Enterprise Strategy Group

SHOWCASE

JUNE 2025

Achieving End-to-end Observability With NetApp Data Infrastructure Insights

Simon Robinson, Principal Analyst; and Monya Keane, Senior Research Analyst

Abstract: Many organizations are facing critical challenges related to data infrastructure visibility. Key issues include expanding IT complexity, data growth, sprawl across multi-cloud environments, and increasingly diverse storage infrastructures. Existing approaches simply fail to provide comprehensive visibility. Fortunately, NetApp Data Infrastructure Insights offers end-to-end, unified observability across storage platforms and cloud providers, powerful visualizations with real-time dashboards, correlation analyses to identify relationships between components, machine learning-powered alerts and anomaly detection, and more. Its capabilities translate into improved customer satisfaction, faster issue resolution, and lower costs.

Overview—The Data Visibility Problem

Data is the lifeblood of every organization. Thus, understanding the data environment has never been more important. That effort includes finding ways to take advantage of advanced analytics and AI to optimize operations, defend against data security threats, achieve compliance, and, of course, manage costs more effectively.

However, according to Enterprise Strategy Group research, 63% of organizations said they regularly encounter visibility issues across their data systems.[1] Those visibility gaps hamper their ability to get the most out of new technologies. They also introduce risks from a security and governance perspective, and they drive-up costs. Moreover, visibility problems often directly impact an organization's operations by interrupting availability (taking key applications and services offline) and making it harder to monitor performance. Organizations report they are being challenged on multiple fronts due to factors such as:

- **Data growth.** Data growth is accelerating with more data types, especially unstructured data. Although storage costs per terabyte continue to fall, most organizations recently surveyed said they had at least one petabyte of data installed on their corporate servers and storage systems.[2] Many organizations have a similar amount stored off premises in cloud, colocation, and/or edge environments. As those data volumes continue to grow, costs continue to rise. Part of the challenge is that some workloads are expanding at unpredictable rates, which is likely to continue or increase as organizations deploy more AI workloads in production. According to Enterprise Strategy Group research, 69% of organizations find data storage infrastructure demands and associated spending to be difficult to predict.
- **Data sprawl.** Data growth across multi-cloud infrastructures is creating data entropy and making visibility into the overall environment difficult, especially with the siloed nature of data storage.
- **Infrastructure complexity.** Six in ten respondents to an Enterprise Strategy Group survey said overall IT complexity has increased over the last two years, with one in five saying the complexity increase was

[1] Source: Enterprise Strategy Group Research Report, *Navigating the Cloud and AI Revolution: The State of Enterprise Storage and HCI*, March 2024. All Enterprise Strategy Group research references and charts in this showcase are from this report unless otherwise noted.

[2] Source: Enterprise Strategy Group Research Report, *Achieving Cyber and Data Resilience*, September 2024.

substantial.[3] In another survey, 68% said IT infrastructure complexity was slowing down the rollout of new digital initiatives.

As the custodians of the data environment, IT teams see many of these challenges prominently within the physical data storage infrastructure. Innovations in storage technology (especially the emergence of all-flash systems) eliminated many of the I/O bottleneck issues that plagued HDD-based storage systems for years. But when a performance issue arises, the finger of blame often still gets pointed at the storage.

The evolution of the storage infrastructure marketplace itself has compounded the challenges. In many organizations, data is scattered across disparate storage silos sold by different storage vendors, and it is accessed by applications across a variety of file, block, and object protocols. As a result, gaining a true understanding of the storage environment in aggregate can be extremely problematic. The challenges manifest in multiple ways (see Figure 1).

Figure 1. Data-related Challenges With On-premises Storage

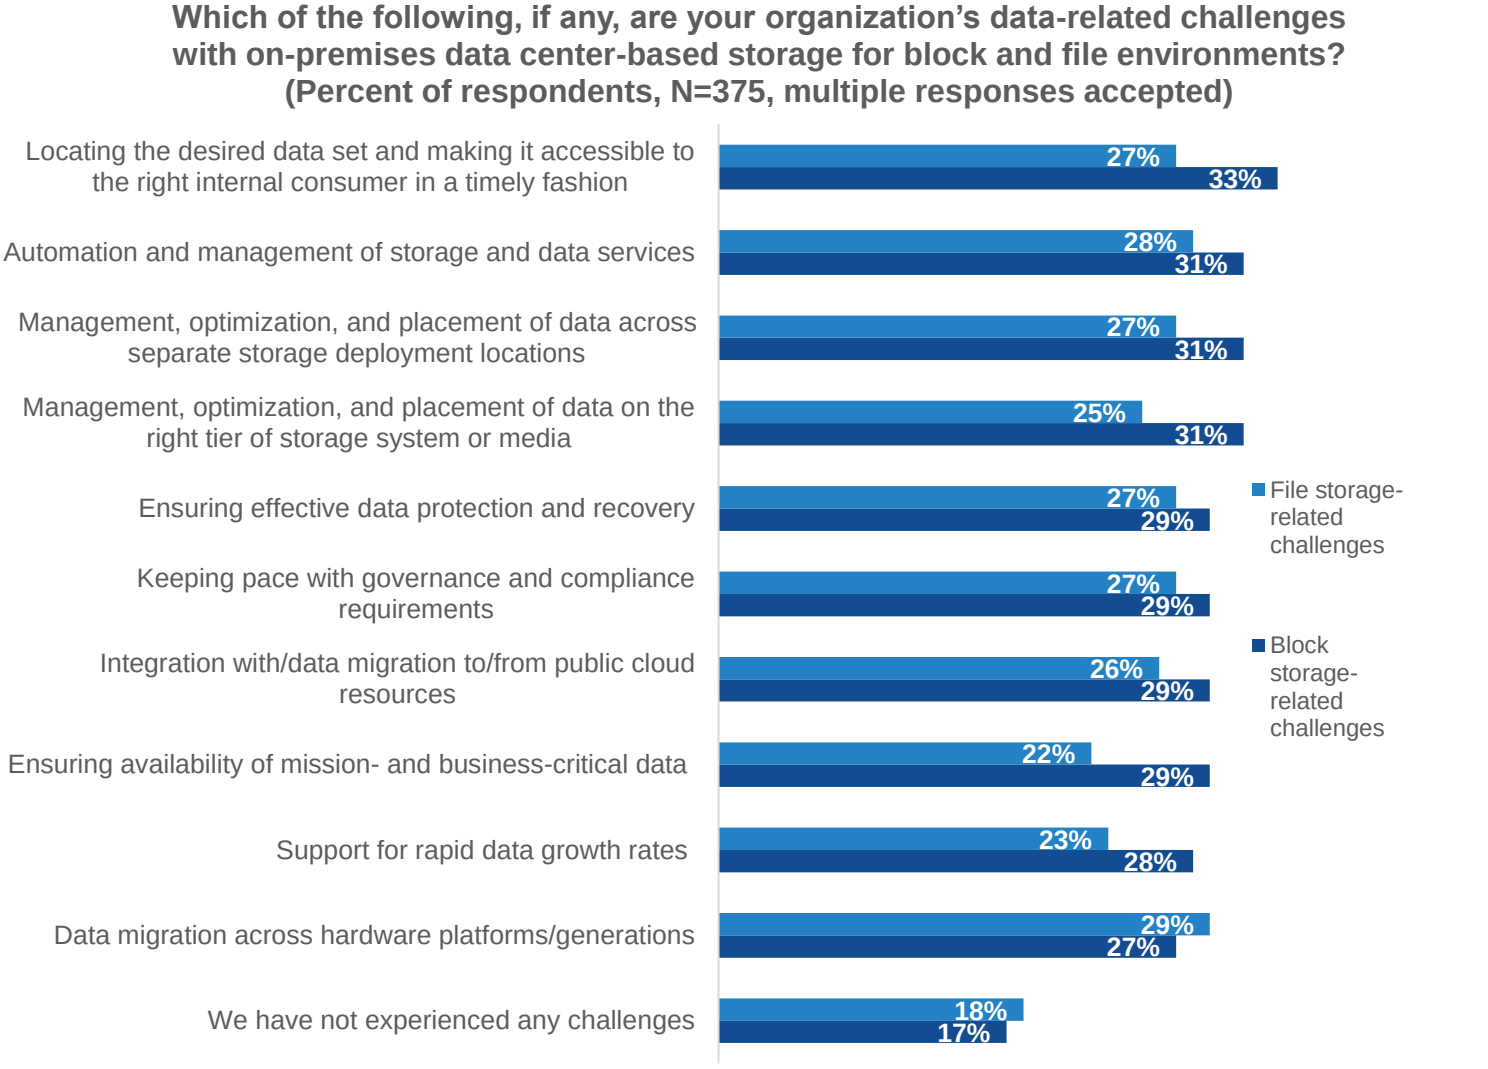

Source: Enterprise Strategy Group, now part of Omdia

[3] Source: Enterprise Strategy Group Research Report, *2025 Technology Spending Intentions Survey*, December 2024.

The key issues for both block and file storage center on locating desired data sets and making them available to the right internal consumers in a timely fashion; automating and managing storage and data services; and placing, managing, and optimizing data across varied storage deployment locations. Additional issues, such as migrating data across hardware platforms and generations and managing the placement of data on the right tier of storage, also feature prominently.

Storage operations are also becoming more complex as data, applications, and services no longer reside within a common infrastructure stack in a single location. They are distributed across a range of on- and off-premises locations, running on diverse platforms, operating systems, hypervisors, and, increasingly, containers. Additionally, many newer applications, especially AI workloads, are growing at rates that are difficult to predict, making capacity planning more complex. Enterprise Strategy Group research suggests that these challenges, which, for many organizations, are longstanding, will become increasingly problematic as they continue to explore the potential of emerging technologies such as AI.

Homegrown and Piecemeal Solutions Are Falling Behind Modern Demands

Organizations have historically faced challenges gaining visibility into their overall data environment due to a lack of tools or capabilities able to support them in the effort. This is not because such tools are unavailable; a study by Enterprise Strategy Group found that 56% of organizations have more than ten observability tools currently deployed.[4] The issue is that those tools focus on one specific part of the environment (e.g., a particular vendor's environment). It's a classic example of not being able to see the forest for the trees.

This tool fragmentation leaves teams struggling to pull together a cohesive view, and many IT organizations resort to writing their own custom scripts or tools to bridge these gaps and monitor their data environments. This may seem like a reasonable direction at the outset, especially given the dearth of alternatives, but this approach quickly begins to break down as the environment grows in scale, diversity, and complexity. Consequently, many organizations are in desperate need of a viable, comprehensive alternative. Dedicated monitoring tools may be good enough for a little while, but they are subpar compared with tools that inherently understand the full storage infrastructure. An organization dealing with these issues needs the help of a dedicated partner team that can support the organization's needs over time.

A Better Approach—Introducing NetApp Data Infrastructure Insights

NetApp is addressing these challenges with a strategic focus on simplifying and unifying data infrastructure. NetApp offers a comprehensive, intuitive product suite to reduce the need for custom-built data visibility tools.

NetApp's core monitoring and observability offering is Data Infrastructure Insights, which is designed to help organizations achieve peak performance and availability via a range of advanced capabilities. The key benefits of Data Infrastructure Insights include:

- **Improved customer satisfaction** by preventing major infrastructure issues before they affect end users.
- **Proactively prevented failures** by reducing mean time to resolution (MTTR) by up to 90%, according to NetApp.
- **Optimized performance and costs** by reducing cloud infrastructure costs by an average of 33%, according to NetApp.

Key characteristics include:

- **End-to-end unified observability.** Data Infrastructure Insights is deeply integrated with NetApp data storage solutions, as well as third-party vendor storage platforms, to provide centralized visibility across a full data

[4] Source: Enterprise Strategy Group Research Report, *Distributed Cloud Series: Observability and Demystifying AIOps*, August 2023.

infrastructure. Support includes native integrations with major cloud providers such as AWS and Azure, providing a unified view of resources across different clouds, as well as storage providers, including Dell, Pure, and IBM. It also extends beyond the storage environment to collect data metrics such as logs from cloud instances, virtual machines, and network devices, as well as performance and configuration data from Kubernetes, VMware, Azure, NetApp Files, and more.

- **Powerful visualizations.** Real-time dynamic visualizations provide views of infrastructure topology, availability, and utilization that enable informed decision-making about resource allocation, scaling, and cost efficiency. IT teams can quickly identify underutilized storage, misallocated resources, and related opportunities to optimize the environment. Preconfigured and customizable dashboards and predefined reports offer immediate insights into the health and performance of the infrastructure. And, as a SaaS-based tool, Data Infrastructure Insights is always up to date.
- **Correlation analysis.** Identifying relationships between metrics and events reduces MTTR by enabling IT teams to troubleshoot, find root causes, and optimize performance (see Figure 2). End-to-end views enable teams to quickly understand how changes in one component or metric can affect others. For example, if a spike in disk I/O latency and a corresponding increase in database response time is detected, Data Infrastructure Insights correlates those metrics to make it easy to identify whether the storage system is causing the performance degradation. Correlation analysis also helps with planning by showing the relationships between different components and can improve resource allocation and capacity planning.
- **Alerts and anomaly detection.** Machine learning algorithms baseline normal activity and then alert the IT team when unusual behavior has been detected.

Figure 2. Data Infrastructure Insights Correlation Analysis

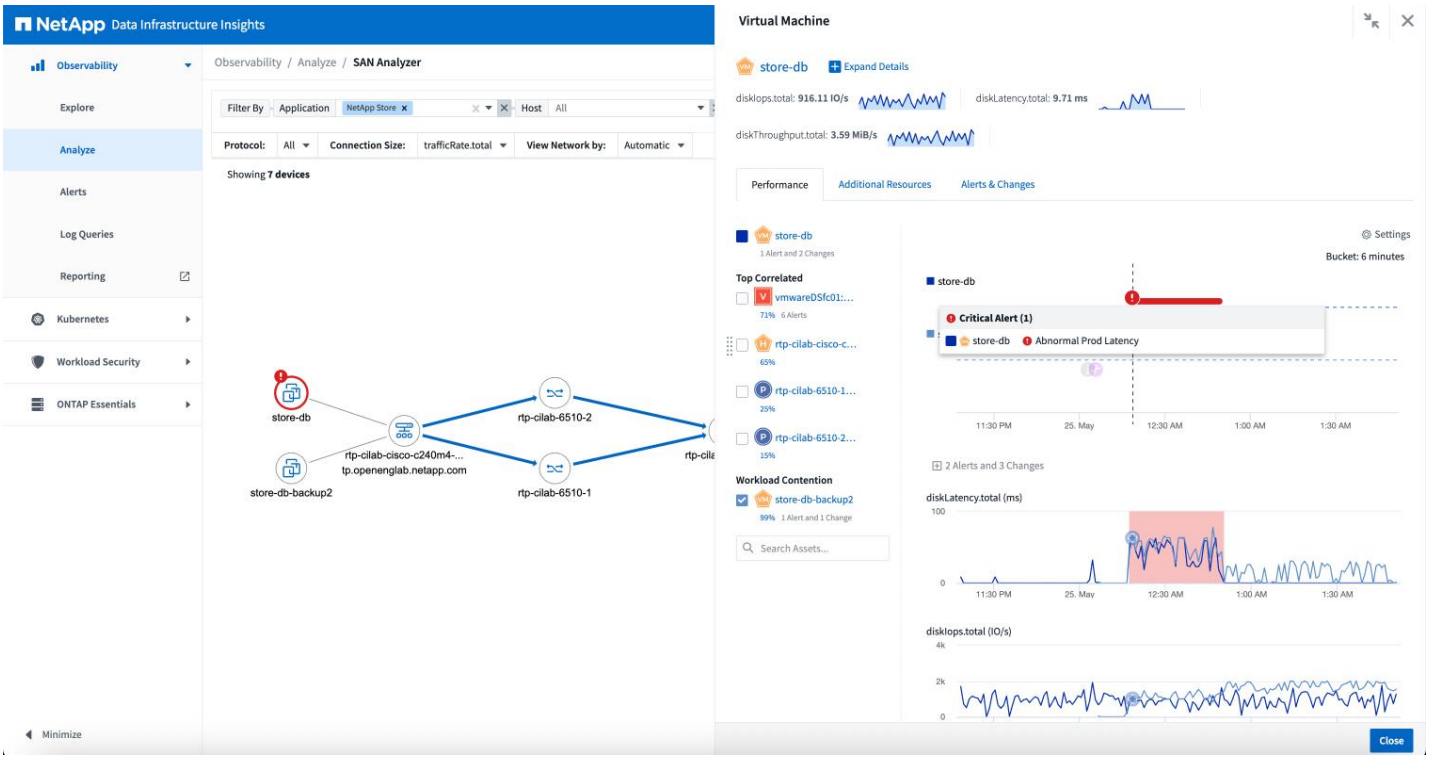

Source: NetApp

Detailed topology views of connected infrastructure resources, anomaly detection highlighting behavioral changes over time, and change analysis pinpointing cause-effect relationships empower teams to troubleshoot with precision.

Conclusion

"Build versus buy" has always been a tricky question for IT decision-makers. Many organizations that opted to build their own storage monitoring and observability tools or cobble together various third-party point solutions found they still didn't have the visibility that they needed across their total storage environment.

Those organizations should be looking for a solution that can offer all the capabilities they need across the full infrastructure, be deployable out of the box, yet be customizable. Such a solution, especially when it is supported by a trusted partner that ensures it is always kept up to date, is going to be far more valuable than something that was "piecemealed" together.

Learn more at netapp.com/DII and request a personalized demo of NetApp Data Infrastructure Insights to see how it can benefit your organization.

About Enterprise Strategy Group
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.

contact@esg-global.com
www.esg-global.com