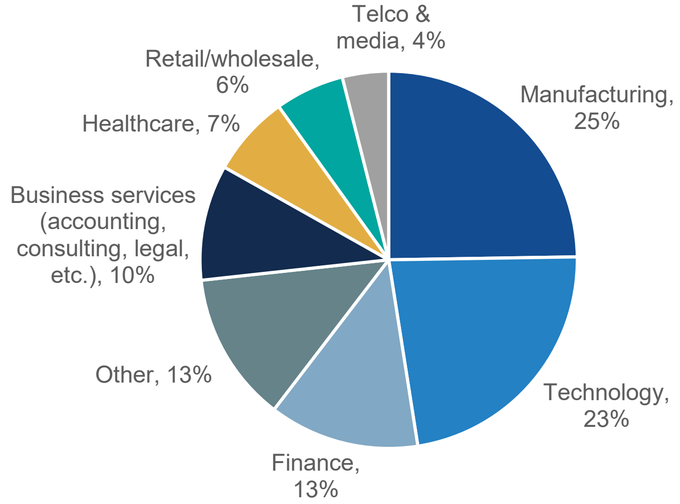




BACK TO CONTENTS


Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.