



WHITE PAPER

Cloud Architects Choose Datacenter Flash to Power Industry's Largest Clouds

Sponsored by: SanDisk

Eric Burgener
Jeff Janukowicz
April 2016

Laura DuBois

IDC OPINION

Cloud architects designing public and private clouds must satisfy business needs while meeting cloud infrastructure scale, performance, and agility requirements. One challenge often faced is migrating existing workloads to public cloud services. Analyzing application portfolios, identifying dependencies and common infrastructure platform components, and assessing migration feasibility with an eye to establishing shared IT service centers for cloud operations are core requirements. However, cloud infrastructure must frequently be designed to support both traditional enterprise applications and next-generation cloud-native applications with very different I/O patterns.

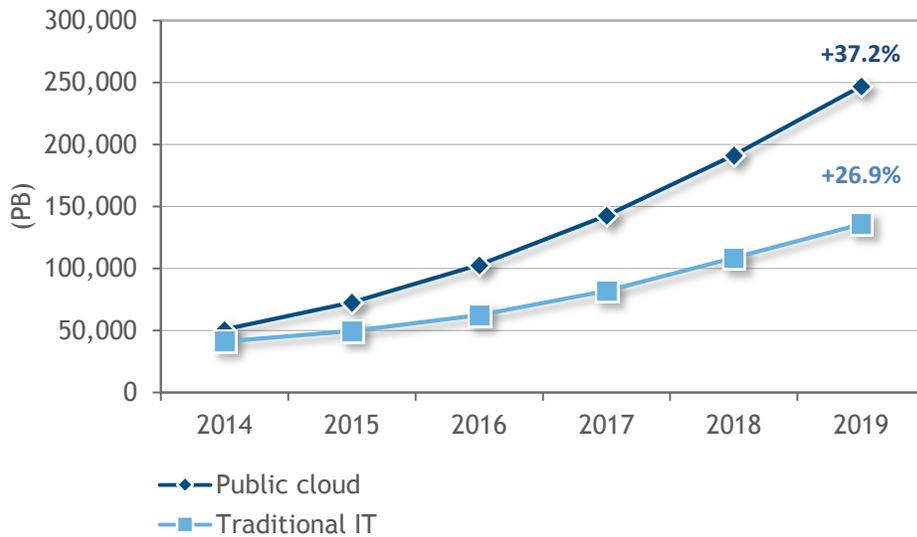
Cloud IaaS and PaaS infrastructure can provide the launchpad for application innovation but must meet varying service- and application-specific performance requirements. Persistent tiers, both flash and HDD, in cloud infrastructure must satisfy consistent throughput, latency, and performance isolation requirements over time. With the understanding that more workloads are moving to flash every day, cloud architects must consider that flash memory performance is heavily workload dependent, with performance variance over time based on block size, random versus sequential I/O, read versus write patterns, temporality, and data throughput requirements. Further, IaaS and PaaS today run mixed workloads with varying application I/O patterns. As a result, cloud architects must build out cloud infrastructure with both existing workloads and new applications in mind, running proof of concepts accordingly.

STORAGE DISRUPTION AND THE RISE OF FLASH

Four disruptive forces are driving change in the datacenter storage infrastructure: cloud architectures, software-defined storage, the convergence of compute and persistence, and the growing adoption of flash memory. Storage capacity for public cloud infrastructure will grow at a 37.2% compound annual growth rate (CAGR) through 2019, while enterprise storage capacity for traditional IT will grow at a more modest 26.9% CAGR (see Figure 1). However, with this capacity growth comes a shift away from the use of commercial SAN and NAS appliances to server-based shared and shared nothing storage architectures that leverage internal or DAS JBOD. The rise in software-defined storage approaches comes as a result of increased focus on cost reduction, automation, and lock-in avoidance, in particular among cloud builders. Last, the growing need to reduce network latency and increase IOPS has fueled adoption of flash memory in server and storage architectures, bringing application processing, analytics, and persistence closer together than ever before.

FIGURE 1

Worldwide Storage Capacity by Deployment, 2014-2019



Source: IDC's Worldwide Quarterly Cloud IT Infrastructure Tracker, 3Q15

TECHNOLOGY TRENDS DRIVING CLOUD FLASH ADOPTION

While a large part of IaaS adoption today is supporting existing workloads, cloud architects must design cloud infrastructure to support existing as well as new application use cases. Cloud infrastructure, either IaaS or PaaS, provides the stem for innovation in the cloud. This innovation is materializing through approaches and processes such as database as a service, analytics in the cloud (analytics as a service [pay as you go]), DevOps, real-time data processing, data lakes, and newly developed applications supporting Internet of Things (IoT) initiatives. Cloud service providers (CSPs) that do not provide customers with the ability to use cloud infrastructure as a launchpad to further developments, insights, and offerings will quickly find themselves sidelined.

While flash performance is workload dependent based on block size, random versus sequential I/O, read versus write patterns, and data throughput requirements, the following trends share a set of common requirements that lend themselves to greater use of flash memory:

- Need for real-time data capture, processing, and analytics
- Elimination of network and storage infrastructure latency
- Consistent and predictable application and infrastructure performance
- Large-scale data processing of complex data sets
- Performance and capacity tiers that can scale independently
- Ability to operate in scale-out cluster architectures with shared or shared nothing designs

These requirements point to the greater use of flash, which is increasingly used in the cloud by IaaS, PaaS, and SaaS providers alike. Persistent tiers in cloud infrastructure must meet predictability,

throughput, latency, and performance isolation requirements, including consistency in IOPS. Flash memory performance is heavily workload dependent, with performance variance over time based on block size, random versus sequential I/O, read versus write patterns, and data throughput requirements.

Another cloud requirement of flash is the ability to set and enforce performance tiers (this allows CSPs to charge more for a premium service that they know for sure they can provide). "Temporality" is a factor in cloud environments. A greater occurrence of hotspots that migrate around in cloud infrastructure with unpredictable workloads means that even if you've highly tuned an infrastructure for a certain I/O profile, that profile changes dynamically in cloud infrastructures more often based on the fact that you've consolidated varying workloads on virtual infrastructure.

The economic model of public cloud IaaS/PaaS places enormous cost pressures on the service provider. Offering a pay-as-you-go pricing model requires ongoing cost optimization analysis and implementation by the cloud architect. This has brought about new approaches to designing and building infrastructure, including moving away from an "appliance/branded systems" approach to a self-built software-defined infrastructure. It has also resulted in greater focus on storage efficiency including RAID optimization for performance, cost, and density. Leading IaaS providers are offering multiple compute instance volumes using RAID-0 disk striping to increase aggregate IOPS or improve sequential disk throughput in conjunction with value-added cloud services, snapshots, and availability zone replication to protect customer data sets.

Many new applications being developed to run in the cloud are employing NoSQL, schemaless databases where time-sensitive I/O completion is paramount. While the use of NoSQL is common across customers, the I/O pattern of NoSQL workloads will differ based on use case – varying from small block, random reads to large, sequential writes of streaming media, for example. Cloud architects must select an underlying persistence tier that can satisfy NoSQL small, random reads and the write streams of NoSQL logging/journaling and indexes.

A longstanding hurdle in datacenter infrastructure and one that is exacerbated in virtualized, shared cloud environments is the I/O bottleneck. Optimization of the CPU utilization resulted in the need for I/O optimization. The use of flash has been a game changer for application developers in eliminating storage I/O as the bottleneck in achieving IOPS and latency objectives. Consider a game developer using IaaS/PaaS to write the next "Candy Crush." The I/O pattern for mobile gaming is lots of small block reads and writes, with the need to adjust for and minimize latency. The sheer randomness of the I/O combined with the variability in the number of concurrent users itself makes latency a paramount concern in this workload. So, while IOPS and throughput are important for many use cases, latency avoidance is mission critical in consumer gaming applications.

Internet of Things Applications

At the root of IoT value is the collection, processing, and analysis of data generated by sensors and actuators embedded within an IoT device. IoT endpoints are performing real-time data capture, collection, and event processing. In IoT applications, scalability, network bandwidth, and response times necessitate a distributed approach. Given the sheer volume of data that IoT devices are collecting, it becomes practical for these devices to perform some level of data processing locally to minimize the amount of data that has to be transmitted across wide area networks. However, distributed IoT data must be transmitted to a centralized application using secure methods. This includes ensuring the authenticity of the data from the sensor to the collector and the eventual transmission from edge to core to datacenter cloud infrastructure. The IoT application must be

designed to ensure secure authentication, authorization, policy/control, monitoring, and analytics routines can run at scale in terms of data processing and analytics, supporting petabytes of data.

As it pertains to flash memory, there are several places for its insertion: First, in the endpoints themselves; second, in an edge data processing location where analog-to-digital conversion is performed; and last, in a centralized datacenter or cloud infrastructure. The IoT ecosystem is compute and data processing intensive and must support ingestion and analysis of large volumes of IoT data. With the endpoints, flash memory must be industrial grade, reliable, and secure and provide as much endurance as possible. Flash memory will be embedded in IoT endpoints that must provide for a long life cycle, handle temperature variation, and support field programmability. In edge locations, where data analytics functions are performed, flash memory affords immediate insight and control of IoT devices and processes.

With IoT analytics applications, the underlying infrastructure must support distributed, commonly cluster-based application architectures with scaling of compute and storage independently. Persistence in the infrastructure is provided by both a performance flash tier and a higher-capacity HDD or SSD tier. Flash is used as a performance tier to support and provide real-time analytics on volumes of IoT-generated data. However, analytics processes must be consistently performant at scale, even while background infrastructure processes such as garbage collection are performed. For IoT to provide immediate insight and take action on endpoint processes, analytics must run close to where the data is collected, while more complex analytics necessitate persistence of IoT data and distribution of it to a centralized cloud infrastructure for deeper computation and analysis.

Enterprise Data Lakes and Hadoop

Finding the value in large volumes of unstructured data is increasingly an imperative for any organization. New types of data, both machine and human generated, include IoT sensor data, Web logs, click stream content, and social media data. To mine this data for insights, Hadoop analytics clusters are complementing existing systems of record such as data warehouses and traditional relational database systems. To satisfy a growing need for real-time processing, in-memory databases such as SAP HANA are complementing traditional databases. The volume of data in and across these systems necessitates a scale-out approach to analytics, compute, and data persistence processes.

Data lakes can be thought of as a corpus of unstructured and semistructured data collected and collated from different sources into a single, unified data pool. A data lake offers multiple access points for data on-ramping including standard network access protocols and RESTful object interfaces. However, more crucially, a data lake supports the storing of data in a manner agnostic to how it is moved into the repository and in a manner that makes it easier for adjacent big data workloads to analyze it. As such, data is stored using open standards rather than proprietary formats. In enterprises, data lakes can be considered a central storage repository for consolidating different types of data.

Data lakes must provide upstream tiering leveraging flash memory for analytics on hot data sets and downstream tiering to higher-capacity tiers. In short, the system should support data movement between tiers depending on I/O activity. While much of the data in a data lake remains on rotational media today, flash is needed to serve queries on hot data and increase IOPS in MapReduce processes. And increasingly, users are using flash not only in primary storage use cases such as block storage for analytics but also as a performance tier in large-scale object repositories.

NoSQL and In-Memory Databases

Traditional commercial and open source relational databases that have been in use for decades are now being complemented with hundreds of new types of NoSQL data stores. These new NoSQL databases include approaches such as columnar (i.e., Cassandra), document (i.e., MongoDB), and key-value (i.e., Riak, Redis) graph and object databases. Much of the innovation in these new databases is being spawned in the open source community, with commercial distributions to follow. While each approach is different, these NoSQL databases share the following: no use of a relational model and a schemaless approach as well as typically being run on clusters of compute resources. NoSQL schemas provide the application developer choice by using a particular database or database features best suited for a given application. The ability to improve I/O performance by handling larger data volumes, reducing latency, and improving throughput is the other reason NoSQL use continues to increase. In-memory databases such as SAN HANA and GemFire are also being used.

The sweet spot for NoSQL databases is in large data sets with high concurrency requirements. With this use case, performance is of paramount importance. To meet this requirement, many new NoSQL databases are being optimized for flash use. This optimization is occurring to allow the application architecture to keep up common demands including tens of millions of requests per hour, with nearly 100% of these requests completed in sub-5ms. While DRAM would meet these requirements, cost is a factor in today's cloud infrastructure. This is where the use of SSD comes in. In fact, most leading NoSQL services offered by Amazon and Microsoft are put on SSDs by default. These demands in NoSQL environment databases have resulted in a greater use of flash memory for caching and persistence to achieve consistently faster performance in application response times over high load.

Next-Generation Applications

The rise of mobile applications and social communications together with the proliferation of cloud computing has spawned the design and development of next-generation applications. These next-generation applications share a set of common attributes: They are designed to horizontally scale, tend to be stateless in nature, leverage open source software, and are designed to run on scale-out infrastructure. These applications are designed with cloud in mind – some run in a public cloud IaaS, while others run in a private cloud. These applications are increasingly composite based, run in containers, and leverage microservices. Designed to be distributed, the application itself, rather than the underlying infrastructure, provides for reliability and resiliency through the geodispersal of data.

The implication for the underlying infrastructure is the need for performance and capacity over resiliency. Performance is achieved through the use of a flash tier, with databases placed on flash block storage and the capacity tier placed on an object store. And while object stores have traditionally been HDD based, the capacity or content tier is now starting to be placed on flash as well.

Continuous Integration/Continuous Development and DevOps

The development of next-generation applications is complementary to the use of DevOps. DevOps is a methodology being used to develop new applications using agile, continuous integration (CI) and continuous development (CD) methodologies. DevOps is at once both a process and a team of people, both developers and operations staff that collaborate through the development life cycle when creating and operating a service. It's very common for next-generation applications to use a DevOps mentality in the building and running of the application. DevOps teams are often formed with line-of-business organizations, as opposed to centralized corporate IT, to build an application and get it to market quickly. Because the development and operations teams are so closely coupled, there is a tendency to use open source

software in not only application development but also operations, specifically in the underlying infrastructure (i.e., OpenStack APIs, Linux, Ceph, and open source databases and middleware).

From a caching and persistence perspective, there is reliance on server-based storage with shared or shared nothing storage designs and leveraging of flash for the database tier and, increasingly, the object-based tier. Docker Swarm enables persistence for production applications for Linux container architectures. And there is growing demand for object stores with larger and larger data sets, in both object size (i.e., high-fidelity medical images) and the expanding storage of more objects, including larger rich media such as 4K video on-demand.

Database as a Service or Cloud Databases

With the advent of cloud computing, businesses are both moving existing databases to cloud infrastructure and deploying new databases on infrastructure as a service. Database as a service (DBaaS) is a broad category that spans different deployment scenarios. One option includes installing database machine images on rented, cloud-hosted server and storage resources. Application and database owners have access to database functions via standard database tools and manage the cloud-hosted database on their own. A different implementation of database as a service is where a service provider provides not only the underlying infrastructure but also the database engine of its choice. For example, Amazon AWS RDS provides users with a range of commercial and open source relational databases. These services include additional management functions such as database patching, backup, retention, geographic data distribution, replication, and recovery. DBaaS can span both SQL and NoSQL schemas; other examples include Google Cloud SQL, Microsoft Azure Cloud SQL, AWS DynamoDB, and OpenStack Trove.

Cloud databases can include both traditional RDBMS and new database schemas, including databases designed to run in shared cloud infrastructure while meeting scalability, availability, and reliability needs. Native cloud databases offer advantages such as increased scalability and accessibility, automatic failover, and fast, automated recovery. Common concerns with cloud-native databases include security and data privacy and data lock-in. In both the DBaaS approach and the cloud-native database approach, the use of flash memory can offer faster time to market or query result and greater use of resources including running of NoSQL jobs by the hour.

Analytics in the Cloud/Analytics as a Service

With applications and data increasingly running in the cloud, a natural extension is to run analytics functions in the cloud, in particular because of the compute- and volume-intensive nature of data analytics today. The leading motivator to analytics in the cloud is in leveraging the scale of cloud computing while accelerating time to insight or time to market. Cloud IaaS allows firms to leverage processing spends and capacity to process large volumes of data. Flash not only makes things run faster but also ensures that large jobs complete the first time they're submitted. Faster batch times allow end users to work with fresher data sets (they can be updated more often with flash when the batch time drops from 8 hours to 25 minutes), which lead directly to better business decisions and improved efficiencies. Analytics in the cloud allows an enterprise to pay by the hour, which makes opex more attractive than on-premises capex investments.

SANDISK DATACENTER FLASH SOLUTIONS

SanDisk datacenter flash solutions include SanDisk InfiniFlash and the SanDisk SATA solid state disk (SSD) portfolio. SanDisk has been a material flash memory supplier for close to three decades, and today, SanDisk datacenter flash powers the top 12 worldwide Internet properties and cloud providers. Cloud architects building native IaaS and PaaS platforms can also benefit from the architectural approach chosen by the largest public and private clouds in the world by leveraging SanDisk as both a technology and innovation provider and a business partner.

SanDisk's InfiniFlash System

SanDisk built InfiniFlash as a next-generation scale-out storage platform enabling the use of flash for high-capacity cloud environments where flash at scale was not economically feasible. From a price/performance perspective, InfiniFlash was designed to slot in between high-performance, high-cost conventional flash arrays and lower-performance, lower-cost HDD options. InfiniFlash not only offers better performance than HDD-based solutions for high-capacity storage use cases but also consistently delivers sub-1ms latencies with 99.9% quality of service (QoS). Further, the economic hurdles that have stymied the use of flash for secondary storage use cases are addressed with InfiniFlash. InfiniFlash enables rack-scale consolidation of storage. Higher-density configurations, enabled by compact packaging as well as storage efficiency services, combine with lower power, cooling, floor space, server, and licensing costs to deliver a reduced total cost of ownership (TCO) over alternative high-performance HDD offerings.

The architectural details of InfiniFlash include:

- SanDisk NAND (designed, engineered, and manufactured by SanDisk)
- Datacenter-grade SSD
- 8 SAS ports, 12Gbps SAS connectivity
- >2.0 million raw IOPS
- 1.2 million IOPS, 4K 100% RR @ <1ms 99.9%
- 880,000 IOPS, 4K 90%/10% RR/RW @ <1ms 99%
- >12GBps throughput
- Can scale to tens to hundreds of petabytes on demand
- Performance at scale for both object and block data throughput requirements
- Disaggregated storage with direct attach to servers via SAS ports
- 512/256/128TB capacity points in a 3U chassis (8TB InfiniFlash cards, 64 per chassis)
- Density optimized for attach to 1U and 2U servers
 - Support for storage disaggregation configurations to enable servers to run storage software with Ethernet/InfiniBand/Fibre Channel physical interfaces to the datacenter

These benefits are driving the adoption of InfiniFlash into cloud use cases such as Hadoop clusters, NoSQL analytics, data lakes, Ceph performance needs, media streaming, and content repositories. Leading cloud deployments of InfiniFlash include Ceph block and object storage deployments for consumer cloud and telco OTT offerings and scale-out file-based storage with Nexenta and other distributed file systems.

SanDisk SATA SSD

SanDisk SATA SSD is a datacenter-grade 2TB SATA SSD based on NAND that is designed, engineered, and manufactured by SanDisk. The SanDisk CloudSpeed family of SSDs is optimized to manage transactional and mixed workloads – to satisfy consumer- and enterprise-grade cloud services. To offer the database and application performance tiers of service, service providers can use CloudSpeed devices to increase compute storage density for efficient datacenter utilization.

CloudSpeed Ultra SSDs are optimized for latency-sensitive application and mixed-use workloads. For predictable, random I/O write performance, this flash storage solution works within existing infrastructure – the SATA interface – and can replace hard disk drives for lower TCO and better datacenter efficiency through server consolidation.

CloudSpeed Eco SSDs are optimized for read-intensive cloud workloads – video/media streaming and content repositories – where sustained data throughput is critical to satisfy end users of cloud services. Designed to replace high-capacity HDDs, CloudSpeed Eco Gen. II enables highly scalable cloud architectures for consumer and enterprise cloud offerings.

The architectural details of SanDisk CloudSpeed SATA SSDs include:

- Onboard capacitors to protect data written to write buffer in case of power failure
- Cryptographically signed firmware to protect firmware from being overwritten
- Guardian Technology Platform for improved flash endurance and data integrity/protection
- Full data path protection
- Support for RAID-0 disk striping configurations to maximize cost savings, performance, and density for MySQL and NoSQL workloads with data protection mechanism (i.e., snapshots, replication)
- Optimizations to support NoSQL workload I/O patterns including both small, random read activity and large sequential write activity

Application developers need to know 99.9% QoS write latency to accommodate application threads that work across compute clusters; average latency does not provide meaningful measures of predictable performance. 24 CloudSpeed SSDs in a two-socket server provide over 700,000 4K write IOPS with enterprise-grade reliability – a projected UBER of only 1 unrecoverable error in 10^{18} reads.¹ Better UBER enables flexibility by some service providers to deploy RAID-0 for NoSQL and MySQL for performance; Availability Zones (AZ) replication and snapshots help ensure data protection. In short, anyaaS is now mission critical, and zero downtime is expected.

CloudSpeed SSDs are used today across ecommerce, NoSQL database, and object store use cases. eCommerce requirements include predictable performance for transaction processing to avoid poor end-user experiences or abandoned shopping carts. SSDs used for NoSQL deployments are applicable for event processing with the acceleration of small, random reads across NoSQL databases and predictable performance for data stream writes. CloudSpeed SSDs support sustained throughput of over 500MBps per SSD, allowing two-socket servers to deliver gigabytes per second of throughput

¹ = 1 recoverable error in 10^{18} bits read, per the JEDEC 64.8 specification, as stated on the CloudSpeed data sheet, available by [clicking here](#).

for applications like video streaming. CloudSpeed SSDs are also applicable for file sync and share, social media (end user-created content) repositories, and rich media files where access time matters.

The benefits of CloudSpeed SATA SSDs in cloud deployments include a 99.9% QoS for random write latency based on a 4KB transfer rate. This enables application developers to create applications with reliable and predictable performance.

SANDISK AS A BUSINESS PARTNER

SanDisk has a 25-year history of flash technology innovation and production. Leading cloud service providers, known by consumers and IT professionals alike, leverage SanDisk datacenter flash to power their Internet properties and cloud services. And with the ongoing adoption of cloud services, these customers are one of the largest growing business units at SanDisk today. SanDisk and Toshiba have entered into a joint venture to run their own NAND chip fabrication. The benefits to the two companies as well as their cloud provider customers are the timely and quality production of NAND and the ability to retool for 3D NAND.

To support its datacenter flash business, SanDisk employs over 8,000 people across R&D, engineering, manufacturing, and supply chain and has over 5,000 patents worldwide. SanDisk advanced development includes ongoing research into 3D memory and innovations that increase flash memory capacity and make new devices possible. SanDisk also has resources for critical SSD flash firmware and SSD controller development to optimize NAND operations for latency, I/O, wear leveling, garbage collection, and endurance. Last, in addition to controlling NAND production, SanDisk also owns and operates its own manufacturing facilities. In short, SanDisk controls everything from the silicon up in design, development, and production.

While spinning disk still remains in most datacenters the same way tape does, there is an undeniable increase in the use of flash memory persistence and a decreasing role for rotational disk. This places flash as a strategic component of any datacenter strategy, but in particular for cloud service providers that must meet performance, automation, and economic requirements. SanDisk has recognized the importance of cloud services to drive the IT economy and works with leading cloud providers as well as Taiwanese ODMs to jointly deliver hyperscale infrastructure. SanDisk serves its customers with a direct sales force and a network of VARs, SIs, and channel distribution partners. SanDisk also has a track record of supporting traditional server and storage OEMs that serve traditional enterprises. This experience proves SanDisk's commitment to supply chain and SSD flash engineering road map timetables.

CHALLENGES/OPPORTUNITIES

While it seems like a daunting challenge for a cloud architect to build the infrastructure to support a range of workloads, use cases, and I/O patterns, much insight can be gained by looking at what the top IaaS providers today are recommending for cloud infrastructure and the use of flash memory. There is an enormous opportunity for cloud architects not only to eliminate the I/O bottleneck but also to deploy flash as the underlying persistence tier; this allows customers to do things with their data that have never been possible before. The addition of flash memory to cloud infrastructure provides a foundation and platform for innovation for use cases such as:

- Database-as-a-service offerings
- Analytics in the cloud/analytics as a service

- New product development models including DevOps and CI/CD
- Real-time data processing
- Internet of Things initiatives
- New mobile applications supporting digital transformation initiatives

CONCLUSION

The reality is that more workloads are moving to flash every day. This has been recognized and embraced by the leading IaaS/PaaS providers today. Cloud architects that design infrastructure with flash memory in mind are best able to capitalize on the cloud use cases identified in this white paper. The application-specific performance requirements of these use cases can be met with SSD- and PCIe-based flash deployed in both shared and shared nothing storage architectures. Persistent tiers in current and future cloud infrastructure will increasingly employ both flash and HDD to satisfy consistent throughput, latency, and performance isolation requirements over time. Flash memory deployed in cloud infrastructure will support mixed workload environments with varying application I/O patterns. SanDisk is a leading datacenter flash systems supplier that cloud architects should put on their short list when running proof-of-concept testing.

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-insights-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 2016 IDC. Reproduction without written permission is completely forbidden.

