



TECHNOLOGY SPOTLIGHT

Storage Solutions for the Transforming Business

July 2017

By Cynthia Ho and Avneesh Saxena

Sponsored by Lenovo

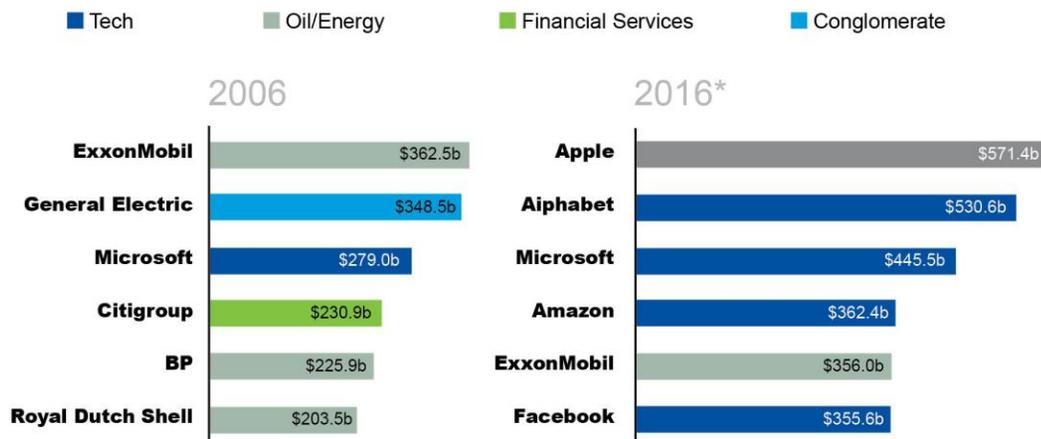
In today's digital era, organizations that have storage and data management as an integral part of their overall enterprise infrastructure strategy stand in good stead to keep pace with the rapid growth of data. This IDC Technology Spotlight takes a closer look at the impact of digital transformation on the enterprise from a storage management standpoint, and discusses the role of Lenovo in this strategically important market.

The Reality of the Digital Economy

The Economist writes that data, and not oil, is the most valuable resource today. Thanks to the digitally savvy world we live in now we are inadvertently generating a digital footprint or shadow of ourselves every second we transact online, post a comment on social media, watch content online, or simply interact with an online service just to get through our day. There is always data generated for humans today, and machines tomorrow (i.e., artificial intelligence [AI]), to analyze patterns of behavior and predict what next. This is the holy grail of data value and monetization, and will continue to see huge investments not only from disruptive born-in-the-cloud companies but also traditional companies like GE, Siemens, Airbus et al that are trying to become data companies. It's interesting to see the validation of the above comments, shown in Figure 1, which compares market capitalization of the most valuable companies in 2006 and 2016. We literally see the big oil companies knocked out of the top club by the new data giants!

Figure 1

Market Capitalization of the World's Most Valuable Public Companies



* as of August 1, 2016

Sources: Yahoo! Finance, Forbes, Statista

Agility is Driving Transformation

There are many ways to look at the impact of digital transformation on IT, but prime among them is the ability to build more *agile IT*. This is a philosophy that cuts through all forms of IT, and builds on the success of leading enterprises which have continuously introduced technology based services and failed fast to taste success. Netflix is among the leading examples of companies that practice the fail-fast-and-succeed mantra. But this requires a trinity of groups that work in tandem:

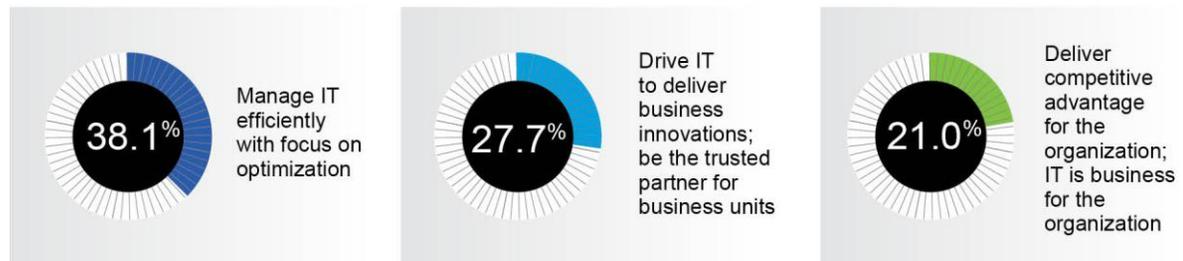
- **Innovation of ideas** that deliver new digital experiences to existing and new customers to drive growth;
- **A DevOps team** that converts those ideas into digital software and services that are readily deployable; and
- **Agile infrastructure** that helps deliver and analyze the impact of services rendered.

Innovation is like breathing for most enterprises today, and they understand the need to build the practice of deep breathing to stay healthier. The more evolved are embarking on building digital teams to transform faster and serve the unmet needs of their customers. These teams comprise business and IT people, and endeavor to help convert business ideas into digital services. IT has been the stumbling block and the main driver for line of business (LOB) executives to consider going straight to cloud in the early stages. Legacy IT was built to serve the unique needs of different business groups, and on the way, it became extremely rigid in the cross-compatibility. This is among the biggest challenges for a CIO who is trying to deliver to new expectations, while keeping the status quo on the existing IT infrastructure. Figure 2 illustrates the key priorities of a CIO that are built off the business expectations to have more streamlined IT systems. The highest rated factor mirrors the frustration of the CIO who is expected to streamline the IT resources to move faster and address the next two priorities on his/her list.

Figure 2

Top 3 CIO Priorities

Percentage of respondents



N=667 Asia/Pacific excluding Japan

Source: IDC's C-Suite Barometer Survey 2016-17

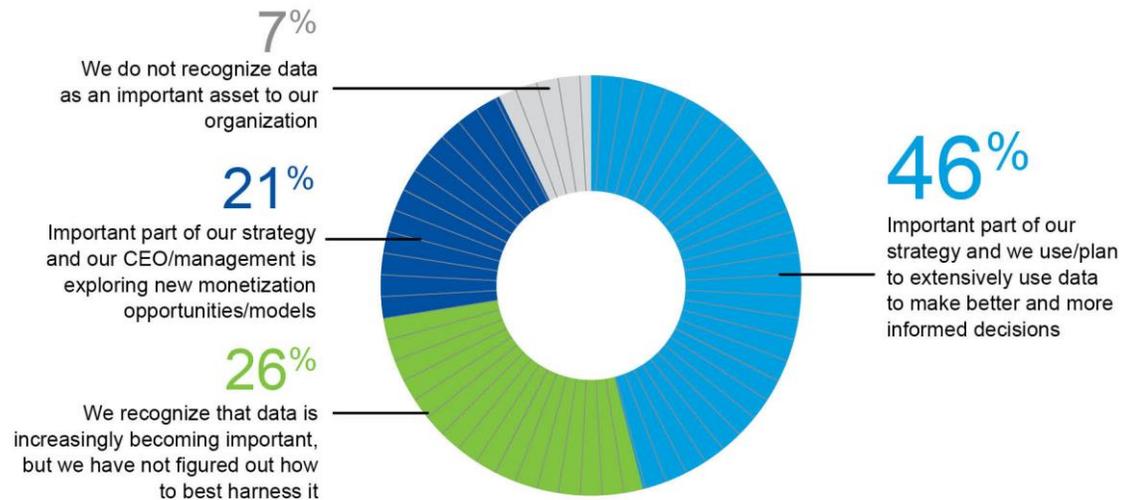
Mining the Gold

There is no denying that data is a strategic asset, as reflected in IDC's annual C-Suite Barometer survey (see Figure 3). However, managing data resources effectively is another matter. The islands of legacy IT within an enterprise stand guard on some extremely valuable data, and most digital endeavors struggle to connect the new services to the old wisdom. To bridge this proverbial chasm between legacy and next-generation IT, CIOs need to avoid finding stop-gap solutions, like compartmentalization of old and new, since that will only increase their pains in the longer run.

Figure 3

How CXOs View Data in the Enterprise

Q: Is data a strategic asset for your CEO and organization?



N=1,472 Asia/Pacific excluding Japan

Source: IDC's C-Suite Barometer Survey 2016-2017

IDC believes that an organization needs to consider embarking on an introspective journey that helps identify and resolve where the data should stay in order to have more seamless access to data. To do this, an organization will need to break its requirements by the type of data they are planning for, and its current affinity to the enterprise workloads that drive that data. For instance, the unstructured steams of data have been the Achilles heel for most administrators, who have failed to predict its impact and hence unable to plan the requisite storage solutions. In the interim, they have reached out to public cloud since the organization has failed to cope with the rise in unstructured steams of data.

Datacenters have traditionally been quite siloed, and not ready to handle modern age workloads. Traditional systems of record, such as enterprise resource planning (ERP) systems are being outpaced by digital transformation workloads: systems of engagement, insight and action. The nature of workloads has also changed tremendously over the year. To support such workloads, the datacenter needs to transform through five stages of maturity: archaic, static, transformation, flexible and finally, dynamic. IDC also noticed the increased awareness in cloud adoption, which has changed the way enterprises deliver IT infrastructures. Organizations won't be able to digitally transform their businesses unless their datacenters and cloud environment maturities are parallel. The underlying IT and facility infrastructure will also go through five distinct stages in parallel: consolidation, virtualization, orchestration and federation. Critically, an enterprise's digital transformation efforts need to mature and evolve at the same rate as its underlying cloud, datacenter, IT and facility infrastructure.

IDC suggests reviewing the data strategy by taking a deeper look at how the data should be stored based on the three dimensions of data access for business benefit, data integration to new apps and monetizable value of the data to the business.

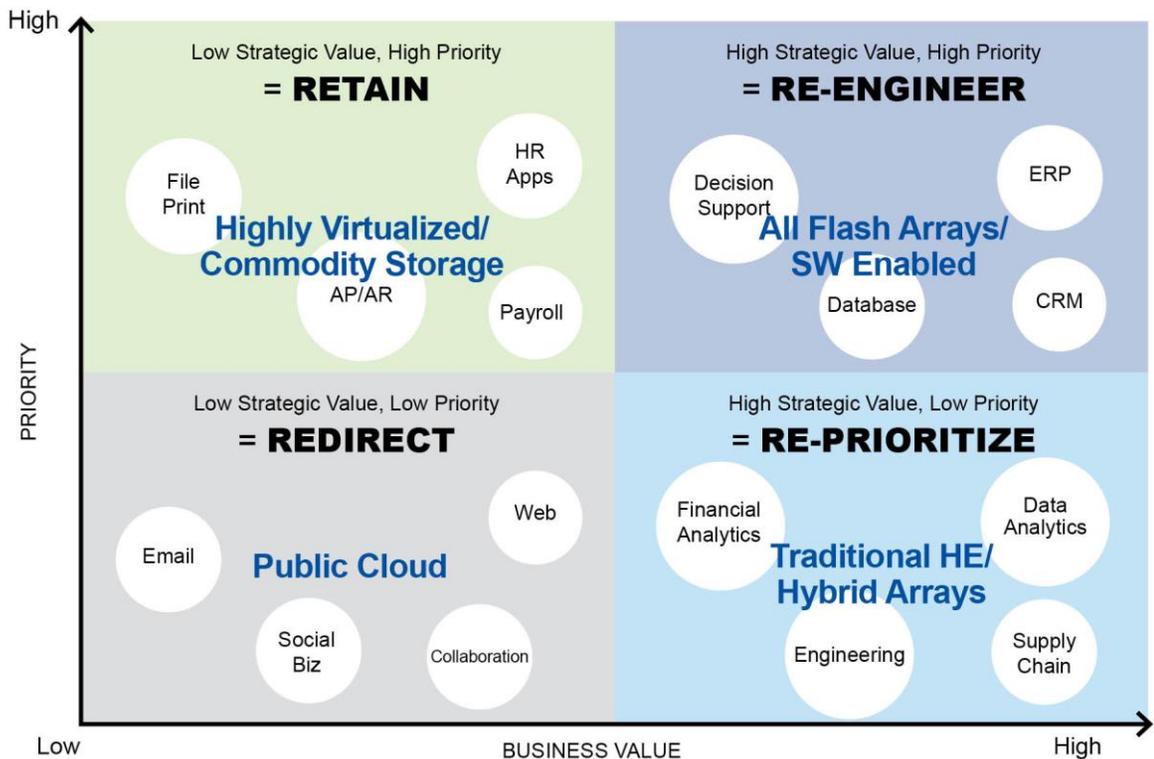
Store-Right: One Size Doesn't Fit All

There is no one solution that fits the different needs of future enterprises. Most of the enterprise workloads today consist of a mix of legacy and next-generation applications. The 3rd Platform workloads exploit the use of mobile, social media, big data/analytics, and cloud which are very different from traditional workloads that are more inward-looking. Legacy workloads run primarily on tailored infrastructure that is designed to offer the best performance for a specific workload, whereas next-generation applications work well in a virtual and shared infrastructure environment. The shift also creates an I/O profile that is very different from what enterprises used to have 10 years ago. Enterprises' I/O profile today are more dynamic and characterized by a variety of read/write ratios, a wide distribution of block sizes, a heavy skew toward random (rather than serial) I/O, significantly less amount of data redundancy (e.g., data de-dupe), and "hot spots" (I/O bands that drift over time). Another key factor is having a policy to deploy storage that adheres to open standards which help in easier accessibility and integration of data. Open standards go a long way in unifying the underlying infrastructure, and avoiding silos that are otherwise inadvertently created due to non-communicating protocols. Lastly, the idea of using cloud as a service changes enterprises' expectations, and enterprises are working to bring cloud economics into the datacenter and bridge on-premises and cloud storage solutions.

To select the right storage solution that addresses business goals, enterprises will need to understand their workload criticalness and characteristics. Traditional disk technologies are not suitable in handling highly random I/O, hence the need for flash technologies to satisfy the requirement. Besides, workloads that require lower latency give rise to the possibility of leveraging inline storage efficiency. Of course, business value, business priority, and financial goals should be considered concurrently in order to come up with a suitable platform choice that suits the whole organization. Enterprises can consider four different guidelines to determine the right storage platform for their organization. The guideline in Figure 4 is based on business priority and business value.

Figure 4

Application Portfolio Drives Storage Platform Choice



Source: IDC, 2017

- **High strategic business value/High business priority:** Tier 1 critical workloads which are extremely time sensitive and require analysis on high volatile information to perform decision making should be re-engineered to all flash array or software enable platform. Workloads example such as: DB, CRM and ERP will fail the business operation if speed is not up to optimal level.
- **High strategic business value/Low business priority:** Tier 2 workloads support many major business applications, such as financial analytic, supply chain and data analytic, require active business data to be securely stored. If end users don't require subsecond response and can tolerate slightly slower response, they should re-prioritize using high-end storage or hybrid array.
- **Low strategic business value/High business priority:** Tier 3 or tier 4 workloads are mostly traditional applications which are not critical enough to stay on tier 1 or tier 2 system. As time goes by, enterprises do not need to access the data all the time and the reads drop off rapidly. Workloads such as file print, payroll, HR apps should be retained on using virtualized storage or commodity storage.
- **Low strategic business value/Low business priority:** Tier 4 workloads are not critical enough but need to be on backup require such as email, social and Web should be redirected to public cloud.

Getting Future Proofed

Enterprises have been focusing their IT investment in infrastructure consolidation, optimization and improving operational efficiency. In the next 2-3 years, enterprises need to prepare another major shift in workloads while expanding their datacenters. According to IDC's FutureScape prediction, by 2019, 25% of organizations' datacenter investments will be supporting next-generation contextual workloads such as cognitive/AI, machine learning, and augmented reality. Many of these next-generation applications deliver context to enhance users experience. Enterprises are hiring coders to build cool applications to engage with their customers which can be an effective tool to drive more profit for their business. Data will eventually be generated in cloud, edge and traditional environments, illustrated in Figure 5.

Figure 5

New Applications Driving Up Storage Requirements

Exa-scale storage	Cloud databases	Edge workloads	Next-gen apps
<ul style="list-style-type: none"> • Geo-distributed key/value based • Software-defined • Runs on user-configured platforms 	<ul style="list-style-type: none"> • No data conformity required • Key/value based • Analytics friendly • Runs on any kind of core/edge compute 	<ul style="list-style-type: none"> • Compute footprint distributed at the "edge" and "core" • Compute units follow a "best fit" form factor • Runs on platforms "embedded" inside "edge" devices 	<ul style="list-style-type: none"> • Stateless/horizontally scalable • Don't assume infrastructure resiliency • New development processes • User experience

New workloads require more power capability to support higher density and capacity in both compute and storage. Enterprises cannot rely on traditional infrastructures to help them scale up and out quickly, while facing stringent budgets internally. We are starting to see datacenter service and workloads failure caused by hardware obsolescence. IDC is predicting that around 30% of enterprises in the Asia/Pacific excluding Japan (APeJ) region will suffer datacenter service failure in the next two years due to mismatches in power delivery and IT workloads, and this will threaten business competitiveness. In a recent IDC survey to understand current datacenter issues, enterprises in APeJ said that when they experienced datacenter downtime due to system failure and human error, latency and performance issues dampened their productivity level. Their existing compute and storage cannot keep up with the increasingly complex technical environment and workloads. Thus, infrastructures and facilities cannot deliver performance optimally due to the mismatch in power and latency. To address the mismatches or gaps in datacenter capability, enterprises will need to find a new balance in their strategies. To build smarter forward-looking datacenters enterprises will need to implement next-generation hardware and software solutions to build a more dynamic "thinking and acting" datacenter. Based on IDC's Asia/Pacific latest enterprise and datacenter survey of 100 respondents, the top 3 enterprise IT initiatives include modernizing core apps, investing in automation and management tools, and adopting software-defined storage solutions.

IDC recommends a framework to help CIOs and IT heads ask the relevant questions as they decide on their enterprise transformation needs. Enterprises should begin asking these questions, as shown in Figure 6.

Figure 6

Key Questions to Help Identify Enterprise Transformation Needs

Agility and scale according to business	Data availability, risk and reliability	Infrastructure openness
<ul style="list-style-type: none"> • Do you have the resources to run your infrastructure to support 10 times more customers than today? • Can your infrastructure ramp up to meet your digital expansion needs? • How is your data stored today, and are you optimizing cost by tiering it, building easy access for business and ensuring it scales on demand? • Are your internal IT resources struggling to address daily network infrastructure and operational issues? 	<ul style="list-style-type: none"> • Is your backup and archival system tuned to seamlessly store and retrieve data? • How reliable is your tiered infrastructure in securing your data from risks ranging from human error to threats? • Have you invested in the most reliable and resilient storage solution, or the cheapest technology that creates islands of disparity in your datacenter? 	<ul style="list-style-type: none"> • How much visibility should you allow and yet maintain high security without impacting innovation? • How do you reduce the cost of products and services in the process? • How do you stay compliant and maintain your speed-to-market competitive advantage as digital transformation continues to evolve?

Source: IDC, 2017

A word on the resting place of the workloads as there seems to be an almost religious debate on choosing between a public or private cloud framework: IDC believes that the natural transition of the workload will lean towards a hybrid model that includes non-cloud, private and public deployments. Hybrid infrastructure sourcing and management capabilities are critical for an agile IT environment

According to IDC’s FutureScape prediction, by 2018, 45% of APeJ enterprises ICT spend will be in a mix of colocation, hosted cloud, and public cloud datacenters. This indicates that there is rarely a one-solution-fits-all scenario to address the pressing issues of today’s increasingly complex IT infrastructure.

Lenovo’s Value Proposition

As data becomes more critical for an organization, and more data is created every second, there seems to be an overt urgency to find the quickest way to store data. Cloud and virtualization seem like a panacea of the rising data problem, as you pay on-demand from an opex budget, or scale on demand from available storage resources. Not surprisingly the traditional hardware vendors have felt the pressure of this trend on their bottom lines and are fast finding ways to stay relevant in the changing landscape. Adding to their woes are new disruptive storage solution providers that are building on a software-defined and hyperconverged approach that sounds both exciting and mould-breaking from the traditional siloed structure.

Lenovo, realizing this changing landscape, formed the Lenovo Datacenter Group (DCG) in April 2017 with the intention of providing different software-defined solutions through unique partnerships. This addition, at the back of the IBM acquired x86 server and networking business, helps Lenovo become an end-to-end infrastructure provider, but with no bindings from existing storage revenue chains. The latter dynamics helps Lenovo in remaining unhinged to agendas that are less relevant for their

customers. Lenovo's SDS offerings come in two flavors: the appliance model and the reference architecture model.

The reference architecture model centers around validating the configuration required to run the partner software. However, the appliance model is a bit more tuned-to-task with both Lenovo and their partners invested heavily from ideation to manufacturing to support in bringing the product to market. Customers will see this as a seamless entity as all customer problems are jointly dealt with. Most of the time, Lenovo spends time with the partner doing a root cause analysis (RCA) to fix the current problem and prevent future ones. Below are the specific SDS solutions that Lenovo has released in the market:

- In October 2016, Lenovo announced appliances with Cloudian to address the growing object storage market, called the DX8200C. While Cloudian had partnerships with both Google and Amazon, Lenovo became the only vendor to offer onsite object storage capability. With Cloudian today, Lenovo can offer data capacity starting at 100TB scaling up to Exabytes of data, with the ability to burst to the cloud as well, that is also 100% compatible with S3 API. This solution comes more into the conversation of management of large size data, and accessibility.
- In that same month, Lenovo announced partnerships with Nexenta on File & Block storage with DX8200N appliances that run the NexentaStor software. The DX8200N delivers unified file (NFS and SMB) and block (FC and iSCSI) storage services, scaling from tens of terabytes to petabyte configurations, and includes all data management functionality by default. This solution is more into the conversation of \$/TB, perfect for large backup, and cloud hosters.
- In March 2017, Lenovo introduced DX8200D based on Datacore. The #1 & #2 SPC-1 benchmark on storage today is held by Datacore running on Lenovo Servers. Both companies have seen many successes in the field, and it is the next logical step in the relationship. Due to its storage virtualization capability, the DX8200D protects existing investments on traditional SAN, and makes it easy to move from traditional SAN to software-defined storage. This solution is more into the conversation of utilizing existing storage assets, and expanding to enterprise capabilities at incremental costs.
- Lastly, Lenovo has also introduced the DSS-G appliance that leverages IBM Spectrum Scale (erstwhile GPFS), which uses massive parallelism in HPC environments, to bring capacity and performance based storage systems. This parallel architecture enables the system to scale into petabytes yet deliver very low latency, high performance, at practical price points.

Lenovo hasn't stopped innovating. The computer maker has also introduced the D3284, a 5U JBOD that supports 84x3.5-inch disks, enabling even higher density solutions for high capacity solutions like Cloudian, Nexenta, and DSS-G. With OpenStack support from Lenovo XClarity, Lenovo is constantly working to bring together an open yet integrated appliance that adapts to enterprises' needs. Lenovo realizes that not all SDS solutions fit all kinds of workloads today. There are several criteria such as data type, performance, latency, capacity, access to open API, hybrid cloud, TCA/TCO and scaling needs to be assessed before selecting the right storage platform.

Table 1 illustrates the workload characteristic and the right SDS within Lenovo portfolio to be adopted.

Table 1					
Lenovo SDS Models and Capabilities					
Capability	DX8200C	DX8200D	DX8200N	DSS-G	DS6200
Data type	Structured & unstructured	Structured, block data	Structured & unstructured, but mostly used for structured	Object, block, structured, unstructured data	Structured data, traditional SAN
Performance	Low performance	Very high performance	Low performance	Very high performance	Very high performance
Latency	High latency	Very low latency	High latency	Very low latency	Very low latency
Capacity	Very high capacity, into 10s PB	from 20TB to 100s TB (block)	High capacity, up to 5PB in 1 cluster	In PBs, very large	Up to 50TB for performance, up to 300TB for capacity
Access to open APIs	100% S3 Compatible	Open Stack support	Open Stack support	Limited	Yes, with Xclarity Open Stack integration
Hybrid cloud	Yes	No (no backup to cloud)	Yes	No	No
Total cost of acquisition (TCA)	Very low	Low	Very low	Mid	Very low
Total cost of ownership (TCO)	Very low	Low	Very low	Mid	Very low
Scaling type	Scale out	Scale out	Scale up	Scale out	Scale up

Source: Lenovo, 2017

IDC's Analysis

If the last decade was the compute era, then this is the data era. Compute cycles and processing capability are secondary since applications are built differently today. Besides, scale has been created through virtualization and cloud, and more time is being spent now in defining how to optimally build and run business outcome focused IT. On the other hand, data is recognized as an important strategic asset, and organizations are vying to become data companies to improve products and services and build monetizable advisory intelligence. This means significant thinking will go into acquiring, storing, analyzing and acting on data collected.

Storage is of paramount importance as enterprises realize they need agile and adaptive infrastructure to store right (cost and scale), have easy access to information (integrated databases), and connect to different systems/devices (open APIs). IDC has observed that enterprises still want to work with bigger vendors for business workloads since they care about availability, resilience and support. Lenovo seems to have recognized this critical need and, without any protecting-the-legacy-installed-base distraction, strategizes to enter this space with focused partnerships aimed at solving important storage issues. While it may be a little late in entering this crowded space, IDC feels Lenovo has what it takes to be a strong storage vendor, provided they focus on the following three aspects:

- **Focus on execution and support:** Differentiation through successful execution and after-sales support to improve customer experience which will help build the value proposition they aim for.

- **Re-establish enterprise infrastructure leadership:** Lenovo needs to re-connect with the IBM x86 installed base and re-assert its leadership in the enterprise infrastructure space. To regain its momentum outside China, it needs to quickly build credibility by citing customer success in the home market, particularly where it has been busy working to transform some large enterprise customers.
- **Build a strong partner network:** Apart from customers, Lenovo has to re-engage with partners that will be instrumental in driving the scale that it is after and taking these solutions to its customers. IDC believes that partners may in fact welcome the opportunity as it offers them ways to sell storage-solution contracts.

A B O U T T H I S P U B L I C A T I O N

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

C O P Y R I G H T A N D R E S T R I C T I O N S

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the Custom Solutions information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document require an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC Custom Solutions, visit http://www.idc.com/prodserv/custom_solutions/index.jsp.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com