

White Paper

Driving Innovation Through Data Storage Modernization

Sponsored by: IBM

Ashish Nadkarni
October 2019

EXECUTIVE SUMMARY

IDC anticipates that in the next two years, the number of "digitally determined" organizations with a fully integrated enterprisewide technology architecture will grow from 46% to over 90%. Through massive tech industry growth bolstered by 3rd Platform solutions, digitally motivated businesses are getting more precise in their messaging, targeting customer needs, and bringing new experiences to market through digital transformation (DX). This is leading to:

- An app revolution with next-generation cloud-native apps being created to support hundreds of DX use cases. Innovation is everywhere and anywhere customers need it.
- Intelligent applications, digital platforms, and technologies that enable enterprises to overtly and covertly interact with customers around the clock.
- Ease of sharing and analyzing data and information, which is fostering the exploration of new ways to utilize and monetize data.

A modern data infrastructure serves as a foundation for organizations seeking to be digitally determined. The defining characteristics of this enabling infrastructure are:

- **Hybrid cloud infrastructure delivered as code.** This approach enables enterprises to take advantage of the benefits of traditional datacenters and public and private cloud services delivered as a common pool of software-defined resources accessible via APIs.
- **Cyber-resilient architecture.** A cyber-resilient architecture enables organizations to implement robust security in their infrastructure, which in turn protects not just the various hardware elements but also the data that lives on these elements.
- **Open by design.** An open design enables deployment flexibility, thereby minimizing challenges with vendor lock-in and platform integration with open source-based application stacks where appropriate.
- **AI and analytics-friendly.** An AI and analytics-friendly infrastructure enables enterprises to incorporate AI-driven data insights and analytics into every process and therefore gain unprecedented agility necessary for their transformation.
- **Consistent data availability and recovery.** An effective storage system protects against unplanned business downtime by providing mechanisms to speed data recovery, simplify management, lower storage costs, and improve data security.

By making timely investments in a modern storage infrastructure from a trusted vendor, organizations can remain in firm control of their future, enabling them to be digitally determined and lead with data as a differentiating asset.

SITUATION OVERVIEW

Infrastructure Modernization in the Context of Digital Transformation

Organizations seeking to be digitally transformed have established investments that deliver or enhance strategic business outcomes. They track investments in people, process, and technology in terms of key performance indicators (KPIs) aligned with their business objectives. IDC's research shows that 46% of organizations worldwide are making the strategic, organizational, technological, and financial decisions that will set them up to digitally transform their organization in the next several years. Leaders in these "digitally determined" enterprises have the intestinal fortitude to make the technological, organizational, and financial decisions required to extend their firm's competitive differentiation well into the future. This is followed by a disciplined execution that aligns digital initiatives, technologies, and road maps across the entire organization. Using the results from an IDC global study of 1,987 digital leaders, IDC has developed a blueprint for digital success that relies on four key elements:

- **A single enterprise strategy.** All digitally determined organizations rally around a single strategy, as opposed to organizations that are trying to coordinate multiple digital strategies rooted in the various lines of business and functional areas.
- **Resoluteness to make the required organizational and cultural changes.** Digitally determined organizations are two times more likely to have digital embedded throughout the organization, as opposed to residing in a central digital group.
- **A long-term investment strategy based on the principle that digital is inherently valuable to the business.** Digitally determined organizations are more likely to fund their digital transformation through a capital budget, as opposed to a short-term funding mechanism.
- **A single integrated digital platform to scale technology innovations.** Digitally determined organizations are focused on scaling digital operations and therefore are working toward a single digital platform.

An unfortunate side effect of digital transformation at many organizations is the creation of islands of innovation, each of which is driven by its own infrastructure. On the front end, they may have innovations around chatbots and mobile applications, and on the back end, they may have innovations around real-time inventory management or self-healing assets. And in between, there are plenty of innovations around data as a service, fraud monitoring, and P2P payments. The challenge they face is connecting these innovations to one another and with their existing enterprise systems to achieve scale.

Applying these principles to infrastructure – and specifically to data infrastructure – has profound implications on how organizations approach infrastructure modernization initiatives. It starts with a single enterprise strategy that lays the groundwork for a long-term investment strategy, which is in turn driven by a single integrated digital platform to scale technology innovations.

Digitally determined organizations are inclined to bring their innovations to scale through a single digital platform, with approaches that vary. The end goal however is a fully integrated enterprisewide technology architecture enabling digital products, services, and experiences while modernizing and integrating the internal IT environment.

This enterprise technology architecture is underpinned by a data (and storage) infrastructure. Organizations ingest data from internal and external sources. Data comes in through connected assets, employees, and connected processes and other API-based data streams. Data also comes in

through external processes with customers and ecosystem partners that are using a variety of technology, such as bots, mobile devices, and IoT sensors. The infrastructure needs to enable insight or actions from this data circulating in the organization before the expiry of its "time to value."

Modern Storage Infrastructure Accelerates Digital Determination

A data storage infrastructure platform built around an intelligent core – where the algorithms, the code, and the models live – enables organizations to efficiently glean insights and actions from the data. This platform enables secure and consistent delivery of resources that are tied to and enabled by service-level objectives, which are in turn tied to the organization's strategic business objectives. Trusted storage solutions are a key element of a data-first infrastructure.

Hybrid Cloud Infrastructure Delivered as Code

Most enterprises today use multiple clouds in their infrastructure. Usually this is a combination of various private and public cloud deployments. This approach allows enterprises to take advantage of the benefits of traditional datacenters and public and private cloud services that support current- and next-generation applications. Shifting toward delivering a true hybrid cloud experience from the storage infrastructure point of view requires a common orchestration and management framework.

A hybrid cloud storage infrastructure enables organizations to:

- **Deliver resources across clouds efficiently and flexibly.** Hybrid cloud storage makes data accessible across multiple clouds and for multiple use cases and enables the use of public cloud for data backups, disaster recovery, and data retention.
- **Deliver infrastructure as code.** Hybrid cloud storage enables enterprises to deliver public and private cloud infrastructure as a common pool of software-defined resources accessible via APIs. These resources can be seamlessly integrated with developer/operations workflows (e.g., DevOps) or as an integral part of business workflows.
- **Enforce governance and service objectives across clouds.** Hybrid cloud storage connects public and private cloud and traditional datacenters with a single security, data governance, and common access model and enables a single-pane-of-glass management layer for all resources independent of where they exist.

Cyber-Resilient Architecture

As organizations ratchet up their reliance on data, they are more exposed to threats – which can have a debilitating effect on their business, never mind the embarrassment and loss of reputation. A storage infrastructure with a cyber-resilient architecture enables organizations to implement robust security in their infrastructure, which in turn protects not just the various hardware elements but also the data that lives on these elements. This can be achieved by deploying pervasive at-rest and in-flight data encryption, role-based access control, cyberattack detection, secure snapshots, and air gap isolation.

Open by Design

An open design enables deployment flexibility, reducing vendor lock-in and increasing ease of integration with open source-based application stacks where appropriate.

AI and Analytics-Friendly

An AI and analytics-friendly infrastructure enables enterprises to incorporate deep insights and analytics into every process and therefore gain unprecedented agility necessary for their transformation. This means that the storage infrastructure must support the different data management requirements for

business, human-generated, and machine-generated data. The requirements for storage and analytics of machine-generated data for machine learning are quite different from the requirements for storage and analytics of deep learning data. For example, machine-generated data usually requires distributed processing at the edge enabled by a scalable, shared repository in a central location (core). The type of storage that may have been deployed traditionally may not fit a machine learning environment.

Consistent Data Availability and Recovery

Data protection and archiving are long-term decisions that can have a profound impact on how effectively an organization can meet its service-level objectives. An effective storage system protects against unplanned business downtime by providing mechanisms to speed data recovery, simplify management, lower storage costs, and improve data security. The architecture must be sustainable to cater to specific decision points made in support of the organization's overall data protection and archiving strategy. It must:

- **Protect physical and virtual environments.** The system adequately and intelligently protects diverse data sets and information types in an organization. In most IT infrastructure, it is common to find virtualization objects (virtual images and containers), one or more relational databases, and user data such as emails, documents, and spreadsheets.
- **Support new applications.** The system scales to host newer data sets and information types as they are created. If an organization implements a new document management system or a database type, the IT staff should be able to leverage the same system to protect these new data types.
- **Store data in an efficient manner.** The system offers the ability to deduplicate and compress primary, backup, and archive data. It can intelligently deduplicate data sets that are archived and backed up while maintaining logical separation between the two data sets.
- **Offer flexibility to meet retention requirements based on business value.** The system is customizable and supports one-off situations that may require custom scripting and/or exception handling. Certain types of data sets may not conform to standard backup or archiving policies and may require special handling from a platform perspective.
- **Ensure data integrity and security.** The system has adequate authorization, authentication, and audit mechanisms in place for the protected/archived data at rest and provides inline and at-rest encryption with key management as an additional level of security.
- **Deliver scalability.** The system scales to support unanticipated growth in one or more data sets. The platform should also be upgradeable in place and minimize disruptions during such upgrades.

Voice of the Business – The IT Decision Maker Perspective

IDC interviewed IT decision makers (ITDMs) across three verticals as part of this study. The storage infrastructure in any organization must enable IT stakeholders to achieve their business objectives. The following sections discuss enabling characteristics of the storage infrastructure that ITDMs consider to be important.

Healthcare

For ITDMs in the healthcare industry, key priorities include improving performance, scalability, and agility of their storage solutions (e.g., the use of NVMe to reduce I/O latency). Storage consolidation is top of mind for them: Interviewed organizations are actively seeking to reduce the number of applications in place and look to data storage solution providers for assistance in consolidating the number of platforms in use. Similarly, they seek a simple solution that can be automated.

Financial Services

For ITDMs in the financial services industry (FSI), top priorities include improvement in storage efficiency (i.e., how capacity is used) and alternatives to SAS/SATA that improve performance without taking on additional cost. They are investing in software-defined networking as they move to new storage solutions and seek vendors that have software-centric solutions instead of hardware-centric solutions.

Many enterprises in this industry are aggressively moving to a hybrid cloud infrastructure and seek vendors that can deliver a seamless data mobility experience between on-premises infrastructure and the public cloud. Enterprises in the FSI move toward a more integrated solution as part of their hybrid cloud infrastructure and believe the days of piecemeal solutions are coming to an end. In the same vein, they consider data infrastructure to be a strategic asset for their entire organization.

As can be expected, security is also a top priority for these firms. Maintaining strong encryption practices to ensure that data is secure even if security is breached is a critical requirement.

Oil and Gas

For ITDMs in the oil and gas industry, a crucial requirement is the ease with which IT can move data and applications to a hybrid cloud infrastructure from traditional IT environments as cost effectively as possible.

Enterprises expect vendors to offer transparency in terms of costs and solutions that maximize flexibility and ease of access to data. They also expect vendors to offer scalable solutions that allow for upsizing and downsizing of storage resources based on demand. Also, by providing solutions that include cloud-based preventive maintenance options, the storage infrastructure can be deployed in various edge locations with tighter service-level objectives in terms of fast recovery, resiliency, and fewer demands on infrastructure personnel.

NVMe Arrays Enable Consolidation

"NVMe protocol on these arrays is finally going to give us the ability to meet the needs of our applications across the board. We used to buy different high-end systems to meet heavy read/write requirements for few applications. Now I'm hopeful that with NVMe-based storage, we can consolidate our disparate platforms."

—Vice president of IT, financial services firm

Overview of IBM Storage

Hybrid Cloud Storage

IBM hybrid flash storage delivers easy data mobility across on-premises and hybrid cloud deployments with container-ready storage. This offering combines IBM all-flash solutions (built with IBM Spectrum Virtualize software) with AI-infused cloud-based storage management and support from IBM Storage Insights and IBM Spectrum Virtualize on IBM Cloud or Amazon Web Services.

Storage for AI and Analytics

IBM Storage for AI and big data enables seamless adoption of AI across the organization. This offering unifies data and metadata across the enterprise to provide a shared data service and a "single source of truth" as a trusted foundation for analytics and data science. As clients mature in their use of AI and analytics, IBM Storage enables the creation of a single, scalable, and flexible data lake built upon IBM's software-defined storage. This enables organizations to collect, organize, and analyze data to infuse AI throughout their business – on-premises and into the cloud.

Cyber-Resiliency

IBM Storage enables multisite resiliency and copy data safeguards. Cybersecurity is built into the product architecture, further enhanced by air-gapped offerings and features such as no overhead compression and encryption.

Modern Data Protection

IBM Spectrum Protect Suite utilizes a policy-driven approach to simplify data backup and recovery for virtual machines, physical file shares, and applications running on virtual and physical machines. IBM Spectrum Protect delivers life-cycle management with low total cost of ownership with backup and recovery performance that can be tuned to meet stringent service-level objectives. IBM Spectrum Protect Plus includes global search and instant data recovery, which improves resiliency, and self-service data reuse, which improves the quality and speed of analytics, development, testing, and reporting. IBM Spectrum Protect Plus is easy to deploy as a virtual appliance, and the solution's agentless architecture is easy to maintain.

Storage Modernization Is an Opportunity for Consolidation

"From an infrastructure perspective, flash and data reduction technologies are two other big targets for modernization. They enable us to reduce the number of applications. We end up with multiple platforms per hospital system. It is challenging to manage all of them, and it is not realistic."

—Vice president of IT, healthcare services firm

CHALLENGES AND OPPORTUNITIES FOR IBM

Organizations struggle with storage modernization primarily because it is akin to retooling engines on a plane when it is in midflight. The storage infrastructure must be modernized in alignment with business objectives, which seek to morph the organization into a digitally determined entity that uses data as one of its sources of competitive differentiation. Key opportunities for vendors and IT buyers working in collaboration with each other include:

- **Operational stability that leads to accelerated go-to-market development.** The shift to becoming digitally determined is driving several organizations to embark on new application development, which is highly custom in nature. It also introduces new operations-developer workflows such as DevOps and CI/CD pipelines. The right storage infrastructure enables not only frictionless delivery of resources, which in turn increases workplace productivity, but also deployment of modern rules of engagement, which discourages the use of "shadow IT."
- **Bridging the skills gap.** Modern storage systems – and however complex they may be – need not create a huge barrier in terms of skill sets. Such systems must be designed so that they are easier to manage and operate. This is, however, also an opportunity to evaluate and transition the broad skills base in an organization toward a structure that can best handle modernization. It is also an opportunity for the organization to shift from a siloed operations structure to cross-functional teams in which the handoffs between teams are minimized.
- **Serviceability.** Enterprises in general believe that improvements in service and support are essential to increase confidence in the storage solution. Cloud-based preventive maintenance options enable deployment of a distributed hybrid cloud infrastructure, including the placement of solutions at the edge.

ESSENTIAL GUIDANCE FOR TECHNOLOGY BUYERS

IT organizations face a unique opportunity with storage infrastructure as they seek to modernize it. The opportunity is to see storage infrastructure as an underpinning for their strategic business initiatives. Of course, this means tackling necessary attributes first. These include:

- **Resiliency and/or reliability.** Investing in technologies that provide a consistent user experience aligned with business use cases and expectations
- **Performance and scalability.** Ability to scale performance and capacity independent of each other
- **Agility and efficiency.** Ability to scale for changing data requirements driven by the business itself undergoing transformation
- **Deployment type flexibility.** Ability to deploy in a hybrid cloud mode, software only, and other type and to move workloads seamlessly between on-premises and off-premises
- **Security.** Ability to encrypt data in flight and at rest and to safeguard data at rest via mechanisms such as role-based access, single sign-on, and multifactor authentication
- **Ease of management.** A single-pane-of-glass approach for managing all resources (regardless of the type or deployment location) to enable automation and workflow integration for distributed access from various teams in the organization
- **Workload alignment and mobility.** Enabling investments in the right infrastructure for the right applications – both current generation and next generation – and the ability to move between them

CONCLUSION

The consequence of a failed digital transformation is far greater than a project failure. Organizations that can align and mobilize their resources will be able to realize the fruits of their innovation efforts. They will respond to customer needs with the rapid creation of new products and services. Those that cannot align and mobilize their resources will be too busy focused on internal processes to bring meaningful innovations to market. Digitally determined organizations are led by management teams that champion difficult decisions in areas such as capital budgets, organizational alignment, integrated road maps, technology investments, and process alignment.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2019 IDC. Reproduction without written permission is completely forbidden.

