

Tapping the True Value of Data

WHY TAKING A VALUE-DRIVEN APPROACH CAN TAKE THE PAIN OUT OF DATA MANAGEMENT AND DELIVER PREDICTABLE APPLICATION PERFORMANCE FOR A TOP-NOTCH USER EXPERIENCE

Not all data has equal value, yet storage and data management solutions treat diverse data uniformly. When an enterprise needs higher performance, most storage vendors accelerate all data indiscriminately, regardless of its value to the business. A recent survey demonstrates that IT professionals recognize the value of managing stored data according to its business value, but less than a third are attempting to do so in a formal manner. A new approach providing improved flash utilization to manage workloads, not capacity, and QoS managed caching, makes it easier to manage data based on business priorities and the value of the data.

Data is increasingly critical to business success. But with growing volumes, many enterprises are struggling to deal with capacity issues and, as a result, overlooking the issue of providing fast access to the information of most value to users. Data centers are increasingly turning to flash memory storage to speed up access to data. This can be an arduous and expensive proposition and still falls short when it comes to prioritizing data with important business value.

Because they are generally managing stored data in a uniform manner, most enterprises in effect are paying a premium to manage storage whether the data is valuable or not. In fact, according to IDC, the volume of data created annually is doubling in size every two years and, largely driven by big data and the Internet of things, will reach 44 zeta bytes in 2020¹. Yet in 2013, IDC says, less than a quarter of data would be considered useful if tagged, and just 5 percent of useful data could be considered especially valuable or “data rich.” Being able to discriminate between high value data and low value data would enable enterprises to manage their information for performance and cost-effectiveness.

¹Vernon Turner, et al, “The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things,” April 2014, International Data Corp.

NOT ALL DATA IS CREATED EQUALLY

In a survey conducted recently by IDG Research Services, 90% of the more than 200 respondents indicate that their organization’s strategy for managing and storing data is either critical or important to achieving business goals. When it comes to achieving those goals, however, improving the customer or user experience is most often cited as a driving force, even over bottom-line profit growth or improved margins.

In some ways, the server virtualization phenomenon that has swept enterprise IT has exposed fundamental weaknesses in storage architectures. As virtual workloads



Strategy for Managing and Storing Data: Impact on Ability to Achieve Key Business Goals

Not surprisingly, in organizations where more than 50% of data is considered mission-critical, managing and storing data is more likely to be a critical priority



increase and virtual DRAM is paged out, multiple applications are forced to access storage for application data in highly random patterns, making it difficult to achieve optimum input/output cycles that are crucial to workload performance. That's a management headache, for sure, but can also impact the user experience in a very negative way.

The root of the problem is quite evident. In the IDG Research Services survey, respondents report that on average, 43 percent of applications are considered mission-critical, as is 50 percent of data. When so much is mission-critical, it becomes exceedingly more difficult to assign segment applications and data based on performance of resources. The issue is exacerbated for smaller companies, which are more likely to consider all data to be mission-critical.

That said, the survey also finds that 62 percent of respondents believe that inefficient use of storage resources is the most negative consequence of managing all data equally. Additionally, 54 percent cite the cost of storing non-critical data as a concern.

In fact, while enterprises are relatively happy with the availability and backup and recovery capabilities, when it comes to factors such as the impact on application performance, data latency and cost of storage, less than half of survey respondents are satisfied with their current storage solutions. This is particularly concerning, given that respondents overwhelmingly cite improving the customer/user experience as their top business objective.

UNIFORMLY MANAGING VARIABLE DATA

There is broad recognition in IT that not all data has equal value to the business, but relatively few are able to manage their data accordingly.

Specifically, the IDG Research survey finds that 94 percent of respondents recognize the value of managing stored data according to the level of priority or value of the data to the business. Yet less than one-third indicate they have any way to formally manage stored data according to the level of priority or value to the business. Another 36 percent say they manage all data equally, while 27 percent indicate they are starting to assign priority levels to stored data but lack any formal system in place. More than half are concerned by the increased cost of managing and storing non-critical data.

"No one has enough money to do everything they want with data, and ultimately they are going to have to make choices over how to effectively manage, access and store it," says Mark Peters, senior analyst with Enterprise Strategy Group (ESG). "IT is under constant pressure to deliver a lower total cost of ownership," he says, and must come up with more effective ways to manage not just the growing volumes of stored data, but also user expectations in regard to velocity and access.

Trying to establish rigid categories of assigned value is not likely to offer much relief, however, as 75 percent of survey respondents say the value even of mission-critical data changes constantly as it ages. Painstaking planning is required to prioritize data into specific categories for performance, availability, reliability, and backup. Meanwhile, unforeseen events and shifting business priorities can impact the value of data, most survey respondents indicate.

ANGST IN VALUE-DRIVEN DATA MANAGEMENT

"The ability to use performance where it is needed and to apply that to specific workloads is the key to efficient storage management," says ESG's Peters. He adds storage solutions generally take a reactive approach, spreading out resources "like peanut butter" rather than focusing on mission-critical needs.

Many enterprises have responded to data and application performance issues by adopting solid-state devices (SSDs) utilizing flash memory, which is faster and lacks the moving parts and performance limitations of hard disk drives (HDDs). Flash memory makes applications faster—up to 10x faster compared to applications relying on hard disk drives—by speeding up the input/output operations per second (IOPS). This kind of performance boost can have huge implications when it comes to achieving key business objectives, but it comes at a cost.

For mission-critical applications and data, flash can provide significant dollar/IOPS advantages over slower HDDs. But because HDDs have an advantage when it comes to the cost per gigabyte (dollar/GB), utilizing all-flash arrays for storage of non-critical data is an expensive use of resources. Hybrid storage arrays that combine flash with HDDs provide a partial solution, but can lead to unpredictable performance because IT can't control whether mission-critical or non-critical data

accesses the more costly flash storage.

While flash prices are expected to continue to decline over time, there are no indications that it can replace HDDs on a dollar/GB basis. "You can use as much flash as you can afford, but at some point it's wasteful," says Peters. "Even if flash prices were very close, it would be a bad use of resources; you don't need a Ferrari if all you do is go to the supermarket once a week."

Flash can be used more effectively, however. "We are seeing flash going through a very standard technology adoption lifecycle," says Chris McCall, senior vice president of marketing with NexGen Storage, which provides third generation flash storage solutions. He says flash is undergoing a transformation similar to how virtualization transformed the use and deployment of servers a few years ago.

"X86 processors became so powerful that a single application could not leverage all the available resources, so along came the hypervisor, which allows you to run multiple applications on the same server and increase server utilization. The more critical change, however, is the way applications were managed changed drastically. Before virtualization, it was about configuring a server to meet application needs – very technology focused. Now with virtualization, it's all about managing workloads, understanding which ones are mission critical and knowing how much resources to allocate," McCall says.

"Today", says McCall, "storage is undergoing the same transformation. Flash is analogous to x86 processors; the industry has offered customers the ability to deploy more flash than customers need. Similar to the advent of the hypervisor, we are at the precipice of a shift in focus from flash technology to how it's managed and improving the utilization of it. The only way to do this is by understanding data, the workload associated with it and what the value of the data is to the business to ensure the right data is in flash at the right time."

RETHINKING STORAGE ARCHITECTURES

With a new storage architecture and policy-based management, IT can more effectively utilize its flash and HDD resources to prioritize high value apps and data and to drive down the effective cost of flash. Obviously, storage vendors design storage Quality of Service (QoS) features into their systems to help prioritize performance by

application. But it's very difficult to do so in a hybrid array where performance varies dramatically depending on whether data is stored in flash or in HDDs, or other types of storage media.

SSDs look like drives to system designers because they retain legacy disk interfaces and rely on disk protection schemes that were designed in the 1980s specifically for high-capacity, slow I/O mechanical disks. By emulating disks, SSDs can be wedged into existing software stacks so they work with existing disk controllers and RAID systems. The problem is that the native features of solid state are hidden by emulation. Performance gains are blunted, protection at the silicon level is ignored, and software interfaces at the cell level are simply not considered.


With flash in the data center becoming pervasive, a new storage architecture is required to fully empower user applications. Painstaking planning is required to prioritize data into specific categories for performance, availability, reliability, backup, etc., and then even more effort is spent designing and implementing storage systems that can meet those goals.

The QoS policy should define performance targets for a volume in terms of IOPS, throughput and latency. This eliminates unpredictability and the need to design storage systems for peak workload performance. For example, an organization can assign a business intelligence application 30,000 IOPS, an order database 25,000 IOPS, and a marketing file share 5,000 IOPS.

Managing where data resides should be a real-time, automatic process of the storage system itself. Based on



Management of Data Storage Today



SOURCE: IDG RESEARCH SERVICES, JANUARY 2015

Taming growing volumes of data

NCS is a 44-year-old, Ohio-based company that initially was founded as a commercial collection agency and has evolved to provide corporate credit services to support a fast-growing client base.

As NCS grew, so did the size and complexity of its IT systems, many of which are proprietary solutions designed specifically for the credit industry. The NCS IT team needed to be able to assign storage services based on the criticality of multiple workloads.

Utilizing a NexGen N5 Hybrid Flash Array, NCS is able to assign specific tiers to each application. "We can change performance in real time, moving a particular data set up into a higher performance service level for a few days when it's critical to the business, then move it back down when that task is complete. It gives us an incredible amount of flexibility," says Michael Frank, manager of the IT services group.

The N5 delivers twice the capacity in half the physical footprint as the prior system and ensures that mission-critical applications maintain consistent, guaranteed performance levels, even during VDI boot storms.

the value of the data, the storage system should locate it in the appropriate location (RAM, flash, HDD, etc.), in real-time, to meet the QoS objectives.

Most all-flash arrays and all hybrid storage vendors have selected SSD technology with a SAS or SATA interface, which severely limits the performance potential of the system. A PCIe interface, however, connects flash more closely to the CPU so it can operate in a manner more akin to memory than to spinning disks.

McCall says NexGen, which was recently spun out of SanDisk to operate on its own, has been developing a storage architecture designed specifically around PCIe flash and storage QoS and service levels that allow IT to instruct the system as to the value of data. A dynamic data path manages data in real-time, making decisions based on QoS performance policies to ensure the data resides in the appropriate location (flash vs. disk), at the appropriate time.

"With PCIe we can run more workloads with a smaller footprint," says McCall. "We can add more flash for performance non-disruptively and scale capacity by adding more disk drives non-disruptively. That way, no matter what workloads you have, the system can grow to accommodate changing needs of users and the business."

As data management evolves into a more value-based prioritization strategy, other vendors are likely to offer other innovative approaches like NexGen's in an effort to address user needs for ready access to the most important data.

MAKING THE MOST DATA

Allocating storage performance, capacity, and data protection policies based on the value of the stored data is widely viewed by survey respondents as improving utilization of storage resources. Managing all data equally, the survey indicates, is a costly way to handle non-critical data, contributes to poor application performance and end-user complaints, and decreases the availability of truly mission-critical data.

In order to achieve business goals, IT must move beyond capacity management and instead manage data and storage in a way that reflects the value and priorities of the business, or will surely fall short on meeting business expectations. ■

To learn more about value-driven data management solutions, which align the cost of storing, accessing and managing data to the business value of the data, please go to:

WWW.NEXGENSTORAGE.COM