

## White Paper

# Why Organizations Need an Intelligent Data Strategy

Sponsored by: HPE

Eric Burgener  
May 2019

## IDC OPINION

---

Digital transformation (DX) is underway in many organizations today, and IDC research notes a strong correlation between those companies that are most effectively leveraging more data-centric business models (digital achievers) and business success. The requirements of this much more dynamic business model are driving the need for infrastructure modernization and, along with it, the introduction of a host of new technologies. The infusion of these new technologies – solid state storage, artificial intelligence/machine learning (AI/ML), software-defined infrastructure, and cloud, among others – is required within information technology (IT) infrastructure to meet evolving business requirements, but it is also making management much more complex.

As enterprises continue to evolve, this complexity will only increase, and a strong argument can already be made that the more human-intensive management approaches of the recent past cannot cope with this increasing complexity. Intelligence will be key for efficient and cost-effective IT infrastructure management, and that intelligence will need to draw on more comprehensive, full-stack monitoring capabilities that go far beyond just storage and are driven by AI/ML and big data analytics. When combined with automation and an overarching data life-cycle management orientation, these technology tools will be able to optimize application service delivery much more efficiently than humans can and will be able to do so in real time – a particularly interesting feature given today's much more dynamic business environments.

These realities point to the need for an "intelligent data platform." Through a combination of more comprehensive monitoring, AI/ML, data analytics, and automation/orchestration that enables human oversight where wanted and/or needed, this will increasingly be the way that successful organizations manage their data assets and IT infrastructure. HPE has introduced such an intelligent data platform with comprehensive coverage and AI/ML-driven real-time optimization that enables intelligent management of the entire data life cycle. IT organizations interested in addressing the added complexity of the modernized infrastructure that successful businesses will require can look to HPE to understand how an intelligent data platform can drive value for their companies.

## IN THIS WHITE PAPER

---

As enterprises undergo DX, they typically also modernize their existing IT infrastructure, leveraging a wide array of newer technologies. While these environments meet the performance, availability, scalability, feature, and agility demands of more data-centric business models, they are very complex and becoming increasingly difficult for humans to manage. A new paradigm of intelligent data

management, driven by artificial intelligence, machine learning, and big data analytics, is required for IT organizations to efficiently manage these environments to get the most out of their investments. This white paper discusses the industry's evolution toward modernized infrastructure and how an intelligent data strategy will be necessary to meet the evolving requirements of today's much more dynamic and data-driven business models. It will then provide a quick review of HPE's intelligent data platform offerings that are designed to meet these requirements.

## SITUATION OVERVIEW

---

Business has undergone significant change in the past decade, and one of the key trends during that period of evolution has been an increase in the strategic nature of IT's contribution to business success. With the advent of new technologies like big data analytics and AI/ML, companies want to take advantage of the wealth of data available to them to help define new products and services, identify new markets, uncover business insights, and improve customer experience. Companies are digitizing their existing internal processes and workflows as well and leveraging the better instrumentation in modernized IT infrastructure to fine-tune everything from new customer acquisition to manufacturing operations and technical support. Improving the efficiency of IT operations to increase productivity and lower costs has also been a big part of the DX companies undergo as they move to more data-centric business and operational models.

Modernized IT infrastructure brings in a wealth of new technologies, such as solid state storage, AI/ML, software-defined infrastructure, and cloud, that are required to meet the evolving performance, availability, scalability, manageability, and agility requirements of much more dynamic business models at costs that are within budget. Hybrid cloud environments where an enterprise's workloads are split across both cloud and noncloud (i.e., traditional IT infrastructure) are quickly becoming the deployment model of choice. This modernized IT infrastructure transcends the capabilities of more traditional, statically defined designs, but it introduces significant management complexities. An argument can already be made that this complexity makes it increasingly difficult for humans to efficiently manage these environments, and the management complexity of these environments as they evolve over time is only expected to increase.

It is already becoming clear that as DX among enterprises continues, human abilities to manage infrastructure will need to be supplemented with more automated operations. We have already seen the broad use of scripting to automate simple, repetitive tasks, but the complexities of modernized infrastructure will increasingly require much more than that. Intelligence, fueled by AI/ML and informed by big data analytics, offers excellent opportunities to address this complexity, and IDC believes that over the next two to three years, we will see broad use of these approaches to make IT more responsive, workflows more reliable, and operations more efficient.

As intelligence is harnessed to address management complexities, global visibility will become very important. Traditionally, management paradigms and troubleshooting mechanisms have been very siloed. Primary research from IDC indicates that storage causes only 9.8% of downtime events, strongly suggesting that any intelligence used to manage the performance and availability of application services (which is really the external "product" of IT infrastructure) must have a global view of the status and impact of resource consumption, failure impacts, and configuration changes across the entire stack. Correlating and understanding the implications of changes across the entire stack will be difficult for humans, but AI/ML-driven big data analytics can handle this complexity, responding in real time to events in a way no human ever could.

As data evolves to becoming a key strategic asset driving business success, IT infrastructure management strategies will also benefit from a cohesive data life-cycle management orientation. As data is captured, stored, protected, used and analyzed, archived, and ultimately disposed of, it will move to different types of platforms that deliver the capabilities needed at each stage. Here again, a management intelligence informed by AI/ML and big data analytics will do a much better job of ensuring that data is optimally managed to meet IT objectives over its life cycle, which for some data is short but for other data can be decades. All in all, intelligence will be critical to managing IT infrastructure most efficiently and cost effectively.

A number of trends and features have come together over the past several years that are enabling the intelligent data platform:

- Data growth continues to explode, and by 2025, IDC predicts that the world will be generating 175ZB of data (a zettabyte is a billion terabytes) a year; this amount will include data collected from mobile computing and social media devices, the internet through Internet of Things (IoT) sensors, and machine-generated data from logs and other process and workflow-related sources; the ability to leverage this data to drive value for the organization will be critical to the success of enterprises going forward, not only to drive market success but also to efficiently optimize internal operations.
- The use of solid state technology in off-the-shelf devices like solid state disks (SSDs) will provide the performance to deal with not only the latency requirements of real-time workloads but also the throughput and bandwidth to move and analyze large data sets in a timely manner; the broad use of solid state technology with general-purpose primary workloads (and now increasingly with secondary workloads) also drives much better resource utilization and more predictably consistent storage response times under load (because of very low latencies).
- The use of AI/ML algorithms to mine data assets will drive value for organizations in terms of defining new products and services, identifying new markets, uncovering business insights, and improving customer experience; AI/ML technology is being used not only in big data analytics workloads but also in enterprise infrastructure to improve performance, availability, and efficiency and lower costs, and some vendors are making it easily available by bundling cloud-based predictive analytics platforms with the base purchase price of their infrastructure products.
- The broad deployment of cloud services as an accepted part of datacenter infrastructure design will offer increased agility compared with more static traditional infrastructure; easy worldwide yet secure access to IT services through a simple, internet-based "as a service" paradigm; the option to offload routine IT infrastructure management to third parties so that IT employees can focus on more strategic projects; and the ability to shift IT assets off the balance sheet for more flexible financing options.

Modernized IT infrastructure leverages these same technologies to enable next-generation applications (NGAs) and drive better business insights, but these same developments and technologies enable the informed intelligence necessary to optimize the management of these environments. IT organizations should challenge their storage vendors to explain how they are leveraging these and other features and technologies to help minimize management complexity.

Vendors supporting an intelligent data strategy will be able to demonstrate how their products and services simplify the deployment of IT infrastructure solutions that fuel success in today's digitally transforming business environments. That strategy must address key requirements:

- **Performance.** Varied workloads and I/O profiles challenge traditional storage solutions to consistently meet service-level agreements (SLAs) that many organizations now offer their own customers. Cost-conscious IT management will pay for performance where it is required but needs broad support for tiered environments and technologies (storage-class memory, SSDs, HDDs, tape, cloud tiering) to meet different performance requirements at the lowest cost.
- **Availability.** In today's internet-driven world, end users expect IT services to always be available. Actual uptime requirements vary by workload, but IT management needs multiple options to craft the right infrastructure to meet availability requirements at the lowest cost.
- **Automated and on-demand management options.** With the rate of data growth and the complexity of managing modernized workloads and IT infrastructure, IT administrators need to lean heavily on automation to free operators to work on more strategic and innovative tasks (rather than the repetitive operations that just "keep the lights on"). At the same time, they need to be able to quickly respond to unpredictable real-time demands, which are much more common today than they were in the past.
- **Hybrid cloud deployment.** Given that those enterprises that aren't already managing hybrid cloud environments will be implementing them in the future, IT infrastructure solutions must be built with the assumption that they will be used in that manner. Issues like common monitoring and management paradigms that span cloud and noncloud environments, simple and efficient data mobility, and broad support for public clouds are needed, and vendor road maps should be keyed to evolving hybrid cloud requirements.
- **Global management intelligence.** For IT organizations trying to manage hybrid cloud environments most effectively, siloed management intelligence is a problem. Sensor and other monitoring data must be captured from cloud and noncloud locations alike and stored in a central repository that leverages AI/ML to optimize management recommendations across the entire hybrid cloud configuration, not just in local silos. Cloud-based predictive analytics platforms would appear to be an ideal solution for this challenge, but for optimum efficiency, they will need to collect data not just from the storage layer but from other infrastructure hardware and software as well as public environments.

## The HPE Intelligent Data Platform

HPE is a \$31 billion vendor of IT infrastructure solutions, including servers, storage, networking, infrastructure software, and technical support and consulting services. The vendor offers significant choice to its customers in terms of deployment models for storage – delivering storage appliances, software-only hyperconverged infrastructure (HCI), converged infrastructure (CI), and cloud-based services in a storage portfolio that includes primary and secondary storage; block-, file-, and object-based platforms; scale-up and scale-out architectures; and a cloud-based predictive analytics platform (InfoSight). The vendor was founded in 1939; is a proven, trusted supplier to enterprises of all sizes; offers storage platforms that guarantee "six-nines plus" availability; and is one of the market leaders in all-flash array (AFA) revenue. AFAs dominate external primary storage shipments, driving almost 80% of all revenue in that market in 2018.

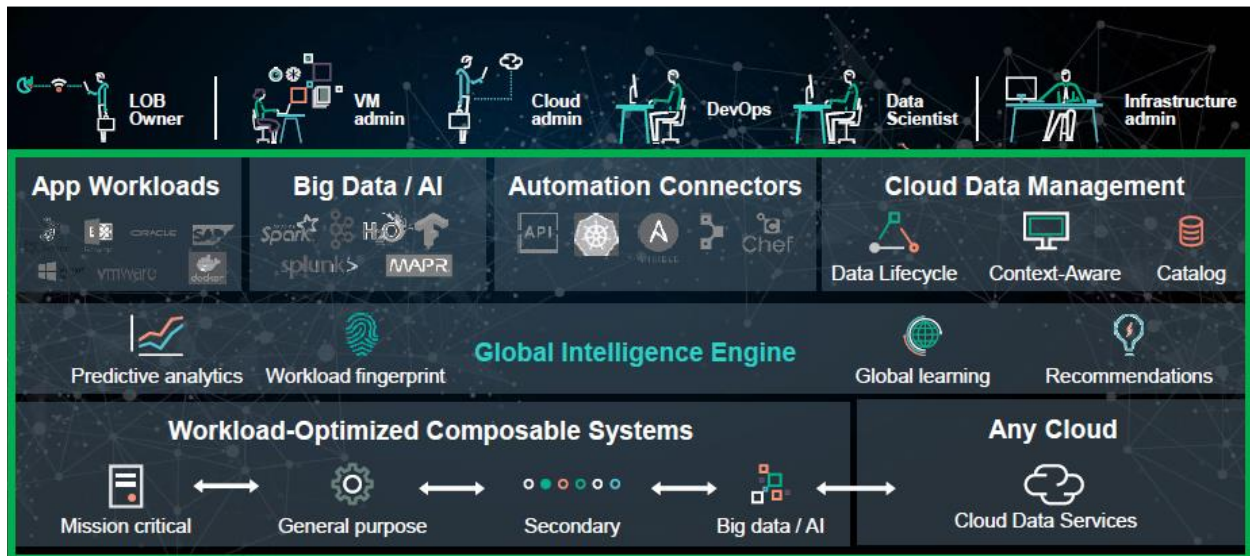
The HPE Intelligent Data Platform starts with an enterprise infrastructure foundation that includes workload-optimized composable systems that span mission-critical, general-purpose, secondary, and

big data/AI applications. HPE brands across these systems include HPE 3PAR StoreServ, HPE Nimble Storage, HPE Nimble Storage Hybrid, and HPE StoreOnce for mission-critical, general-purpose, and secondary storage. HPE Apollo systems coupled with BlueData Software (which can run anywhere) provide the foundation for AI/ML-driven big data analytics. HPE SimpliVity, HPE Synergy (coupled with storage), and HPE ProLiant (coupled with storage) provide the foundation for private clouds. HPE also supports the leading public clouds, including Amazon, Microsoft, and Google, and offers HPE cloud-based storage services as well (HPE Cloud Volumes, HPE Cloud Bank). Native data mobility across these systems, regardless of whether they are deployed in cloud or noncloud environments, enables data and application mobility for a variety of use cases spanning data life-cycle management, hybrid development and test, hybrid analytics, and modern data protection solutions.

Each of these independent systems are instrumented with sensors that comprehensively monitor system and workload metrics and generate a significant amount of telemetric data, all of which is collected and securely transmitted to HPE's cloud-based Global Intelligence Engine (GIE) (see Figure 1).

**FIGURE 1**

**The HPE Intelligent Data Platform**



Source: HPE, 2019

With its centralized repository of data, the Global Intelligence Engine leverages global, AI/ML-driven learning across more than 100,000 systems (spanning servers and storage) to predict and proactively resolve issues before they occur. It goes well beyond this with workload fingerprints (to profile the behavior of applications) and global learning to make recommendations to keep infrastructure optimized, delivering performance, capacity utilization, availability, efficiency, and/or cost goals. As conditions evolve with I/O surges, workload addition or subtraction, and failures and other events, this approach can make real-time adjustments to ensure that business objectives continue to be met.

The GIE is based on HPE's InfoSight cloud-based predictive analytics platform – the pioneer and the most mature representative of this technology in the industry today. InfoSight was architected by

Nimble Storage, an enterprise storage vendor HPE acquired in 2017, and that platform drove significant value for customers through its ability to help administrators optimize existing environments for performance, capacity utilization, availability, and efficiency at lower cost. When HPE acquired the vendor, Nimble Storage had gone through significant learning because of the extended use of the platform across its installed base and had already started to extend it beyond storage to enhance the value it provides. After the acquisition, HPE began to extend the technology to its other product lines, and InfoSight now supports ProLiant and Apollo servers as well as 3PAR StoreServ and Nimble Storage arrays and has announced plans to extend it to other product lines, including its SimpliVity HCI platform, in the near future.

The comprehensive orientation of the GIE contributes significantly to its value for customers. To provide maximum value, the GIE extends its scope far beyond just the enterprise infrastructure layer. In addition to covering HPE hardware, it also covers storage infrastructure software, virtualization platforms, and applications. Storage operating systems and management platforms, data protection software, and orchestration tools fall under its optimization purview. It detects and monitors virtual machines (VMs), providing the source of VM performance anomalies for rapid troubleshooting, suggesting VM migration for better workload balancing, and enabling the identification of dormant VMs that can be removed to optimize resource utilization. Its instrumentation provides visibility into popular enterprise workloads like SQL Server and others as well. This sort of full-stack visibility is critical to the Intelligent Data Platform's predictive, optimization, management, and remedial capabilities and is a significant differentiator relative to alternative approaches that depend much more heavily on older, much more manual approaches.

The HPE Intelligent Data Platform includes automation connectors built on RESTful APIs that enable comprehensive automation capabilities that support more traditional bare metal approaches as well as container environments and infrastructure-as-code deployments. Popular supported tools include Kubernetes, Ansible, Chef, and Puppet. The HPE Intelligent Data Platform can accelerate and bring application-aware protection for traditional business applications (SAP, Oracle, SQL, Exchange, etc.), VM farms, and NGAs based on containers, and it can be used to build data pipelines that leverage technologies like big data analytics and AI/ML with tools like Spark, Splunk, Elastic, Kafka, H2O, and MapR. Building on GIE and native HPE data mobility capabilities, the HPE Intelligent Data Platform will in the future extend to enable software-as-a-service-based data management with context awareness to manage data optimally to IT objectives across the entire data life cycle.

To provide maximum agility in building and optimizing hybrid cloud environments, customers can acquire the HPE Intelligent Data Platform components outright or purchase the entire Intelligent Data Platform offering as a service under the HPE GreenLake program.

## FUTURE OUTLOOK

---

DX is the future for successful organizations, and those organizations that cannot successfully make the transition to more data-centric business models will operate with heightened risk. IDC has identified two distinct profiles among organizations in the DX era – those that have been able to successfully move to more data-centric business models and are flourishing (digital achievers) and those that have not (and as a result are experiencing significant challenges in terms of important business metrics like revenue growth). The IT vision of the digital achievers emphasizes a focus on security, uses automation to extensively reduce management complexity and better enable scalability, leverages hybrid cloud strategies to reduce costs and improve agility, and consciously prepares their



infrastructure for the future by leveraging intelligent data strategies. It is significant that these digital achievers are much more likely to deploy newer storage technologies; in fact, they are:

- 25x more likely to deploy software-defined storage
- 23x more likely to deploy converged infrastructure
- 17x more likely to deploy AFAs
- 46x more likely to deploy unstructured storage platforms
- 11x more likely to deploy modernized data protection infrastructure

Given the scale and complexity of modernized IT infrastructure, CIOs will need to leverage AI/ML, along with big data analytics, to drive efficiencies within their organizations and manage data effectively across its entire life cycle. Even if they do not think consciously about deploying an "intelligent data platform," they will be using many of the underlying technologies that fuel such a strategy: big data and predictive analytics, workload fingerprints, deeper instrumentation across the entire hardware and software stack, automation and orchestration, AI/ML-driven dynamic optimization, and agility enablers like virtualization, containers, and cloud. These digital achievers will be converging on an intelligent data strategy, and those that adopt the mindset consciously will be able to reap the benefits of that strategy for the organization that much sooner.

## CHALLENGES/OPPORTUNITIES

---

DX requires a change in mindset for IT organizations, and it is already clear that there are those that are succeeding (digital achievers) and those that are failing (digital laggards). The actual path and implementation details will be different for each enterprise as they seek to optimize the benefits of a more data-centric approach for their particular industry and business. CIOs and their teams will be challenged to think outside the box about the best way to digitally transform their organizations for future success.

Successfully becoming a digitally transformed enterprise is harder for established organizations because they can be constrained by prior investment decisions and depreciation schedules. New enterprises often start in the cloud, delaying any onsite infrastructure decisions until later in the company's life, while existing enterprises need to take other issues into consideration. Opportunities for moving forward in the DX journey arise when new workloads are deployed and at technology refresh but can also be market driven for existing workloads and infrastructure. Having an intelligent data strategy already mapped out that looks at data life-cycle requirements can determine the best time to make these decisions and help ensure deployment success.

Performance, availability, security, compliance, agility, and cost requirements all feed into those decisions, but it is pretty clear that as enterprises make that transition, they will be moving toward technologies like solid state storage, software-defined infrastructure, and hybrid cloud. They will also be leveraging the significantly greater amounts of data available to them through better instrumentation, AI/ML-driven optimization, and global learning across both cloud and noncloud environments to improve efficiencies and drive value for the enterprise. This is what the IT organization of the future will look like.

## CONCLUSION

---

The management complexities of the modernized IT infrastructure required to meet the needs of digitally transformed organizations are undeniable, and they are only increasing over time. Given the limitations of human-driven management, it is clear that a new form of operational intelligence is required to manage these environments most efficiently within prescribed budgets. Many of the technologies being used by digitally transformed organizations – solid state storage, AI/ML, big data analytics, orchestration frameworks, and cloud – can be leveraged to help tame this infrastructure management complexity. To move to such an intelligence-driven management paradigm, IT organizations will need to deploy an intelligent data platform that has the performance and full-stack visibility to make AI/ML-driven optimization and administrative decisions in real time to meet predefined IT and other business objectives. Over the next two to three years, most digital achievers will move in this direction. Informed intelligence will become key for effective and efficient modernized IT infrastructure management in these successful companies.

HPE has introduced its Intelligent Data Platform, making the company an early entrant in this space. HPE is building on a strong foundation of proven expertise with its InfoSight cloud-based predictive analytics platform but has now extended its capabilities with the addition of a Global Intelligence Engine that provides full-stack visibility and is geared toward managing data from a life-cycle perspective. The promise of the HPE Intelligent Data Platform is to provide customers a real-time, AI/ML-driven management capability across their entire hybrid cloud infrastructure that will dynamically optimize IT environments to get the most out of IT resources while meeting an organization's evolving performance, availability, scalability, manageability, agility, and cost goals over time. With this intelligence-driven vision of IT infrastructure management, HPE is jumping ahead of its competitors with an offering that should drive real value for IT organizations.



## 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.

