

## White Paper

# HPE Primera Resets Expectations for High-End Storage

Sponsored by: HPE

Eric Burgener  
June 2019

### IDC OPINION

---

As enterprises undergo digital transformation (DX) and modernize their existing information technology (IT) infrastructure, a new set of requirements for high-end storage is emerging. Today's much more dynamic business models, looking to leverage data as a strategic asset to better inform business decisions, demand an agility that is difficult to achieve with legacy storage designs – particularly with broad public cloud usage setting high expectations in this area. Furthermore, while legacy high-end storage systems are architected with significant redundancy to minimize the impact of failures, they are unable to address problems arising outside the storage platform (which actually drive over 90% of application downtime). High data growth, combined with the hybrid cloud blueprint for datacenter infrastructure, is creating significant management complexity, and it is clear that fewer and fewer organizations will be able to rely entirely on humans to ensure that IT infrastructure meets service-level agreements (SLAs) most efficiently and cost effectively.

New technologies like solid state storage, artificial intelligence and machine learning (AI/ML), services-oriented software designs, big data analytics, and cloud present opportunities for vendors just like they have for customers to improve their products and services. Solid state storage has come to dominate external storage spend for primary workloads, yet there are few high-end platforms that have been specifically designed and optimized for solid state technologies – particularly the newer ones like NVMe and storage-class memory (SCM). AI/ML, fueled by big data analytics-oriented monitoring that goes far beyond just storage, offers the promise of simplifying the complexity of managing high-end storage infrastructure; yet we have seen these technologies used sparingly for these types of systems. High-end storage operating system (OS) upgrades contribute to this complexity through their still-monolithic designs, while we are seeing most next-generation application (NGA) workloads being developed using a services-based architecture that enables easier, faster, and much more frequent updates with far less risk. Cloud offers agility and flexible consumption options but stands as a distinct storage "silo" that is managed separately with limited visibility into on-premise infrastructure, creating hybrid cloud management complexity.

The maturity of these technologies presents an opportunity to design a high-end storage platform that offers cloud-like simplicity and yet still can meet all the other high-end storage requirements in the areas of performance, scalability, availability, and functionality. With its new HPE Primera high-end storage platform, Hewlett Packard Enterprise (HPE) has done just that. HPE Primera is a modern high-end platform that brings the cloud experience on-premise while delivering the performance, availability, and scalability needed for mission-critical workloads.

## IN THIS WHITE PAPER

---

As enterprises undergo DX and modernize their existing IT infrastructure, a new set of requirements for high-end storage infrastructure is emerging. Driven by heightened requirements for performance, scalability, availability, agility, efficiency, and manageability, customers are looking for vendor offerings that differentiate themselves in these areas. Good opportunities exist for vendors that take advantage of newer technologies like NVMe-based solid state storage, AI/ML, services-oriented software designs, big data analytics, and cloud to address the evolving requirements. This white paper discusses how market dynamics in the era of DX are driving new high-end storage platform requirements and how these newer technologies can be harnessed in high-end storage system designs to meet evolving business needs while freeing up IT resources to perform more innovative and strategic tasks. It then turns to a discussion of HPE Primera, a new class of storage that leverages these newer technologies and promises to change the definition of high-end storage for the industry.

## SITUATION OVERVIEW

---

Over the past several years, the IT industry has been undergoing significant and disruptive change. The big data era is providing significant opportunity for enterprises to better inform business decisions by leveraging big data analytics and AI/ML technologies. As a result, successful businesses have become much more responsive to evolving market conditions, and their underlying IT infrastructures have been forced to become much more agile as well. Next-generation applications (scale-out workloads built around mobile computing, social media, big data analytics, and cloud technologies) are primarily the platforms delivering this agility, but businesses themselves are putting significant demands on storage infrastructure for better performance, higher availability, increased scalability, improved efficiencies, and easier manageability.

These requirements are driving the infusion of technologies that represent a noticeable departure from how and where legacy storage infrastructures were built. Those enterprises that have been able to successfully digitally transform their businesses to more data-centric models are significantly more likely to invest in newer storage platforms like all-flash arrays (AFAs), software-defined storage, and converged infrastructure (CI). In deploying these more modernized IT infrastructure solutions, these same enterprises are also taking advantage of other related technologies, including:

- Flash and other solid state technologies like SCM, along with system architectures and storage OSs that are optimized for NVMe and other forward-looking technologies
- The combination of big data analytics and AI/ML, harnessed in conjunction with automation, that helps offload routine administrative tasks from operational personnel, not only to free them up to focus on more innovative and strategic tasks but also to improve the speed and reliability of operations in general
- Cloud-based services to provide increased agility in IT operations, help offload routine IT infrastructure management to third parties for non-mission-critical workloads, and shift IT assets off balance sheet for more flexible consumption model options

The availability of cloud-based services in particular is driving significant change in IT organizations. It has changed expectations about deployment time frames, made it easier to securely share data across a variety of different constituencies, and introduced IT management to new ways of thinking about easy scalability, frequent and nondisruptive upgrades, and technology refresh for IT infrastructure in general. With end users becoming more used to this "as a service" experience, many

IT organizations are striving to approximate this model for on-demand, always-on, and easy-to-use application services for noncloud infrastructure as well.

## Today's High-End Storage Experience

High-end enterprise storage systems scale to millions of IOPS and petabytes (PB) of capacity, deliver "six-nines" of availability, and offer a full complement of storage management or "data" services that enable dense mixed workload consolidation while still meeting individual application requirements. Granular, multitenant management is a clear strength for these platforms. Despite the advantages, there are a number of challenges to be overcome with these high-end storage systems:

- **Don't address application service availability.** Even if a storage system is 100% available, that does not necessarily translate into application service availability, and that is the real "product" that IT delivers to its constituents. Given that storage only causes 9.8% of the downtime in a typical datacenter, IT clearly needs visibility into the other layers in the stack that deliver the application services, and when administrators must use many tools to accurately diagnose and remediate issues, they present more complex challenges and take longer to resolve.
- **Designed for experts.** The "flip side" to management flexibility is complexity, and these systems often require very sophisticated storage management expertise to be able to install, configure, and manage them to meet business objectives over time, but with storage management tasks migrating away from dedicated storage administration teams toward IT generalists, this knowledge can be lacking.
- **Require constant tuning to address "noisy neighbor" concerns.** Application changes, which happen on a regular basis as new applications are added or capacity usage grows, raise concerns about the ability to predictably meet performance requirements across a set of mixed workloads. If systems do not offer sufficient capabilities to meet service-level agreements, then manual tuning is required to rebalance the system configuration, and this retuning is generally time consuming and requires relatively sophisticated storage knowledge that many administrators managing storage today don't have.
- **Impose disruptive, time-consuming, and/or risky upgrades.** Technology upgrades to these systems can be disruptive, time consuming, and fraught with risk. Moving to the next technology generation still generally requires a forklift upgrade, manual data migration (which not only can take a long time but also often impacts application performance during the migration process), and incremental upgrades to firmware (which can often be performed online, but the monolithic storage OS designs require the entire operating system to be updated each time any kind of software update is performed), and pre-validating and forecasting the "impact" of these types of upgrades depend on human involvement – both of which add time and risk.
- **Have outdated technology "baggage."** System architectures originally designed during the hard disk drive (HDD) era will not necessarily support the massive parallelization necessary to optimize the performance and capacity utilization of the solid state media that has already come to dominate high-end systems shipments, and this factor can hinder workload consolidation and limit a system's ability to deliver predictably consistent performance.

Each of these topics helps foster the perception that high-end storage systems are harder to manage, creating concerns not only about administrative productivity but also about sufficiently skilled resources to manage them. IDC has noted over the past five years that storage administration responsibilities are changing, with IT generalists (e.g., virtual, Windows or Linux administrators) taking over more of the storage management tasks. Low-end and midrange storage system vendors have focused particularly on making storage management easier, but this is primarily being done through

the use of wizards and at the expense of the type of administrative flexibility available with high-end systems. With new technologies becoming available, there is a significant opportunity to design a new high-end storage array to address management complexity concerns while still meeting all the other high-end storage requirements.

## Toward the Modernized Infrastructure Definition of High-End Storage

With agility as a key requirement for today's digitally transforming enterprises, the design of high-end storage systems needs to change. The more purpose-built, statically defined scale-up designs of the past need to give way to more software-defined, modular scale-out designs. Systems that were originally designed and optimized assuming the use of HDDs need to be optimized to take advantage of the latest solid state storage technologies (NVMe and NVMe over Fabric). These systems need to leverage AI/ML technologies, in combination with big data analytics and automation, to improve and simplify routine management tasks at the scale of today's computing environments. They also need to be designed to operate in hybrid cloud environments, since this is clearly the blueprint for how IT infrastructure will be deployed going forward. And they need to offer all this while continuing to deliver the rich storage management functionality that enables secure, multitenant management from a proven storage OS that can scale to millions of IOPS, petabytes of capacity, and "six-nines plus" availability.

Other design considerations will play into this new definition as well. The customer experience based around broad public cloud usage offers many features that are of interest across many different workload types, but public cloud is not appropriate for all applications. As new workloads get deployed, IT organizations must always consider the best deployment location (cloud or noncloud). For certain workloads, performance, availability, security, regulatory/compliance, and data access considerations can drive an on-premise deployment decision. This does not mean, however, that the increased agility; always-on availability; simple, secure sharing of application services; easy scalability; nondisruptive multigenerational technology upgrades; and "pay as you go" pricing would not also be attractive for on-premise "owned" infrastructure. High-end storage solutions that deliver true enterprise-class performance, availability, and security, along with the enterprise-class data services of a mature, proven storage OS, and that also offer the agility features of the "on-demand" experience are an excellent fit for the requirements driving infrastructure modernization in most companies.

In digitally transformed organizations, the direct contribution of the IT organization to drive bottom-line business benefits is increasing overall availability requirements. Recent IDC primary research indicated that more than 90% of enterprises consider 11-50% of their workloads mission critical, with 70% of companies managing these workloads to at least "four-nines" of availability. 9.4% of enterprises are actually managing their mission-critical workloads to "six-nines" of availability or greater. While public cloud usage is extensive, 93.3% of enterprises prefer to keep their mission-critical workloads in on-premise infrastructure and 83.0% of enterprises deem storage platform availability capabilities a key purchase driver. What this means for customers modernizing their IT infrastructure to host multitenant, on-premise workloads is that storage systems must not only offer a full complement of "availability" capabilities (which include features such as RAID, multipathing with transparent failover, snapshots, online replacement of failed components, nondisruptive upgrades, and various replication options including stretch clusters) but also allow those features to be applied selectively at the application level as needed.

AI/ML and big data analytics are not just finding their way into applications to drive better business decisions, they are also increasingly being harnessed by enterprise storage vendors to automate

systems management and improve the efficiency of IT resource utilization. While hybrid cloud environments provide attractive options for IT organizations, they also introduce management complexities that will be increasingly difficult for humans to navigate. Forward-looking vendors have added significant instrumentation to their systems, collecting sensor and other data that is then analyzed using AI/ML and harnessed to increase performance and availability, predictively identify looming failures, better inform performance and capacity planning, speed problem resolution, and improve the efficiency of system configuration and resource utilization. These types of solutions are often referred to as "cloud-based predictive analytics" platforms, and the best of these work from a global view of infrastructure that includes not just storage but also servers, networks, virtual machines, storage infrastructure software, applications, and public cloud environments, leveraging this much more "aware" view of IT infrastructure to make more optimized and efficient management decisions.

## Resetting the Definition of High-End Enterprise Storage: HPE Primera

HPE is a \$31 billion solutions provider of IT infrastructure, including servers, storage, networking, infrastructure software, and technical support and consulting services. HPE offers significant choice to its customers in terms of deployment models for storage, delivering storage appliances; software-only, hyperconverged infrastructure (HCI); converged infrastructure; 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 the industry's most mature cloud-based predictive analytics platform (InfoSight). HPE was founded in 1939, is a proven trusted supplier to enterprises of all sizes, offers storage platforms that can support "six-nines plus" availability, and is one of the market leaders in all-flash array revenue. AFAs dominate primary external enterprise storage shipments, driving almost 80% of all revenue in that market in 2018.

HPE is a long-established player in enterprise storage, with hundreds of thousands of systems installed over many decades of servicing the high-end storage needs of enterprise customers. In introducing its new high-end storage system, HPE Primera, HPE is melding the simplicity of Nimble Storage, the mission-critical heritage of 3PAR, and the intelligence of InfoSight into a single high-performance, highly available, and highly scalable large enterprise storage solution. HPE has designed a new system architecture around a number of new features and technologies to provide a platform that delivers management simplicity without compromising its ability to support an enterprise's most mission-critical workloads:

- **Scale-out architecture to support massive parallelism.** As a storage protocol technology, NVMe supports at least three orders of magnitude higher parallelism than SCSI, and HPE Primera embraces this with a design that is highly optimized for solid state storage. The ASIC (which is still at the core of HPE's high-end system design) has been re-architected to optimize internode concurrency, and each system supports up to four controllers, each with two Intel Skylake CPUs and up to four ASICs. The ASICs provide zero detect, SHA-256, X/OR, cluster communications, and data movement functions. Other capabilities, such as inline data reduction (compression), will be run in either a QAT chip or a controller CPU, depending on maximum real-time efficiency (as determined by the system's AI/ML-driven self-optimization that is constantly evaluating I/O profiles, available resources, and overall system load). In HPE's new "all-active architecture," all controllers and all cache are active all the time, driving low latency and high throughput. Other features, such as on-disk data protection, have been modified for this highly distributed design; the self-healing erasure-coded data layout varies based on the size of the system and is adjusted in real time for optimum performance and availability. Up to four HPE Primera nodes can be added in a single system.

- **A modular, service-centric design for the storage OS.** The HPE Primera storage OS is not a monolithic software platform. It is instead built around a service-centric design that allows any feature to be added or modified without requiring a recompile of the entire OS. This is the way all NGAs are being designed today because it enables faster, more frequent updates that are easier to implement and significantly less risky to perform. All of the proven, trusted features associated with high-end HPE storage – RAID, thin provisioning, snapshots, quality of service, replication, and so forth – have been implemented as independent services for the Primera storage OS. This approach will allow HPE Primera to be upgraded faster, easily, more frequently, and with less risk than other high-end storage systems.
- **AI/ML-driven self-optimizing systems management.** HPE Primera features very comprehensive and granular monitoring, all the data from this is uploaded to HPE InfoSight (part of the HPE Intelligent Data Platform). HPE InfoSight uses AI/ML to evaluate system performance, capacity utilization, energy consumption, and other metrics and combines that data with data collected from other "objects" such as servers, networks, virtual machines, storage infrastructure software, applications, and cloud storage to generate a truly global view of a system or workload state at any point in time. The more global view provides a better overall opportunity to optimize to specific objectives and/or troubleshoot infrastructure issues. Certain performance optimization and other real-time management decisions are done by the system itself through the array's own AI/ML engine (e.g., any performance tuning efforts that may have to be done after more storage capacity or a new workload is added) and based on policies set up by administrators, while for others it makes recommendations for suggested changes (when administrators want to ensure a human is involved).
- **Streamlined multinode hardware design.** Internode communication requires no separate cables – all nodes plug into a passive backplane for much simpler configuration and higher reliability. The HPE Primera node-level building block is a high-density design that can house up to 1PB of effective capacity in 2U (or 2PB in 4U), with additional external storage capacity expansion available in both form factors. Performance and capacity density have been significantly increased (e.g., the throughput and storage capacity that HPE used to deliver in a 19U system can now be delivered in a system that takes up only 4U). The HPE Primera building block is available in two sizes – 2U24 (with two controllers) and 4U48 (with four controllers) – and includes eight dual-purpose (SAS/NVMe) disk slots per controller pair. Each controller node includes up to 12 host ports per node (48 total) with 25GbE or 32Gb FC connectivity and redundant, hot-pluggable controllers, disk devices, and power and cooling modules. This building block node is designed as the hardware foundation for storage systems that can deliver 100% availability. In fact, as part of the HPE Primera announcement, HPE introduced a 100% availability guarantee on every single Primera system that does not require stretch cluster configurations or even a maintenance contract.

## How HPE Primera Changes the Customer Experience of High-End Storage

Leveraging the previously mentioned design tenets and technologies, HPE Primera focuses on three areas that will change how customers view high-end storage. The first that administrators will notice is the "on-demand" experience HPE Primera delivers for on-premise storage infrastructure for everything from deployment to provisioning storage, expanding a system, upgrading it, and paying for it. Initial systems can be installed in as little as 20 minutes, systems can be expanded nondisruptively in 10 minutes, and they support data in-place upgrades to next-generation controllers and storage devices with new media types (as part of the HPE Timeless Storage program).

The system generally follows an "API first" management strategy, with prebuilt automation for VMware vCenter, Virtual Volumes, and the vRealize Suite. Storage can be provisioned in literally seconds, and

storage OS updates (which are nondisruptive as well) require only a single click and take on the order of five minutes (and do not require a reboot). System updates are all pre-validated prior to installation by looking at configurations across the entire installed base (using HPE InfoSight) to identify predictive signatures for that particular update to minimize deployment risk. Under HPE GreenLake, any of HPE's storage platforms can be purchased through a pay-per-use subscription model that is fully International Accounting Standards Board (IASB) 2019 compliant (to ensure that infrastructure remains off balance sheet).

The second is how the HPE infrastructure stack is architected to deliver *application-aware resiliency*. All HPE Primera systems include Recovery Manager Central (RMC), a software solution that provides application-aware data protection for Oracle, SAP, SQL Server, and virtual machines. RMC provides a number of options for cost-effective backup/disaster recovery data storage and rapid recovery with out-of-the box integration for HPE StoreOnce, HPE Nimble Hybrid, and the public cloud (through HPE Cloud Bank Storage).

Application-aware predictive analytics, based on AI/ML-driven global learning, which leverages HPE InfoSight, keep systems operating reliably and at peak efficiency, even as workloads evolve and system configurations change over time. This platform collects data not just from the storage (storage on average only causes 9.8% of measured system downtime across the industry) but also from other layers, including servers, networks, virtual machines, storage infrastructure software, applications, and cloud storage – all of which are leveraged by the Global Intelligence Engine to identify problems before they can occur and troubleshoot and rapidly resolve those problems that do arise. Cross-stack analytics can identify noisy neighbors, perform host and memory analytics, troubleshoot application latency issues, and identify inactive as well as top-performing virtual machines. After an issue has occurred once in the installed base, AI/ML-driven signature matching is performed to predictively identify any other systems that may be at risk for that incident.

In addition, system features like host multipathing, dynamic erasure coding, transparent failover, hot-pluggable field-replaceable components, snapshots, replication, and stretch clusters (which HPE refers to as Peer Persistence) provide flexible "defense in-depth" strategies for cost effectively configuring workloads to meet individual high-availability and application-specific resiliency requirements. HPE also now supports secure, cloud-based quorum witnesses, which simplify and lower the cost of Peer Persistence configurations for those customers that require stretch clusters.

The third is how HPE Primera makes it easy to apply *predictive acceleration* to ensure that workloads meet the SLAs (performance and availability) of the classes in which they are placed. The embedded AI/ML automatically predicts and optimizes performance based on workload needs, even as the overall load on the system is evolving because of I/O spikes, component failures, added resources, data growth, or new workloads. Data is collected and analyzed in real time, anomalies and hot spots are detected, saturation impacts are evaluated, performance is predicted, and the system undertakes remediation to ensure that SLAs continue to be met. The all-active architecture makes the system resistant to hot spots regardless of I/O load, thereby enabling it to deliver predictably consistent latencies even as I/O load is scaled and despite variances in demand. The extreme parallelism of which the all-active architecture is capable allows HPE Primera to get much better utilization out of low-latency solid state technologies like NVMe, NVMe over Fabric, and SCM. To support the much higher degrees of concurrency of which solid state media is capable, the system employs a new "lockless" data integrity mechanism, a new feature in HPE's Generation 6 ASIC.

The HPE Primera system architecture demonstrates an intelligent approach to leveraging newer technologies and concepts in storage infrastructure to create highly scalable, highly available systems that can largely self-manage themselves to policies established by administrators. For those less comfortable with such levels of automation, HPE Primera offers the option to use the system to generate recommendations, which must be approved by a human before they are implemented. With the advent of big data analytics and AI/ML, many types of systems across all industries that used to require close human monitoring and management will be moving toward more automated operations. The new HPE Primera system is already built to ease customers into this way of managing systems at their own pace.

## CHALLENGES/OPPORTUNITIES

---

A few other storage vendors are leveraging AI/ML engines to help manage and optimize their systems, but none have made the commitment to the same extent that HPE has with HPE Primera. HPE owns the cloud-based predictive analytics platform that started all this back in 2010 – InfoSight that was obtained through the 2016 Nimble Storage acquisition – and the company has more experience driving value for customers using AI/ML-driven big data analytics than any other vendor. There may be a reticence on the part of some customers to delegate routine tuning and other administrative operations to an AI/ML-based engine because they are used to a more "hands-on" approach to storage management, but primary research on this topic by IDC in late 2018 paints a different picture. 73.8% of survey respondents indicated that they were interested or very interested in leveraging autonomous operations to streamline storage management and 85.7% of survey respondents highly valued the use of AI/ML in this capacity. IDC believes that with IT administration becoming more complex in this era of hybrid cloud, the direction taken by HPE with HPE Primera will become a foundation requirement for enterprise storage within two to three years. In differentiating HPE Primera from the competition, one of the challenges HPE will face is to make clear to customers and prospective customers the nature of the value that the company is driving in terms of efficiency of operation, easier management, and administrative cost savings with this self-optimizing, self-managing system.

With HPE Primera, HPE has very consciously gone after several key objectives. The first is to design a high-end storage system that takes full advantage of the performance and capacity of newer solid-state storage technologies – in particular by designing an architecture that supports massive parallelism. The second is to transcend the capabilities of humans to optimally manage such a system by combining comprehensive cross-stack monitoring, big data analytics, an AI/ML-driven Global Intelligence Engine, real-time self-optimization, and automation to simplify management without giving up performance, availability and flexibility. And third is to create more of an "on demand" experience for on-premise infrastructure for everything from deployment to on-demand provisioning; system expansion; low-touch, real-time policy-based management to SLA goals; nondisruptive (and potentially more frequent) upgrades; integration with multicloud environments; and IASB 2019-compliant subscription pricing options. This clearly provides an opportunity for HPE to explain just how HPE Primera is meaningfully differentiated from the competition.

## CONCLUSION

---

Enterprises that are successfully navigating the DX journey are those that will be around in the future. As part of their DX journey, these companies are all modernizing their existing IT infrastructure to meet evolving requirements for better performance, higher availability, increased scalability and agility, and



improved efficiencies in energy and floor space consumption and management. Newer storage technologies like NVMe, SCM, big data analytics, AI/ML, services-based software designs, and cloud offer significant opportunities for vendors to improve their products and services in ways that drive meaningful differentiation for their customers. Traditional array designs struggle to meet the storage infrastructure requirements of digitally transformed enterprises, and the needs of IT infrastructure modernization will drive the emergence of a new class of high-end storage platform that leverages these technologies.

With its introduction of HPE Primera, HPE has designed a platform to deliver simplicity while at the same time meeting all the other high-end storage requirements in the areas of performance, availability, scalability, and multitenant management capabilities. With HPE Primera, HPE has intelligently integrated solid state storage technologies, big data analytics-driven AI/ML on a global level, a services-oriented storage OS, and a streamlined hardware design with a scale-out architecture capable of massive parallelism to produce a high-end storage platform that dynamically self-optimizes to meet established performance, availability, efficiency, and cost objectives. HPE is the first established enterprise storage vendor to leverage these technologies to this extent, and it has done this in a way that will set a new customer experience bar for high-end storage. IDC expects this design approach to drive significant value for customers and predicts that we will see more enterprise storage vendors moving in this direction with their solutions as well over the next two to three years.

## 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.

