Enterprise Strategy Group | Getting to the bigger truth.™

RESEARCH HIGHLIGHTS

The Four Infrastructure Essentials for AI/ML Data Pipeline and Data Lake Environments

Scott Sinclair, Senior Analyst

SEPTEMBER 2021

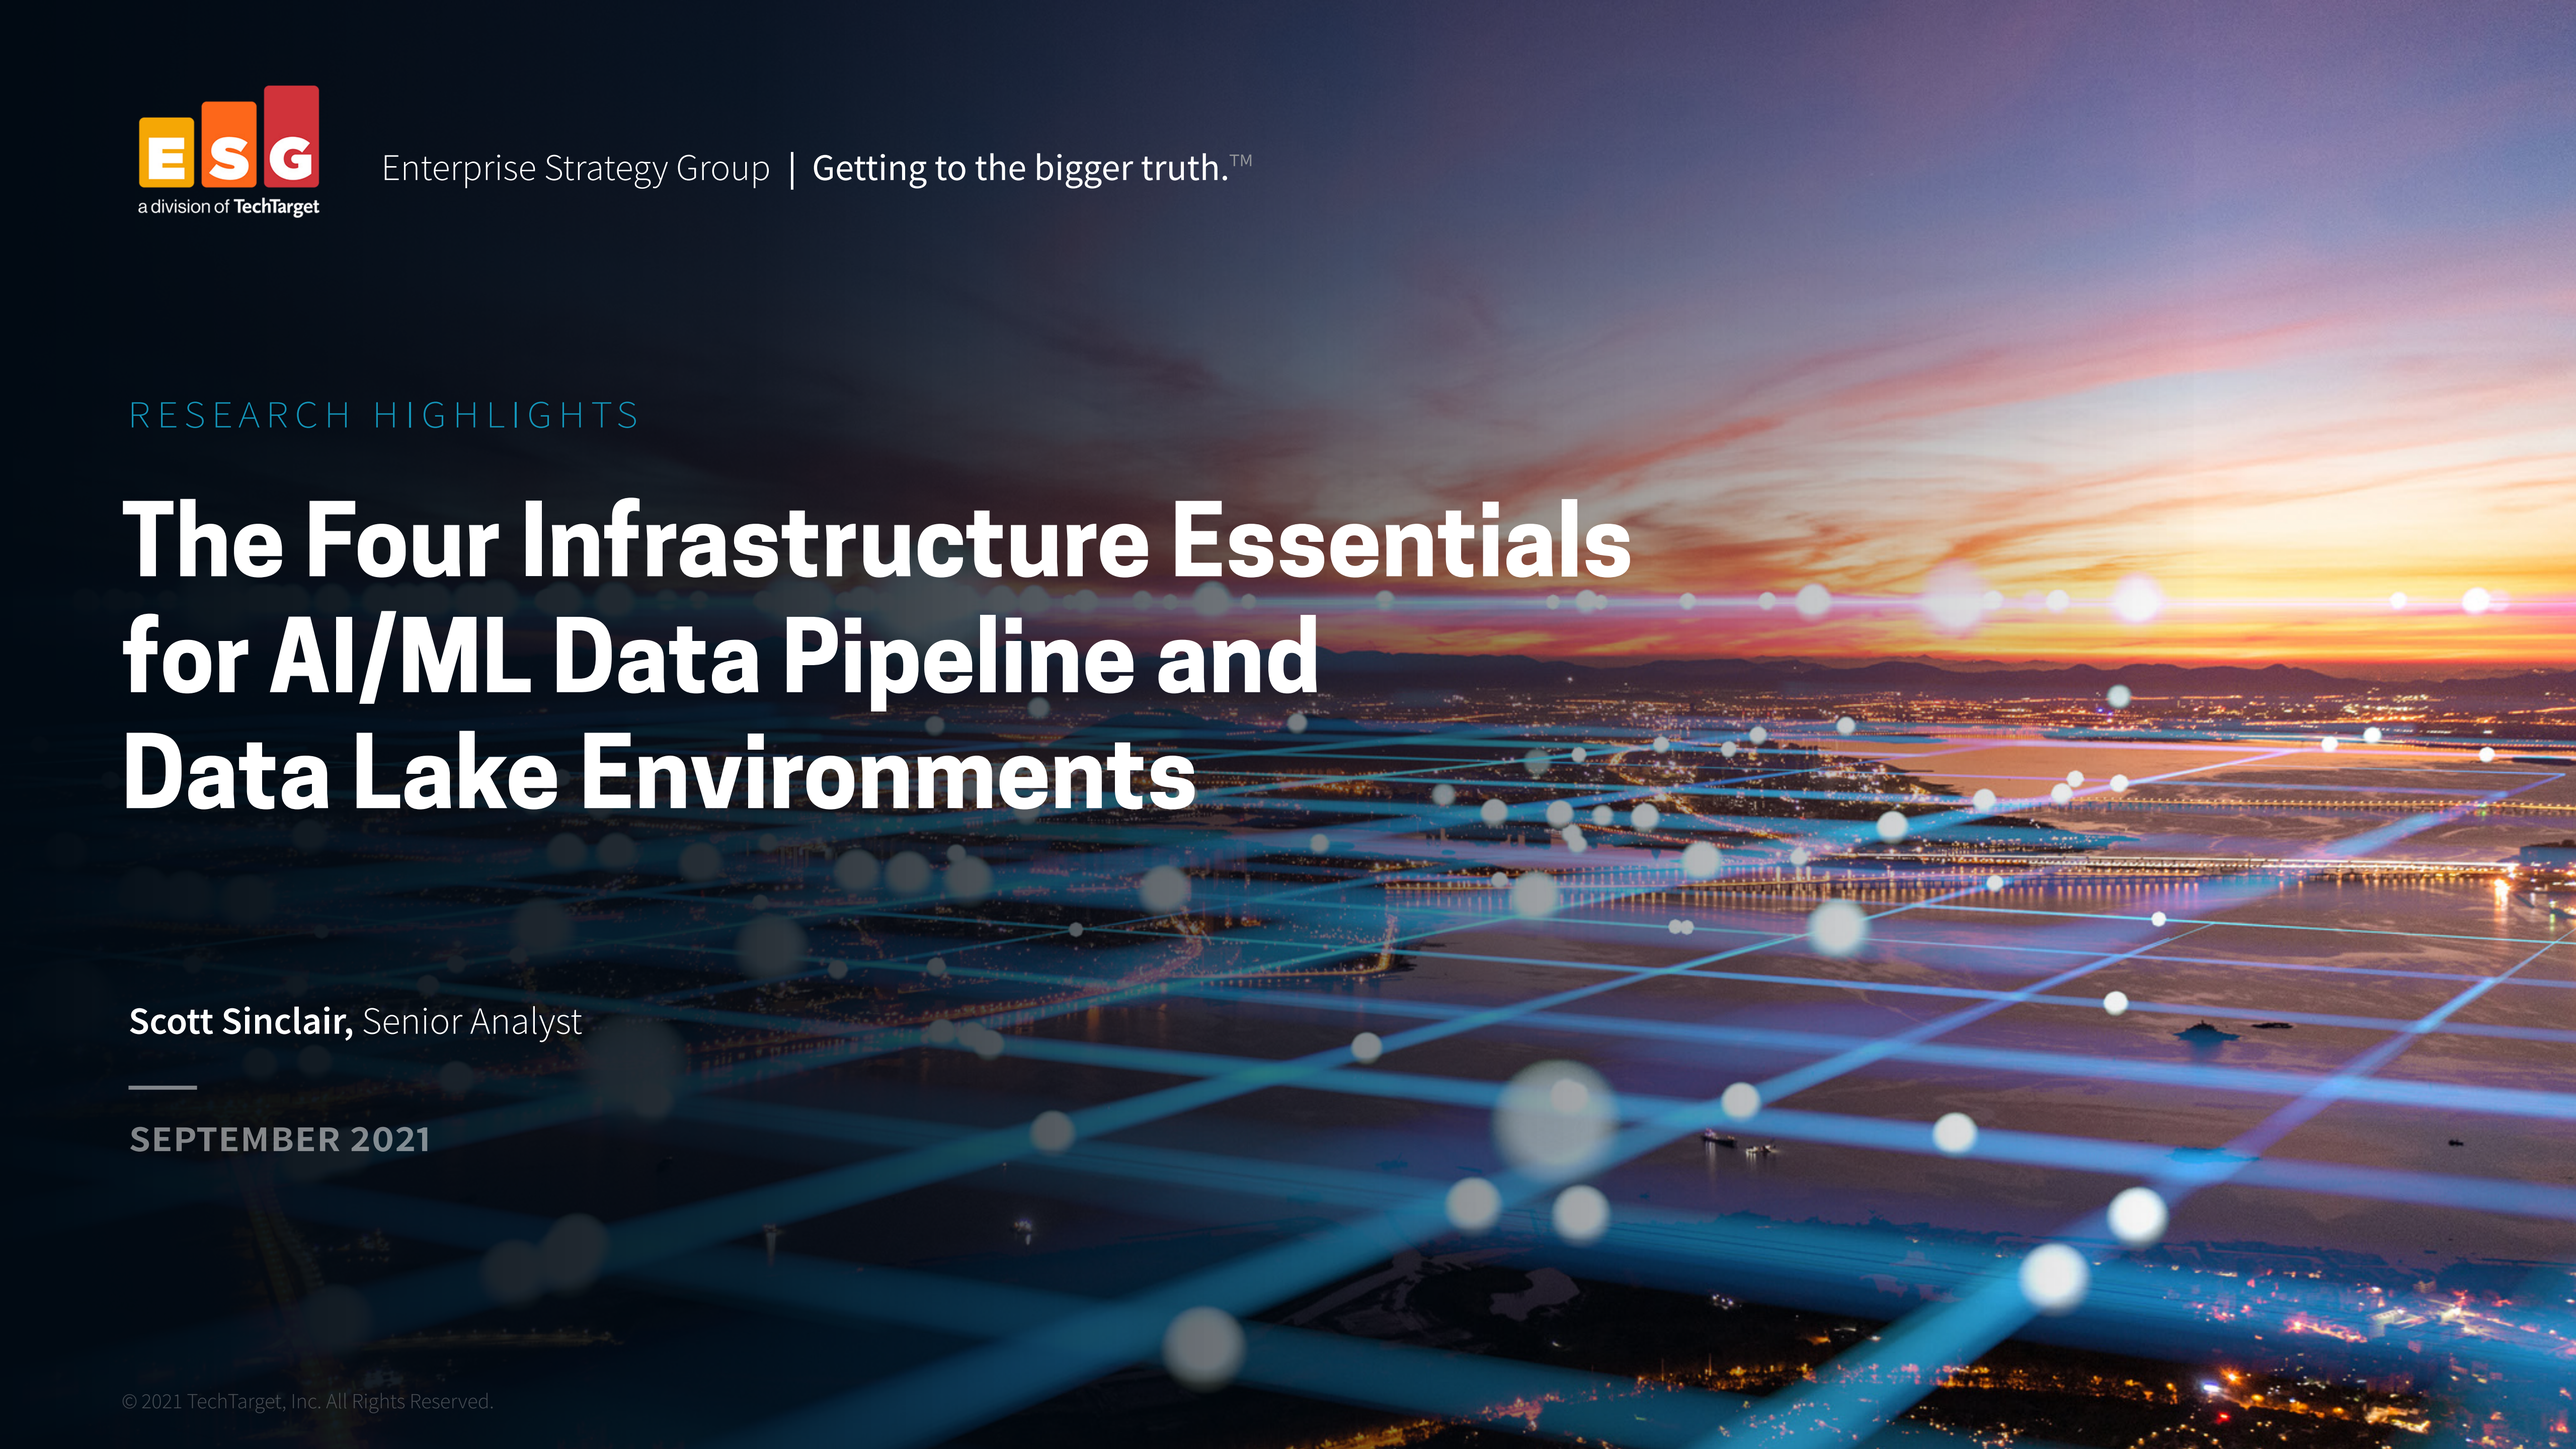

Contents

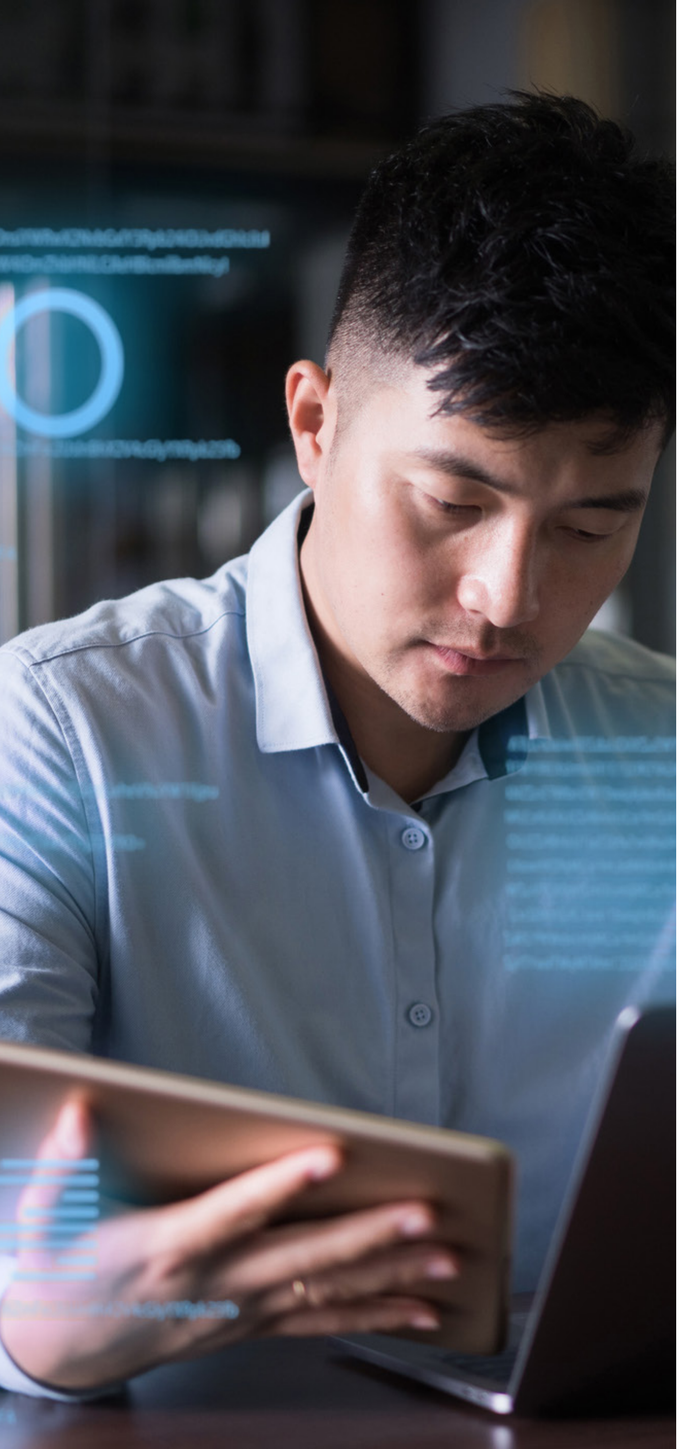

Overview

Artificial intelligence and machine learning (AI/ML) initiatives continue to play a larger role in transforming modern business operations and customer engagement. When organizations achieve success with AI/ML initiatives, they often quickly ramp investment, expanding projects to target multiple objectives. The resulting data pipeline infrastructure often scales quickly, resulting in a massive and often disaggregated infrastructure environment. As a result, ensuring the organization has the right data pipeline infrastructure environment is essential for success.

To gain insight into the state of production-level AI/ML workloads and their supporting infrastructure environments, ESG surveyed 325 IT professionals at organizations in North America (U.S. and Canada) who are familiar and involved with evaluating, purchasing, and/or managing IT infrastructure associated with AI initiatives for their organization. This research was intended to understand details about the breadth of IT infrastructure that comprises modern data pipeline environments and to provide insights into the priorities and business objectives organizations have for their AI/ML initiatives.

Note: Totals in figures and tables throughout this eBook may not add up to 100% due to rounding.

ESG RESEARCH FINDINGS:

Four Infrastructure Essentials for Data Pipeline and Data Lake Environments

- Maximizing Infrastructure Performance and Utilization
- Data Management and Security/ Governance
- Hybrid/ Multi-cloud Capability
- Data Durability/ High Availability

Back to Contents

Businesses Often Find Success Quickly with AI/ML

Businesses Approach AI to Accomplish Multiple Objectives

Organizations on average identified 3.8 objectives as the most important objectives for their investment in AI/ML.

Regardless of the objective, organizations were able to achieve value quickly--97% in the first year and 82% in the first 6 months.

82%
saw value from AI/ML in 6 months or less.

Top Objectives for AI/ML Investments

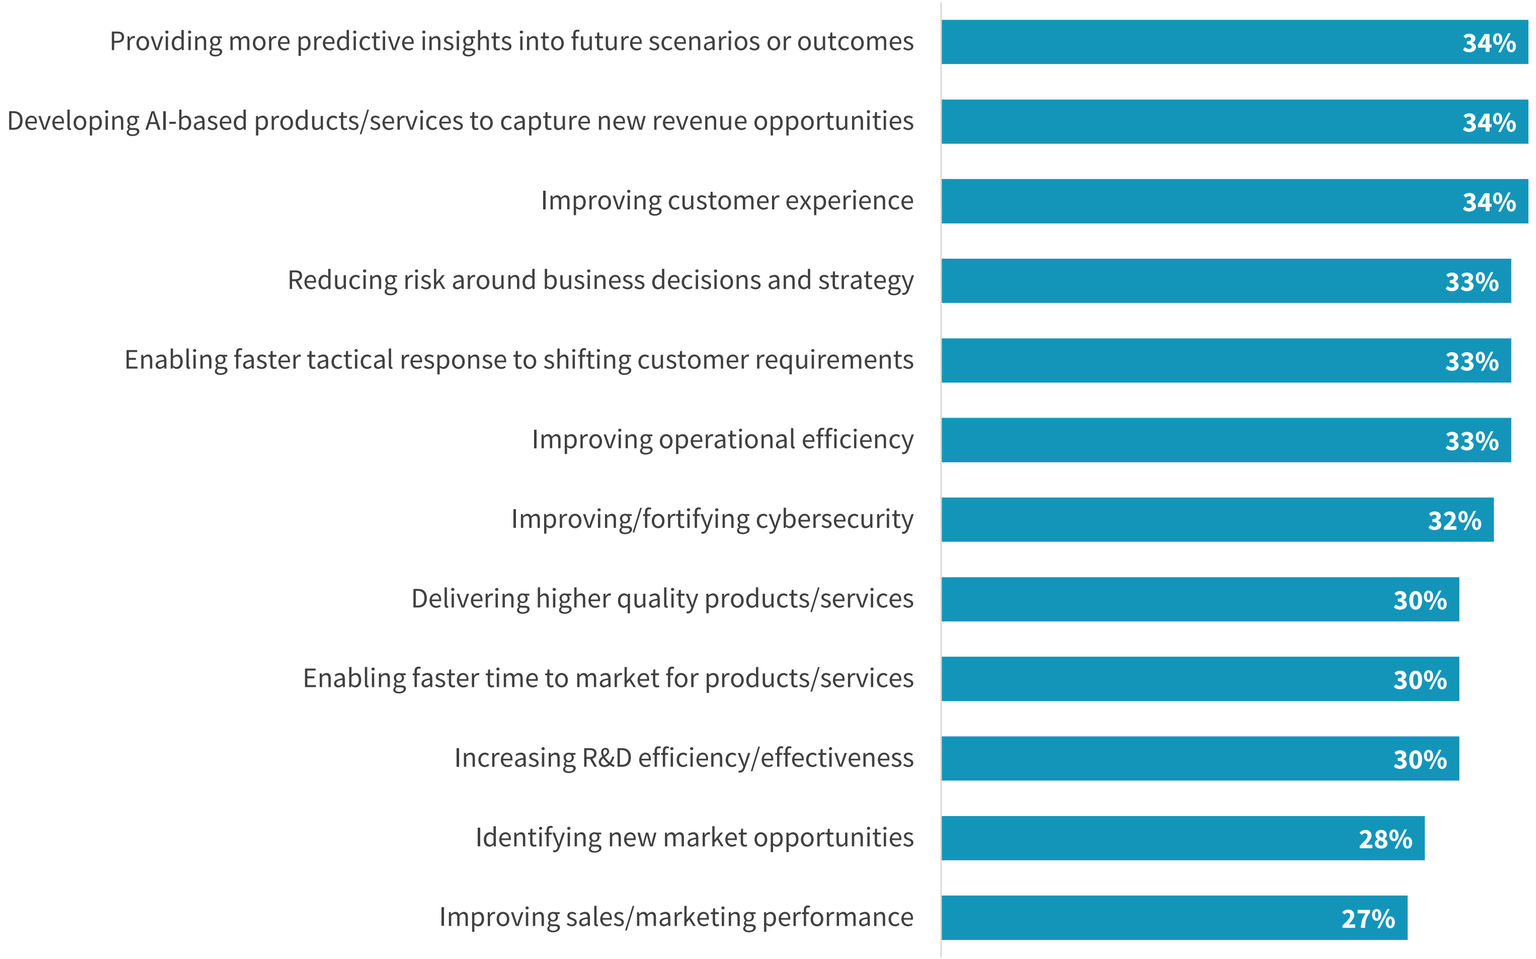

Back to Contents

AI/ML Projects Regularly Met or Exceed Expectations

ESG asked IT decision makers to identify whether their AI/ML initiatives met their expectations for each of their organization's predefined AI/ML objectives. The following statistics reveal for each objective the percentage of organizations that identified that their AI/ML initiative met or exceeded their expectations in that area.

Percentage of organizations where AI/ML met or exceeded their expectations by AI/ML Objective

93%
Higher quality products or services

90%
Developed AI-based products and services to capture new revenue opportunities

86%
Reduced risk around business decision and strategy

94%
Improved/fortified cybersecurity

93%
Improved sales/marketing performance

91%
Enabled faster time to market for products and services

98%
Identified new market opportunities

88%
Increased R&D efficiency/effectiveness

98%
Improved operational efficiency

90%
Enabled faster tactical response to shifting customer requirements

93%
Improved customer experience

92%
Provided more predictive insights into future scenarios or outcomes

Back to Contents

Data Pipeline and Data Lake Infrastructure Decisions that Help Fuel AI/ML Success

Data Pipeline for Production AI/ML Environments Span the Edge, the Data Center, and the Cloud

Locations in Use by Data Pipeline Phase

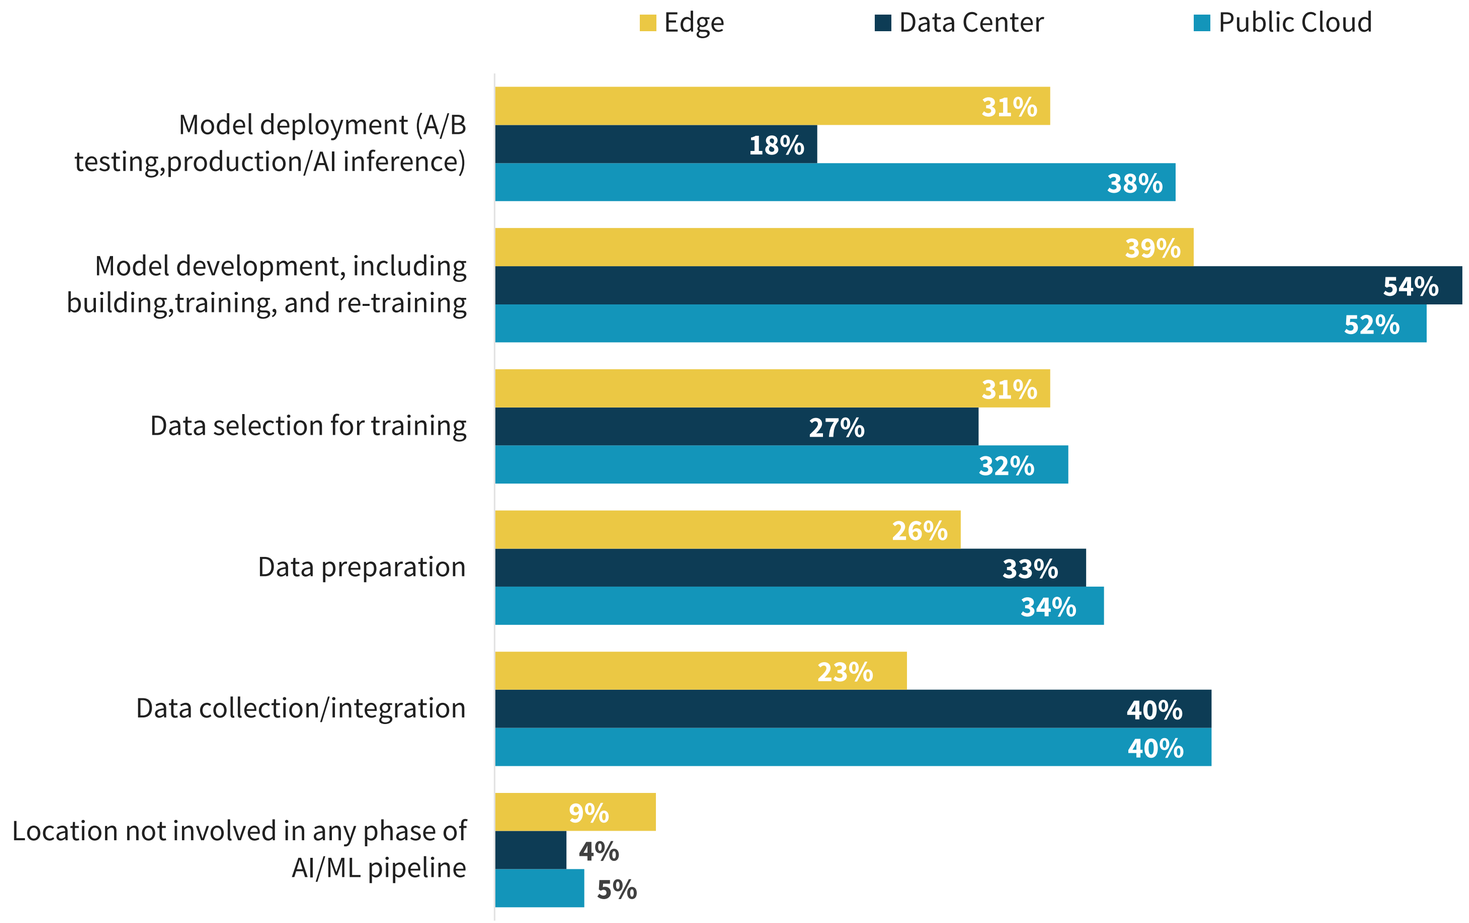

> The distributed nature of data pipeline environments **fuels the need for data and application portability across multi-cloud, data center, and edge locations.”**
>
> - Scott Sinclair, Senior Analyst

Back to Contents

Data Lakes are Increasingly Foundational Data Layer to AI Success

Location of Data Lake Environments

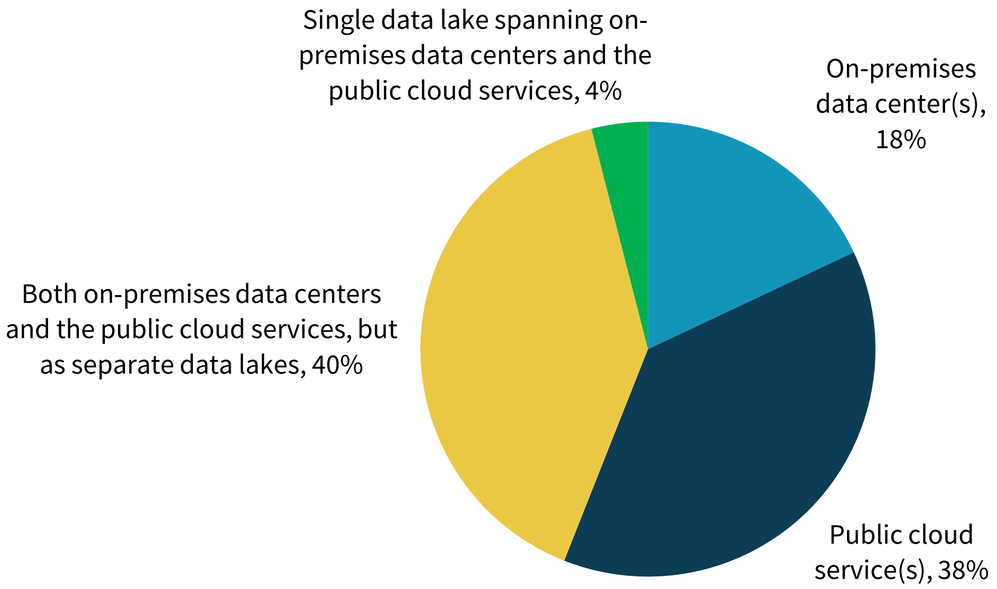

Amount of Flash Storage Deployed to Support Data Lake Environments

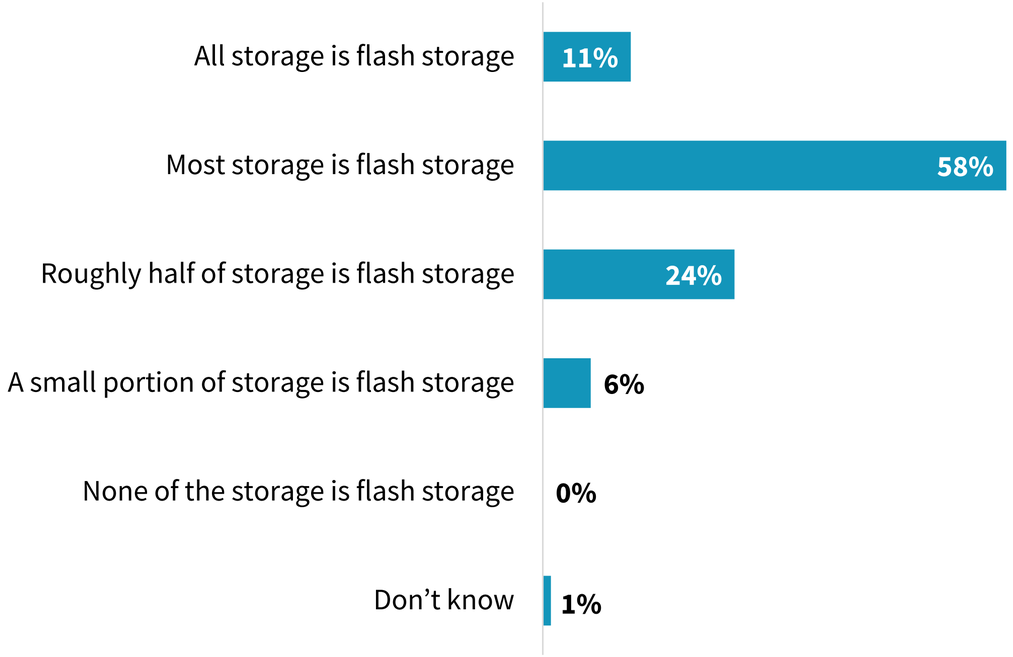

Distributed Nature of Data Lakes Fuels Need for Hybrid/Multi-Cloud Capability

69% of Data Lake Environments for AI/ML are All or Mostly Flash, Highlighting the Pervasive Need for Performance

Back to Contents

Defining the Four Infrastructure Essentials for Data Pipeline and Data Lake Environments

COMMON WEAK LINKS AND INFRASTRUCTURE CONSIDERATIONS

26% identified resource sharing and 25% identified GPU processing as among the weakest links in their AI/ML environnent.

And 17% identified maximizing hardware/ infrastructure production utilization and 16% identified lowest possible latency as among the most important considerations for production infrastructure solutions that support AI/ML model development.

Over half of organizations with production AI/ML identified leveraging the cloud (52%) and data center (54%) environments for AI/ML model development.

And 19% identified hybrid/multi-cloud capability as one of the most important considerations for production infrastructure solutions that support AI/ML data preparation, and 20% for production infrastructure solutions that support AI/ML model deployment.

100% of the participants in this study identified that at least some of the data in their data pipeline was sensitive data.

And 19% identified data security/governance as one of the most important considerations for production infrastructure solutions that support AI/ML model development and 18% for production infrastructure solutions that support AI/ML model deployment.

22% identified data storage as one of the weakest links in their AI/ML environnent.

And 15% identified data durability/high availability as one of the most important considerations for production infrastructure solutions that support AI/ML model deployment.

AS A RESULT, ORGANIZATIONS MUST PRIORITIZE

Maximizing Infrastructure Performance and Utilization

Hybrid / Multi-cloud Capability

Data Management and Security/Governance

Data Durability / High Availability

Back to Contents

Four Infrastructure Essentials for Data Pipeline and Data Lake Environments

Designing the right infrastructure environment helps fuel continued success with AI/ML initiatives. Multiple infrastructure capabilities offer benefits but among production environments, hybrid/multi-cloud capability, maximizing hardware/infrastructure utilization, data durability/high availability, and data security/governance are often among the top considerations.

Maximizing Infrastructure Performance and Utilization

AI/ML environments can quickly scale in both performance and capacity. Maximizing utilization is key to keeping costs under control. This requirement extends to processing components, such as GPUs as well.

Hybrid/Multi-cloud Capability

Data pipeline and data lake environments typically require massive scale and often span multiple locations. Infrastructure that can simplify the management of multiple sites and locations offers value.

Data Management and Security/Governance

100% of the participants in this study identified that at least some of the data in their data pipeline was sensitive data. Security and governance must be top priorities for any AI/ML environment.

Data Durability/High Availability

The data in this environment delivers business value and, for organizations with AI/ML in production, these workloads are often viewed as business–crucial. Data must be resilient and always available.

Back to Contents

The Dell Technologies Approach to AI Environments

Dell Technologies Delivers Scale and Portability for AI Environments

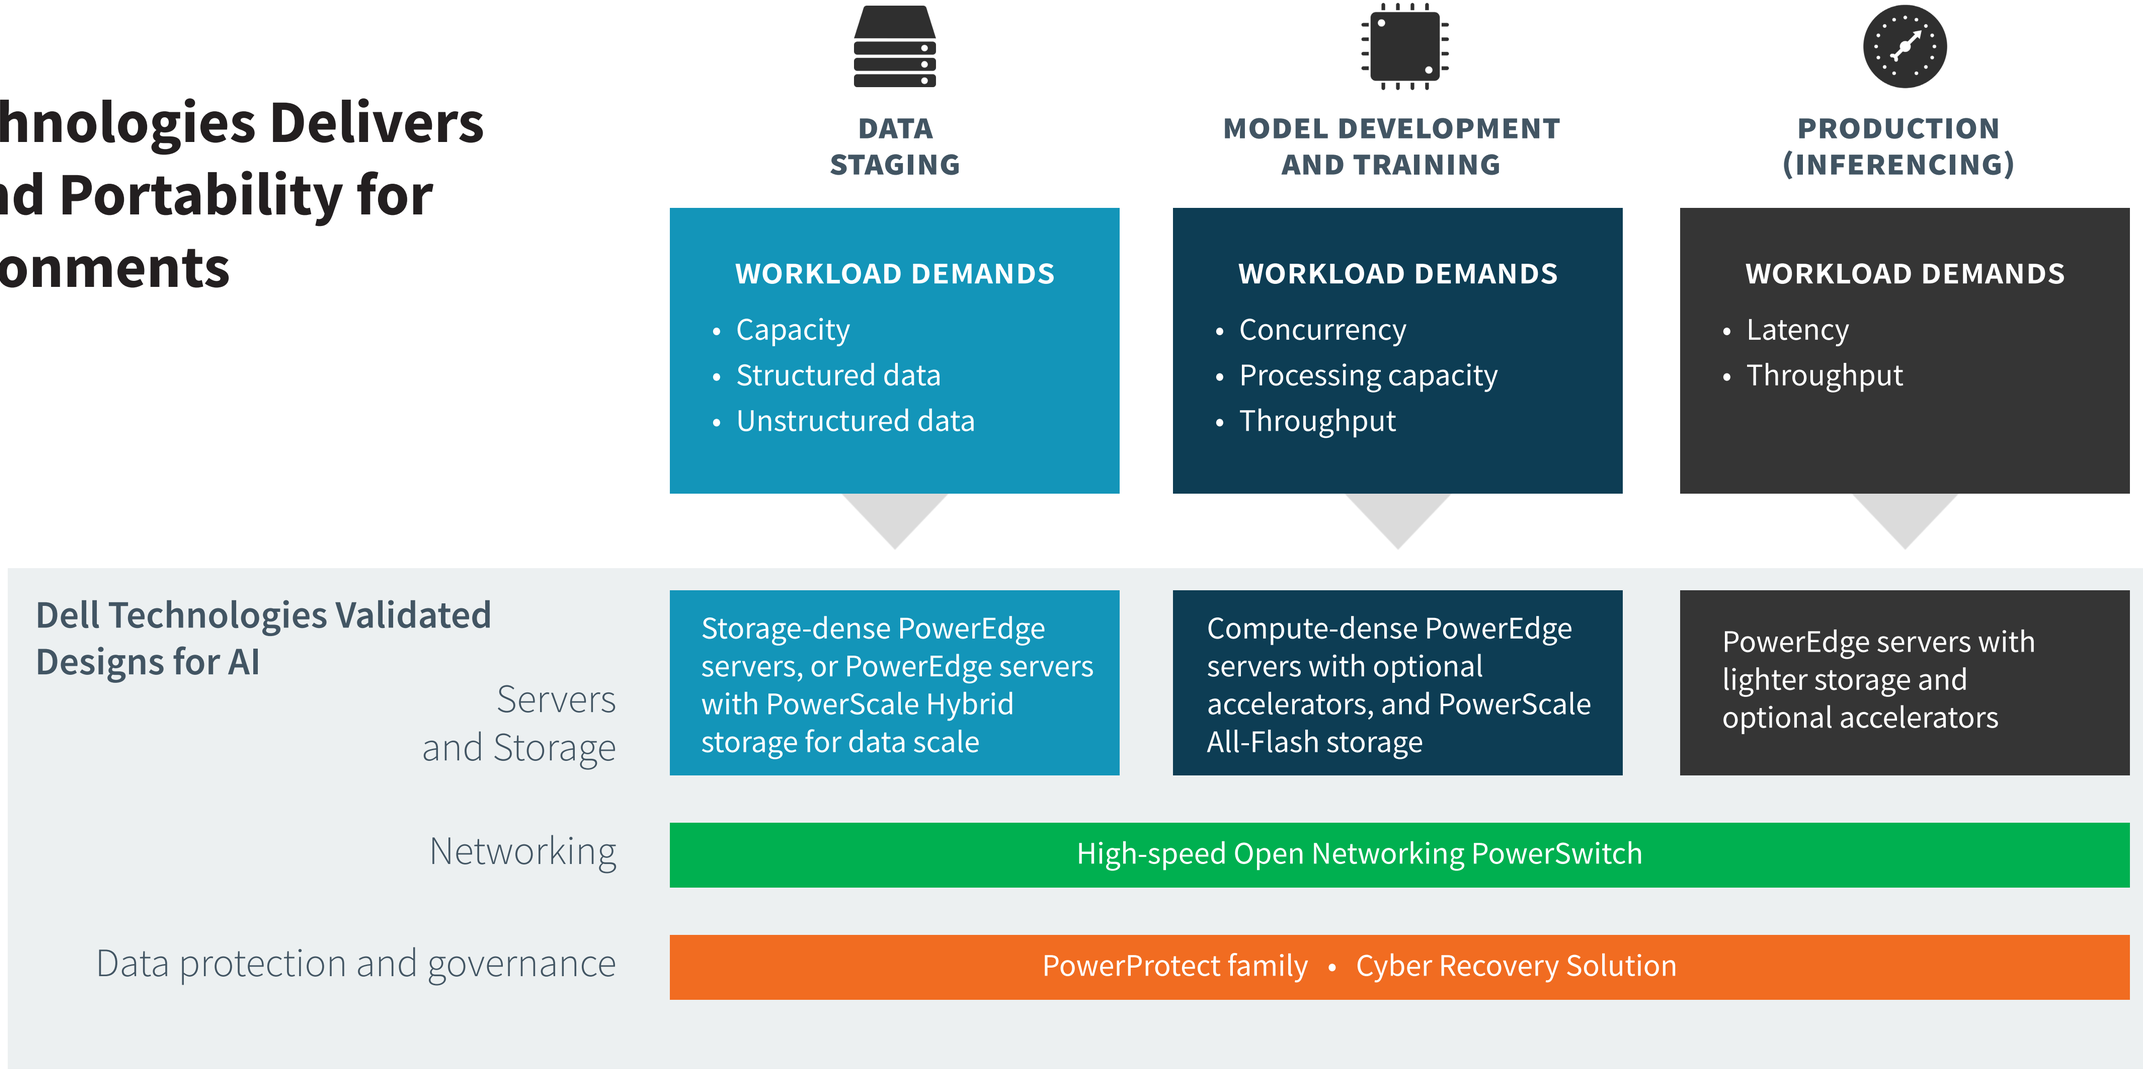

Back to Contents

How Dell Aligns with the Four Infrastructure Essentials for Data Pipeline and Data Lake Environments

Maximizing Infrastructure Performance and Utilization

Dell EMC servers and storage offer multiple high-performance options. For example, the PowerEdge XE8545 claims to provide 6-7x the machine learning performance over today's accelerators with PCIeGen4 A100 SXM4 GPUs with the highest peer-to-peer bandwidth. PowerScale is designed to eliminate AI I/O bottlenecks for high performance, concurrency, and scalability, delivering faster training and validation of AI models, higher model accuracy, improvements in data science productivity, and maximization of ROI for compute investments.

Hybrid/Multi-cloud Capability

Dell Technologies offers multiple cloud storage options for both multi-cloud use cases and native cloud experiences. PowerScale for Multi-cloud is a scale-out NAS cloud storage service offering built on PowerScale, delivered as a managed cloud service offering that can be directly connected to the major hyperscale cloud providers through a very fast, direct interconnect. PowerScale is also offered as a native Google Cloud service empowered by Dell services and enterprise SLAs.

Data Management and Security/Governance

Dell EMC servers and storage offer multiple data management, security, and governance features. For example, Dell EMC PowerScale centralizes governance processes, including data management, data security, data compliance, data protection, and drive encryption to help with compliance requirements such as PCI, HIPAA, Sarbanes-Oxley, FISMA, or GDPR. Dell EMC PowerScale also includes Dell EMC DataIQ, a powerful data management tool that simplifies the ability to find and understand data across all of your file and object data platforms, including non-Dell EMC Storage platforms, across on-premises and cloud environments.

Data Durability/High Availability

Dell EMC servers and storage are designed for enterprise environments with business and mission-critical levels of high availability to ensure that IT infrastructure will stay available and predictable as demands evolve and scale.

Back to Contents

Conclusion

Every AI initiative begins with a business challenge or opportunity. Data scientists convert a use case into a data science problem and then develop a data science solution, at which point IT takes over to make the solution production-ready. These initiatives typically require very large data sets needed for training the AI model. These large volumes of data need to be moved in and out of various systems while at the same time complying with the organization's data governance and ensuring data availability and durability. Without the right infrastructure foundation, organizations lose precious time and money.

Designing the right infrastructure environment requires multiple capabilities but for production environments, hybrid/multi-cloud capability, maximizing hardware/infrastructure utilization, data durability/high availability, and data security/governance are often among the top infrastructure consideration factors .

Dell Technologies

Businesses encounter many hurdles taking AI/ML initiatives from proof of concept to production and success is often determined by having the right infrastructure.

Dell Technologies can support businesses' end-to-end AI needs from proof of concept (POC) all the way to large-scale production, with a range of systems for every AI scenario, allowing businesses to grow their capabilities at their own pace as their needs shift and as their data sets grow.

LEARN MORE

DELL Technologies

CUSTOMER PROFILE

Medacist
The Key to Healthcare Automation

"As we grew from 600 to over 2,000 clients, we were able to minimize growth in management overhead while delivering the scalability and performance our increased workloads required thanks to our Dell Technologies infrastructure."

"With an infrastructure that combines PowerEdge servers for compute horsepower and PowerScale storage for high throughput which fuel our AI & analytics innovation engine, Medacist has achieved remarkable performance improvements. Automation and consolidated processing of client data allows our algorithms to deliver highly accurate analytics, which our clients count on."

"Our clients can improve patient safety or even save lives because the Dell Technologies solutions support analytics based on massive amounts of data."

David J. Brzozowski Jr.
Chief Technology Officer, Medacist

Research Methodology

To gather data for the research report cited in this ebook, ESG conducted a comprehensive online survey of IT and business professionals from private- and public-sector organizations in North America (United States and Canada) between September 17, 2020 and September 26, 2020. To qualify for this survey, respondents were required to be IT professionals familiar and involved with evaluating, purchasing, and/or managing storage associated with AI initiatives for their organization. All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents.

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 325 IT and business professionals.

RESPONDENTS BY NUMBER OF EMPLOYEES

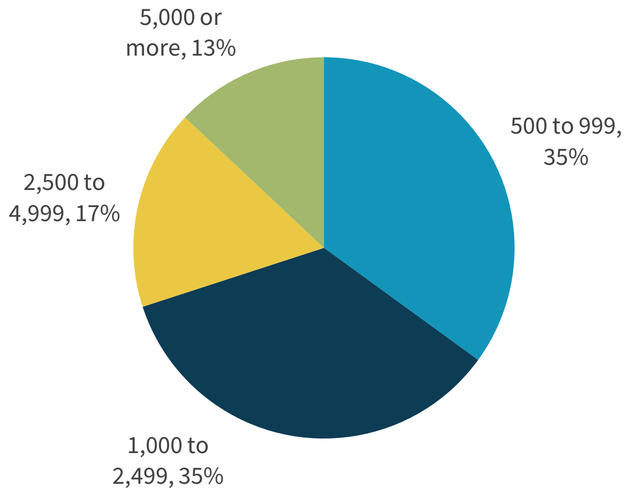

RESPONDENTS BY INDUSTRY

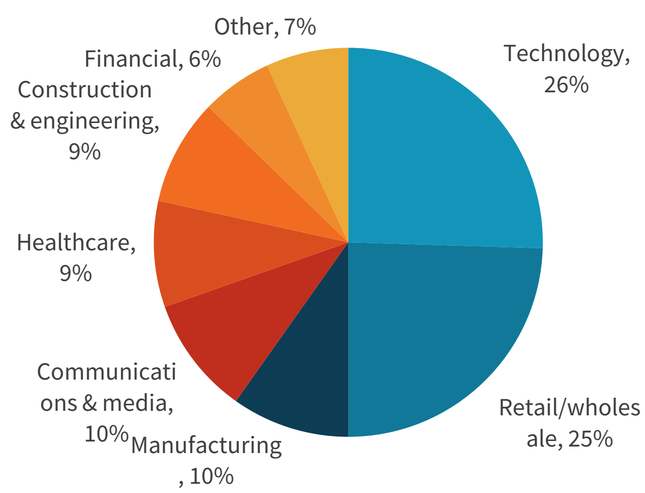

RESPONDENTS BY AGE OF ORGANIZATION

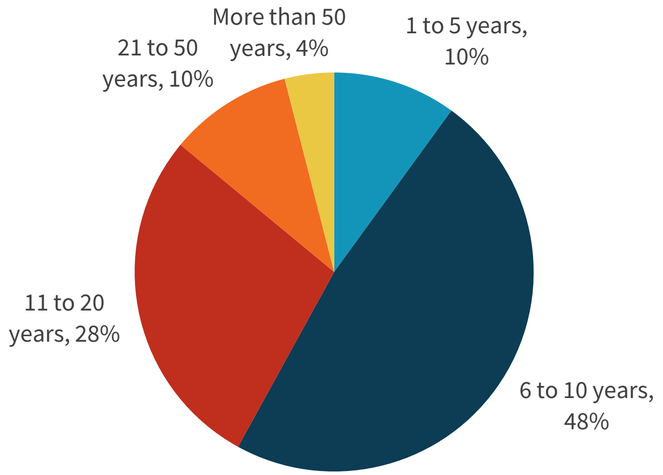

Back to Contents

Enterprise Strategy Group is an integrated technology analysis, research, and strategy firm providing market intelligence, actionable insight, and go-to-market content services to the global technology community.