Executive Summary

Cloud Computing Study

Cloud growing pains: AI fuels adoption even as workload repatriation accelerates

FOUNDRY

Twenty years after Amazon Web Services put cloud computing on the map, the sector is grappling with maturity. The years of cloud hyper growth are now behind it, and customers are adopting a more nuanced attitude towards its promise and pitfalls.

Foundry's 13th annual Cloud Computing Study dramatizes the contrasts this transition creates. On the one hand, organizations are accelerating their shift to the cloud, focusing primarily on strategic workloads such as artificial intelligence and on process improvement and cost efficiency. Conversely, many are also repatriating workloads to on-premises infrastructure in areas where the cloud is seen as adding little value.

AI has become the primary catalyst for investment, architecture decisions and organizational change. More than one-third of the 527 survey respondents have added AI/machine learning engineer positions to their workforce as a result of their cloud investments. Nearly 30% have added an AI platform engineering role. This continues a trend documented in last year's survey, in which 36% of decision-makers said AI is the most sought-after cloud capability.

Among the other top findings of this year's survey:

- Hybrid cloud has emerged as the preferred deployment model, offering flexibility, control, and data sovereignty features that are increasingly important, particularly in regulated industries.
- Multi-cloud adoption is growing despite operational complexity. Cost savings is the top reason, cited by half of all cloud decision-makers.

- Neoclouds, a new breed of AI-first cloud provider that offers high-performance graphics processing units rented as a service, are gaining traction as AI infrastructure alternatives. Nearly three-quarters of respondents are currently using, testing or evaluating neocloud options.
- Cost control, security and skills shortages continue to challenge organizations in their cloud usage.

Respondents to this year's survey reported an average company size of 7,014 employees and an average IT budget of $156 million. Half the respondents came from North America, 31% from the Asia/Pacific (APAC) region and 20% from Europe/Middle East/Africa (EMEA). All are involved in the cloud computing purchase process and have at least one application or a portion of their infrastructure in the cloud.

Cloud adoption accelerates

There is little doubt that cloud infrastructure is the default platform for new IT initiatives. This year, 74% of cloud decision-makers said their organization accelerated cloud migration over the past 12 months, up from 70% in 2025 and 63% in 2024. The 74% who default to using the cloud for new capabilities is unchanged from 72% in 2022.

The cloud is being adopted for sound business reasons. This year, 73%

74%

of ITDMs say their organization has accelerated its migration to the cloud.

70% in 2025 → **63%** in 2024

of respondents agreed that cloud capabilities helped their organizations achieve increased and sustainable revenue over the past 12 months. That's up from 60% who said that in 2024.

Together, these figures underline a new reality in enterprise IT: The vast majority of organizations are now "cloud-first" companies. They no longer see cloud platforms as a way to trim expenses but as an enabler of speed and agility.

Although cloud budgets remain substantial, there are signs that organizations are growing more disciplined in their spending.

73%

report that cloud capabilities have driven sustainable revenue gains over the past 12 months.

71% in 2025 → **60%** in 2024

$156M

Average IT budget over the next year

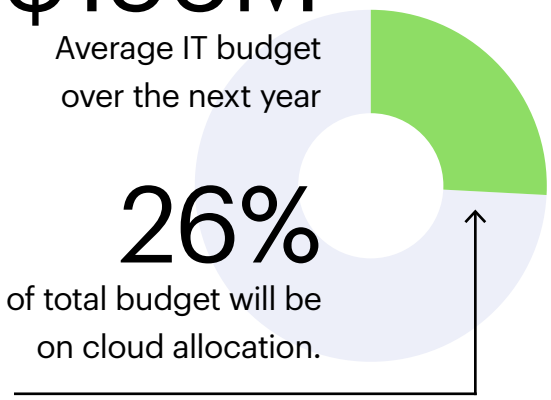

26%

of total budget will be on cloud allocation.

In this year's survey, decision-makers told us that 26% of their total IT budget will be allocated to cloud computing within the next year. That figure is down from 28% last year and 32% in 2022.

This doesn't indicate that organizations are abandoning cloud – far from it. Rather, they are entering a more mature phase where optimization matters as much as expansion. Evidence of this is the emergence of disciplines such as FinOps and cloud cost management tools, whose markets are growing by more than 15% annually, according to most estimates.

AI is the top growth area among cloud capabilities that organizations are seeking in 2026, identified by 56% of respondents.

This isn't news to cloud platform providers, who have been diligent about improving cost transparency and introducing their own tools to help customers manage budgets.

Deeper dive

Large companies (defined as those with more than 1,000 employees) plan to allocate 29.7% of their budgets to cloud computing over the next 12 months, compared to 21.3% for smaller companies (those with fewer than 1,000 employees). Large companies also plan to increase cloud spending the most.

Healthcare and financial services report the largest planned cloud budget allocation; the services industry reports the smallest.

APAC firms report both the largest planned increases and decreases in cloud spending over the next 12 months.

Eighty percent of North American and APAC respondents agree that "Our cloud strategy is helping to accelerate the adoption of AI," compared to 68% in the EMEA region.

IT leaders are much more positive about the cloud's bottom-line impact; 79% agreed, compared to 49% of non-IT executives.

AI now drives cloud strategy

Artificial intelligence has taken the world by storm, going from an academic discipline to a mainstream business imperative in less than four years. AI is now the top strategic cloud capability, cited by 56% of survey respondents, and well ahead of number two, security at 47%.

38%

are deploying or have fully deployed hybrid clouds, seeking cost savings, flexibility, and performance improvements.

A dominant 75% said they understand their cloud provider's AI roadmap, which is good news for the hyperscalers who have invested heavily in data center capacity for AI training and inference.

AI is reshaping workforce planning; 36% of organizations have added AI/machine learning engineer roles to optimize cloud investments (50% in high-tech companies), 27% have added AI platform engineers, and about 20% have opened requisitions for cloud architects, cloud software engineers and cloud developers.

The emergence of specialized AI infrastructure teams underscores the perceived value of the cloud as a target for AI investment. With Gartner expecting worldwide AI spending to grow 47% in 2026 alone, the ability of organizations to find and retain relevant talent will separate the leaders from the laggards.

Deeper dive

AI tops the list of strategic priorities at all organizations, regardless of size.

More buyers at larger companies said they understand their cloud provider's AI roadmap than at smaller ones (82% versus 69%).

Preferred deployment models

AI is also impacting deployment choices. Hybrid cloud is the AI deployment platform of choice for 38%, up from 29% in 2022. That's followed by private cloud at 28%, multicloud at 14%, and public cloud at 13%. Just 6% of respondents said they do not plan to deploy hybrid or multicloud architecture, down from 18% in 2022.

These results indicate that organizations are balancing AI innovation with growing governance and security requirements, making hybrid clouds an appealing compromise. They also reflect increased interest in sovereign clouds, which store and process data in compliance with local laws, privacy

According to IT leaders, the most appealing benefits of hybrid and multicloud architectures are:

1. Cost savings/ optimization **(50%)**
2. Greater platform and service flexibility **(48%)**
3. Performance/latency optimization **(47%)**

regulations, and security standards. International Data Corp. forecasts sovereign cloud spending to climb from $250 billion in 2027 to more than $400 billion before the end of the decade.

The top perceived benefits of hybrid and multicloud architectures are cost savings (cited by 50%), greater platform and service flexibility (48%), performance (47%), and improved business continuity (45%).

But there are challenges as well, led by cloud management and security costs (34%), operational complexity (33%), additional tooling costs (33%), and skills and staffing costs (26%).

Cloud provider policies create their own set of challenges to multicloud adoption. Security tops the list of concerns (cited by 46%), followed closely by operational complexity (43%), and data transfer costs (41%).

The costs and risks of moving large volumes of data securely between cloud providers are clearly a factor here. So are egress fees, which cloud providers charge customers to move data out of their cloud. By some estimates, egress fees account for up to 15% of cloud costs and are among the leading causes of budget overruns.

Complexity has long been a barrier to multicloud adoption. Cloud providers

Why are IT buyers moving to the cloud?

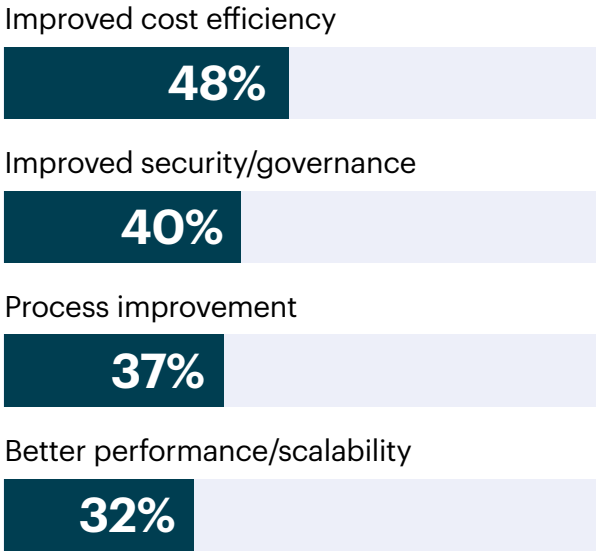

have little incentive to make it easier for customers to shift workloads to competitors. This is despite Foundry research finding that 84% would increase spending if multicloud environments were easier to operate. This dynamic may also explain why data-intensive operations like backup, database, data integration and AI workloads are the top candidates for repatriation to on-premises locations.

This year's study was the first to document the importance of industry-specific clouds, which all major platform providers now offer. The concept clearly resonates with customers, who ranked the importance of industry clouds at 7.8 on a 10-point scale, on average. The top benefits are perceived to be improved cost efficiency

(cited by 48%), improved security/governance (40%), process improvement (37%) and better performance/scalability (32%). With most research firms expecting double-digit annual growth for the next several years, industry clouds are an important new market factor.

Deeper dive

North American organizations are more likely than EMEA organizations to move AI/machine learning applications to on-premises infrastructure, by a margin of 42% to 26%.

Nearly 40% of new applications at larger companies will be cloud-native over the next 12 months, compared with 27% at smaller ones.

Fifty-five percent of large companies have seen performance improvements from adopting multicloud architecture compared to 34% of small companies.

Nearly four times as many IT leaders as business leaders would significantly increase spending with cloud providers if multicloud environments were easier to operate (35% versus 9%).

Cloud AI opportunities and obstacles

AI is transforming the cloud's value proposition from an infrastructure modernization platform to an innovation enabler. This requires new architectures, skills and governance models.

Cloud providers are poised to get the lion's share of the AI business from enterprise customers for some good reasons. The perceived strengths they bring to the market include their wide variety of model platforms and tools (cited by 47%), security (46%), AI-optimized infrastructure (36%), and cost optimization (29%).

But this isn't a winner-takes-all market. In fact, just 13% of respondents to the Cloud Computing Study said they plan to run AI workloads primarily in a public cloud. Hybrid cloud is the preferred platform by a wide margin (38%), followed by private cloud (28%) and multi-cloud (14%). These results probably reflect the importance customers place on keeping tight control over both data and cloud costs.

Customers' top reservations about going with a fully public cloud AI option are data security/privacy/compliance (cited by 50%), integration with existing systems (39%), lack of AI expertise (35%), and managing AI across multiple clouds (32%). Just 15% cited vendor lock-in as a concern (18% in EMEA), indicating that customers have become more sophisticated about bridging public and private cloud options and managing multiple vendors.

Experience with neocloud providers

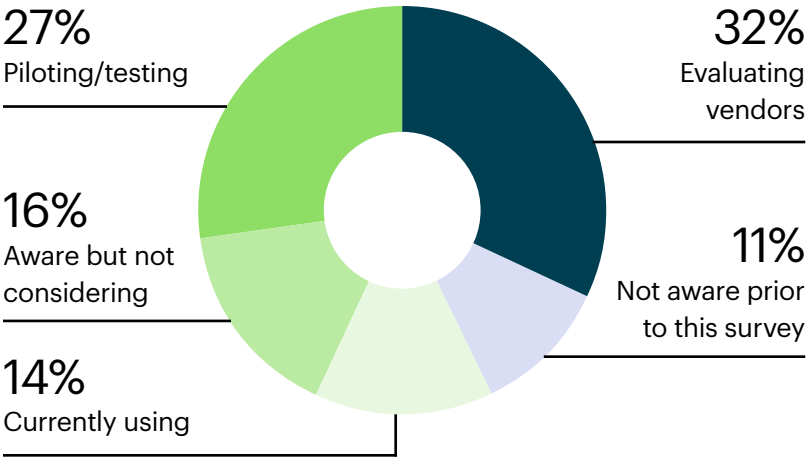

For those considering the cloud for the majority of their AI training and inference, an intriguing new option has emerged: neoclouds. These are specialized providers focused exclusively on delivering graphics processing units as a service on high-performance infrastructure tailored for AI workloads.

Unlike traditional hyperscalers that sell thousands of general-purpose services, neoclouds are designed from the ground up for AI. They keep their infrastructure catalogs lean and generally offer GPU rental rates up to 80% lower than traditional hyperscalers. Examples of neocloud vendors include CoreWeave, Lambda, Vultr and Nscale.

Although the neocloud market is barely three years old, it has rapidly captured customer attention. One in seven survey respondents is currently using a neocloud provider, 27% are piloting, and 32% are evaluating. Only 11% reported being unaware of the neocloud option.

The top use cases are AI experimentation and model development (55%), AI-specific data processing (47%), AI training (46%), high-performance workloads (37%), and production AI applications (36%). These findings indicate that neoclouds cover a wide range of needs.

The growing population of neocloud providers underscores the value of need driving innovation. With GPUs in short supply and selling for up to $40,000 per unit, neocloud providers have emerged to fulfill short-term demand and provide specialized AI infrastructure that delivers significant value. AI is thus creating new cloud market segments.

Deeper dive

Just 8% of EMEA buyers were unaware of neoclouds before the survey, compared with 15% in North America.

Forty-eight percent of APAC respondents are using or piloting neoclouds, compared with 35% of North American respondents.

Cloud repatriation trend growing

The most interesting finding of this survey is that cloud migration and cloud repatriation are both on the rise. That isn't a contradiction but rather an indication that organizations are taking a more sophisticated approach to workload placement than a few years ago. The cloud conversation has evolved from "move everything" to "put every workload in its place."

Respondents report that, on average, 22% of workloads have already been repatriated, and another 24% are planned for repatriation. One-third of responding organizations have moved at least one-quarter of their workloads back on-premises, down from 37% in 2024.

The workloads most likely to move back in-house are storage, backup and file services (75% have moved or plan to move), databases (74%), AI/ML workloads (72%), data integration (72%), and websites and web applications (69%). It's noteworthy that the 2022 study found the top business driver for moving applications to cloud computing was disaster recovery. Changing regulatory and privacy considerations have now made it the top reason to move back on-premises.

66%

of ITDMs say their organizations have moved or plan to move applications or workloads out of the cloud because of:

- Security concerns **(60%)**
- Budget/cost control **(42%)**
- Performance/reliability issues **(42%)**
- Data sovereignty **(40%)**
- Compliance concerns **(31%)**

The principal reasons decision-makers cite for repatriation are security concerns (60%), budget and cost control (42%), performance and reliability (42%), data sovereignty (40%), and compliance (31%).

Repatriation isn't evidence of cloud computing's failure. Rather, it shows that organizations now understand the pros and cons of cloud platforms and their workloads well enough to place each where it delivers the best balance of cost, performance, security, and AI readiness.

There are several likely reasons repatriation is trending.

AI economics are changing the math. Running large-scale AI inference can

become extremely expensive in public cloud environments. For organizations with predictable AI workloads, dedicated infrastructure is often cheaper over the long term than paying cloud consumption fees.

Data egress charges are a sticking point. As noted earlier, the cost of moving data out of the cloud can be high, particularly in large-scale scenarios such as backup and AI model training. These functions may be handled more cost-effectively on-site.

Security and sovereignty concerns have intensified. Although cloud infrastructure is considered highly secure, concerns over disclosure of sensitive AI training data, sovereign AI requirements, and data residency regulations worry customers.

Expected cost savings never materialized. Many organizations have discovered that the cloud can reduce capital costs but increase operational expenditures. Cost control was the second-most-cited reason for repatriation, at 42%.

Hybrid cloud has made repatriation easier. Ten years ago, moving a workload back from the cloud was viewed as a failure. Today, hybrid architectures are mainstream. Shifting workloads dynamically between environments is now considered a best practice rather than a strategic reversal.

Enterprises have become better cloud buyers. Perhaps the most important factor driving repatriation is buyers' greater savvy about their cloud options. "Lift-and-shift" migrations and "cloud first" policies have mostly gone by the boards. Organizations are now scrutinizing which workloads truly benefit from cloud features like elasticity and which can operate under more stable utilization patterns.

The future is not all-cloud or on-premises. It is a carefully optimized mix of environments. As cloud enters young adulthood, it is encountering buyers who are savvier about their options and quicker to pull the plug if perceived value doesn't match reality.

Deeper dive

APAC buyers are more likely to increase spending with cloud providers who make multi-cloud operations easier than their peers in EMEA by a 35%-to-20% margin.

The same is true of larger companies versus smaller ones by a 36%-to-22% margin.

Larger companies are somewhat more likely to repatriate applications than smaller companies.

Obstacles remain

The reasons enterprises don't deploy cloud services have changed little over time.

94%

of ITDMs have experienced some obstacles with their cloud adoption over the past 12 months.

Why? 47% of IT leaders cite budget challenges and cost

Cost, security and skills were the top three barriers cited this year, the same as in 2022.

Early perceptions that cloud platforms are cheaper than on-premises infrastructure have largely evaporated. The reality is that the long-term cost of cloud infrastructure often exceeds that of captive data centers. The value of cloud is more strategic: scale, business agility, and tooling that promotes innovation.

Nearly half (47%) of respondents cited high costs as a barrier to adopting cloud computing initiatives. Top concerns are data transfer charges (31%), pricing changes (30%), networking fees (29%), unpredictable costs (29%), and compliance and regulatory overhead (25%).

Lack of visibility into cloud usage is also a concern. Unsurprisingly, 74% of customers have or will create dedicated cloud cost management roles. Beyond costs, security and compliance concerns persist. Top challenges reported by survey respondents include protecting cloud data (63%), securing cloud resources (55%), and regulatory compliance (41%). While cloud platforms are considered highly secure, the industry operates under a shared responsibility model that invests significant responsibility in customers to control access and secure identities, applications and data. Security concerns may reflect more customers' reservations about their own resilience than about their providers'.

Cloud providers have recently positioned their platforms as enablers of modernization, but the survey results indicate that customers are still struggling with this task. Integration with on-premises systems is cited as a challenge by 56% of buyers. That's followed closely

Cloud adoption has slowed because of security and compliance-related obstacles such as:

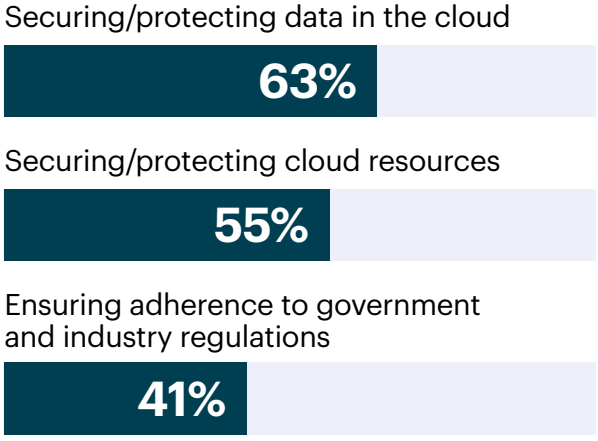

Top new roles to support cloud investments in 2026

- AI/Machine Learning Engineer
- AI Platform Engineer

by legacy modernization complexity (50%) and data portability (49%).

Skills shortages also continue to hold back cloud adoption. Over half (54%) of buyers noted difficulty finding cloud security expertise, followed closely by shortages of people with skills in cloud management (50%), cloud architecture (43%), and cloud development (43%). These figures are sharply up from the 2022 study, in which 33% reported a shortage of cloud management skills/expertise, and 30% struggled to find cloud development skills.

The key takeaway is that cloud computing continues to suffer from many of the same adoption obstacles today as it did more than a decade ago. In fact, 94% of decision-makers said they've experienced obstacles to their cloud adoption over the past 12 months, virtually unchanged from 96% in 2022. Technology continues to advance faster than organizational capabilities, creating persistent operational and governance challenges. AI is only going to increase those pressures.

Deeper dive

Cost is a bigger impediment to cloud adoption at smaller companies than at larger ones (52% to 41%).

Integration with on-premises systems is a barrier for 62% of larger companies, compared to 47% of small ones.

Large companies are more likely than smaller companies to create roles dedicated to managing cloud costs, at 82% versus 66%.

Defining themes

Three major themes dominated the results of this year's Cloud Computing Study.

AI is redefining cloud strategy

AI is now shaping enterprise cloud strategies. The pace of this change has been dramatic. AI wasn't among the top five cloud growth areas in the 2022 study, but by 2024, it had become the number-one planned cloud investment area. This year, it is the top strategic cloud priority for 56% of respondents. Nearly four in five (78%) said cloud is helping to accelerate their adoption of AI.

In 2026, the top new roles being added to support cloud investments are AI/machine learning engineer and AI platform engineer. Neither was on the top five list in 2024.

Cloud has shifted from an infrastructure modernization story to an AI enablement story. Given the massive investments hyperscalers and neocloud providers are making in data center capacity, this dynamic is unlikely to change anytime soon.

Flexibility is more important than standardization

Hybrid cloud, multi-cloud, industry cloud, neocloud and repatriation are all workload-centric strategies, meaning that their choice is driven by a complex set of factors spanning cost, performance, security, and control. The vast majority of organizations are now using more than one cloud provider, indicating that deployment decisions are becoming more sophisticated and nuanced.

Economics and governance matter more than ever

At the same time, cost control, security and operational simplicity have become as important as innovation in driving decisions.

Customers don't perceive significant differences between the three major hyperscalers. Asked which public providers lead in AI capabilities, 58% cited Microsoft, 53% Google and 52% AWS, a virtual tie.

With customers less likely than ever to put all their chips on a single vendor, cloud providers will be challenged to demonstrate superior capabilities in cost management, data protection, and enabling responsible AI development.

The organizations gaining the greatest value from cloud computing are no longer simply those pursuing a cloud-first strategy. They are pursuing an AI-first, workload-optimized plan that uses cloud, hybrid infrastructure and on-premises resources, with each delivering the greatest business value.

Cloud computing has grown up. Diversification, specialization, and AI innovation will drive the industry's next stage of evolution.

About the survey

The 2026 Cloud Computing Study was conducted to measure cloud computing trends among technology decision-makers including: usage and plans across various cloud service and deployment models, investments, business drivers and impact on business strategy and plans. The study was fielded March–April of 2026 and is based off the responses of 527 global IT decision-makers that are involved in the purchase process for cloud computing, and their organization has, or plans to have, at least one application, or a portion of their infrastructure, in the cloud.

Examining the marketplace

Research is an invaluable way for marketers to better understand customers and prospects, with the goal of building quality connections. At Foundry this is one way we are focused on building bridges between tech buyers and sellers. Our first-party relationships with the most important tech buyers and influencers around the world, allows us to apply value across our customers marketing stack. Our research portfolio explores our audiences' perspectives and challenges around specific technologies—from analytics and cloud, to IoT and security—and examines the changing roles within the IT purchase process, arming tech marketers with the information they need to identify opportunities.

To see what research is available, visit foundryco.com/tools-for-marketers. For a presentation of full results from any of these studies, contact your Foundry sales executive or go to foundryco.com/contact-us.

Buying process

Each year we take a deep dive into the enterprise IT purchase process to learn more about who is involved and who influences decision-making, what sources purchasers rely on to keep up to date with technology—and throughout the purchase process—and how they want to engage with the vendors they are working with. Visit foundryco.com/customerjourney for more information.

Buying process studies

- Customer Engagement
- Role and Influence of the Technology Decision-Maker

Technology insights

Each year we explore the technologies that are top of mind among our audiences to understand the business challenges, drivers, and adoption within the enterprise. These research studies are designed to help IT marketers understand what their customers are focused on and where the market is moving.

Role and priority studies

- CIO Tech Poll: Tech Priorities
- State of the CIO

Technology-specific studies

- AI Priorities
- Cloud Computing
- Security Priorities
- Partner Marketing

Stay in touch with us

Email: Sign up for Foundry's newsletters and receive media and marketing trends as well as our proprietary research, product and event information direct to your inbox. Go to foundryco.com/newsletter.

Social: For research, services and events announcements, visit us on LinkedIn.

Find it all on foundryco.com.

About Foundry

Foundry's vision is to make the world a better place by enabling the right use of technology, because we believe that the right use of technology can be a powerful force for good.

Foundry is a trusted and dependable editorial voice, creating quality content to generate knowledge, engagement and deep relationships with our community of the most influential technology and security decision-makers. Our premium media brands, including CIO®, Computerworld®, CSO®, InfoWorld®, Macworld®, Network World®, PCWorld® and Tech Hive®, engage a quality audience of the most powerful technology buyers with essential guidance on the evolving technology landscape.

Our trusted brands inform our global data intelligence platform to identify and activate purchasing intent, powering our clients' success. Our marketing services create custom content with marketing impact across video, mobile, social and digital. We simplify complex campaigns that fulfill marketers' global ambitions seamlessly, with consistency that delivers quality results and wins awards. Additional information about Foundry is available at foundryco.com.