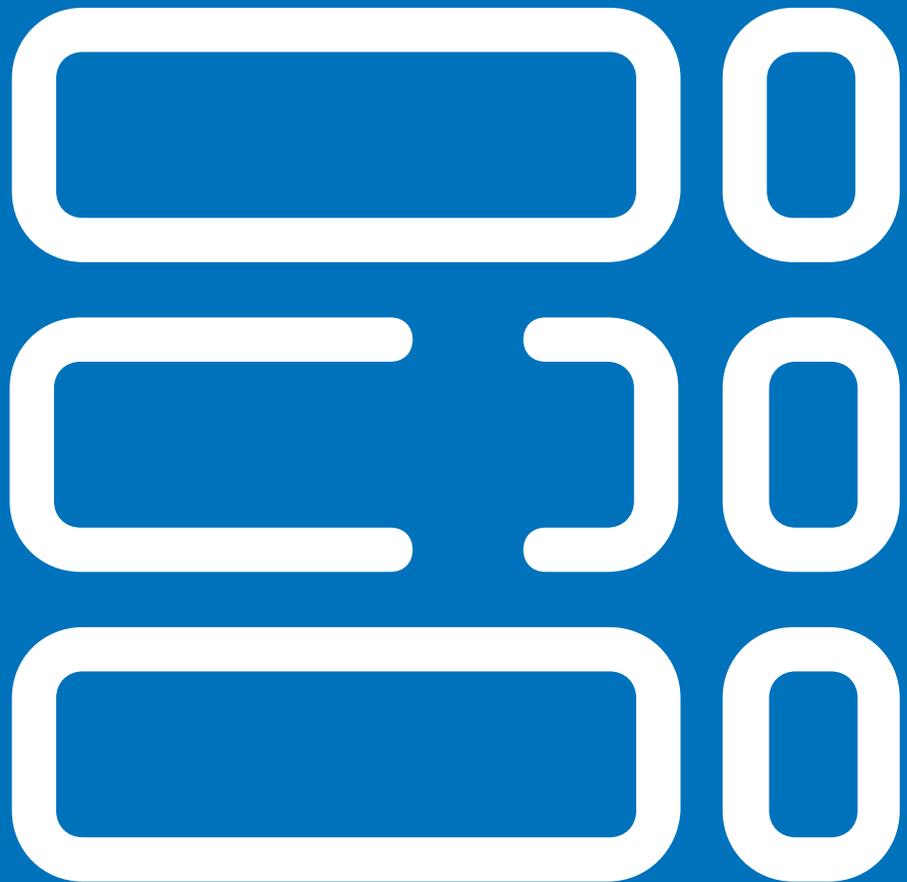


GAME-CHANGING INSIGHTS AT THE SPEED OF FLASH

PeerPathfinder Report



BASED ON ACTUAL USER EXPERIENCES & OPINIONS

This paper was co-produced by IT Central Station and 451 Research on the commission of IBM

CONTENTS

Page 1. Part I: Why CIOs Embracing Analytics Should Care about Flash

Meet the need for speed with the winning partnership of analytics and flash

Breaking through the bottleneck of disk performance

Page 2. Disrupt and thrive with the power of flash-driven analytics

Page 3. Part II: A Flash Revolution in Progress

Flash takes center stage

The falling price of flash

Page 4-5. Flash reduces operational costs

Page 6. Part III: Why IBM for Flash?

IT economics

Flash advantages

Ease of setup

Page 7. Ease of use

Scalability

Speed and performance

Page 8. Conclusion

Part I: Why CIOs Embracing Analytics Should Care about Flash

MEET THE NEED FOR SPEED WITH THE WINNING PARTNERSHIP OF ANALYTICS AND FLASH

New forms of data analytics software are delivering breakthrough levels of insights across multiple industries. CIOs may not be aware that to fully realize the power of this new software, and to fully achieve a return on their investment, they need to carefully consider the infrastructure requirements.

Businesses can only fully realize the benefits of next-generation analytics when the results of data queries are delivered rapidly, allowing organizations to respond quickly to shifting business conditions, or to drive disruption themselves. Analytics is a driving force within digital transformation, and is revolutionizing the speed at which competitive environments change. To survive and thrive in this accelerated climate, businesses must be able to take advantage of the opportunities identified by analytics as quickly as possible. In some cases, the opportunities may be short-lived. In others, the penalty for not reacting rapidly to a threat or potential disruption may be heavy, or even existential. In either case, the analytics needs to be fast.

Winning CIOs will understand the impact that this requirement will impose on their data storage systems. Because analytics involves computationally complex and intense data processing, the need for fast analytics translates into a need for rapid access to data at speeds that cannot be sustained by disk storage. This need is met by flash. Replacing

disk with all-flash storage slashes the time it takes to answer analytic queries from days to hours, or even minutes. Flash has already transformed performance for multiple mainstream enterprise applications. Because the performance of analytics is even more dependent on fast access to data than other applications, it is even more dramatically transformed by flash storage. For many organizations, it will be essential to achieving the full value of their investment.

BREAKING THROUGH THE BOTTLENECK OF DISK PERFORMANCE

Alongside the new forms of software, a major factor in the increased adoption of analytics has been the ever-increasing processing power of servers, as well as the use of processors designed specifically for analytics. Indeed, the availability of unprecedented levels of low-cost processing power has been a driving factor in the uptake of deep learning, or artificial neural-network based analytics.

But while processing power – and network bandwidth – has advanced significantly, the performance of disk drives in data storage has been nearly stagnant for many years. For a large number of datacenter applications, this has been a considerable problem, and has made storage a bottleneck in overall performance. The problems created by the poor performance of disk are further exacerbated for analytics, which is computationally intensive. Disk performance is limited by the mechanical constraints of rotational speed, and the speed at which read-write arms can move across a disk drive platter. Disk drive makers reached the practical limit for rotational speed almost two decades ago.

This means that disk storage not only significantly reduces the speed at which analytics queries can be answered, but it also impacts the depth of data analysis. Analytics can be an iterative process, in

which the results of one query lead to further queries. Fast results from queries speed the process of drilling down with further, refined queries to achieve insights. As a result, enterprises need to process and analyze large volumes of data at frequencies dictated by business needs. These needs can be met by flash storage.

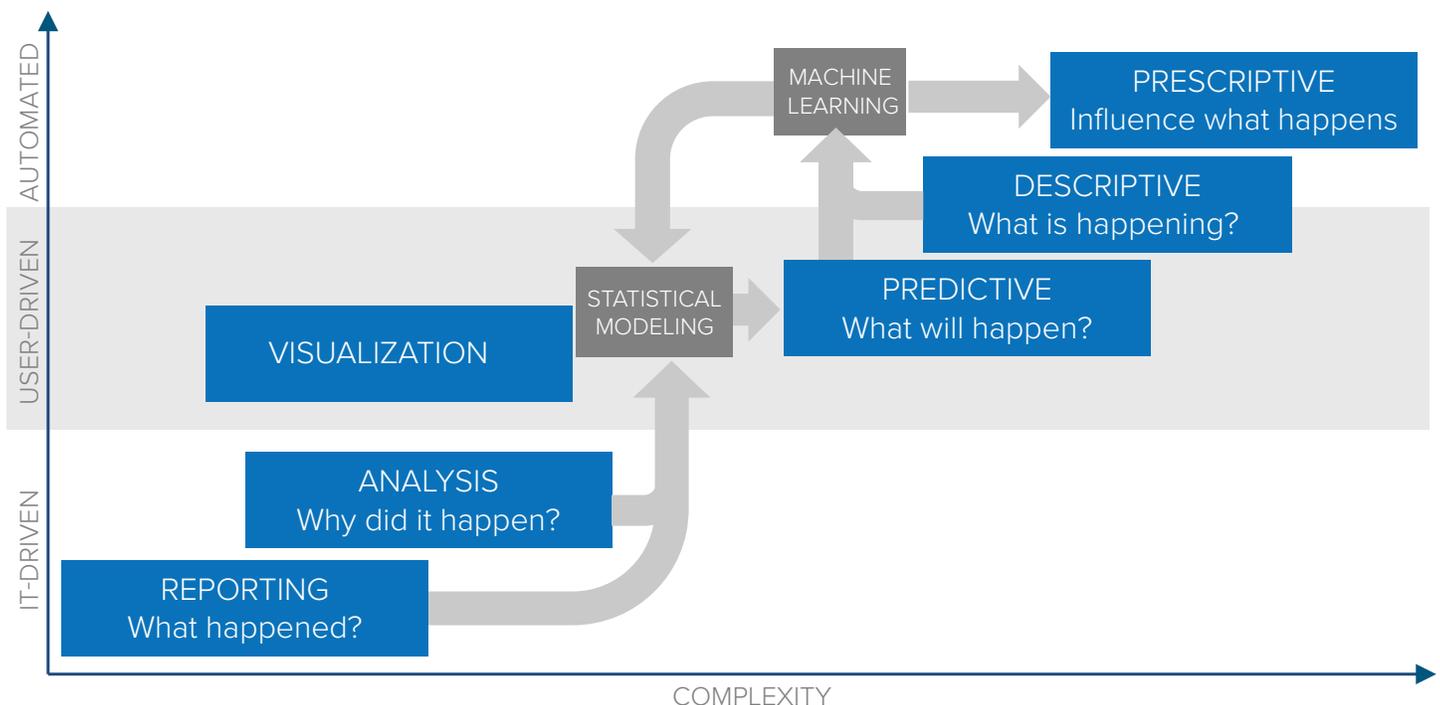
Flash outperforms disk storage up to a thousandfold in terms of the latency or time taken to handle requests to access random data blocks. Random IO performance is the Achilles' heel of disk, and highly randomized IOs are a strong characteristic of analytic processing. Data throughput when handling streams of sequential data is also significantly faster for flash storage than for disk. That reduces the time taken to load analytics systems with what can be very large volumes of data when switching between analytic datasets.

DISRUPT AND THRIVE WITH THE POWER OF FLASH-DRIVEN ANALYTICS

The new forms of data analytic software combine recent developments in data and analytics platforms, such as Hadoop and Apache Spark, with the power of artificial intelligence and machine learning. The result has been a step change in analytics capabilities. Alongside those developments, greater automation of the underlying processes has moved analytics out of the realm of statisticians, mathematicians and data scientists and into the hands of less-skilled, but more business-savvy decision-makers.

The result is that more enterprises are using predictive analytics, which is in the process of superseding historical reporting and analysis approaches. In addition to enabling users to make forecasts, predictive analytics is prescriptive (see Figure 1). To put it another way: having already

The Adoption of Machine Learning and Deep Learning



Source: 451 Research

asked what happened and why, enterprises are now able to ask what is happening now, and how they can influence what will happen next. Predictive analytics is set to become a fundamental part of a corporation's software arsenal, driving competitive advantage and new disruptive business models by delivering forward-looking insights through the use of machine learning and deep learning, two forms of artificial intelligence.

451 Research believes that alongside understanding the importance of ensuring that datacenter infrastructure is optimized to accelerate speed to insight, winning CIOs will be those who take the entire predictive analytics process into consideration. Specifically, enterprises need to create cultures that promote data-driven decision-making, and to provide self-service analytics to business users, as well as to data analysts and data scientists.

Part II: A Flash Revolution is in Progress

FLASH TAKES CENTER STAGE

Since the first commercial disk drives were launched by IBM over half a century ago, the storage of data on magnetic disks has become a cornerstone of IT. Although disk will remain a major presence in datacenters for many years yet, its role in IT has been changing dramatically over the last few years. Disk is being relegated into a secondary role as an enterprise storage medium, and is being displaced by flash.

In datacenters, flash is quickly becoming the default storage medium for primary or active, working

data, while disk has become the means of storing infrequently accessed data that is not performance-sensitive, such as archives and backups. Put simply, flash is the new disk, and disk is the new tape.

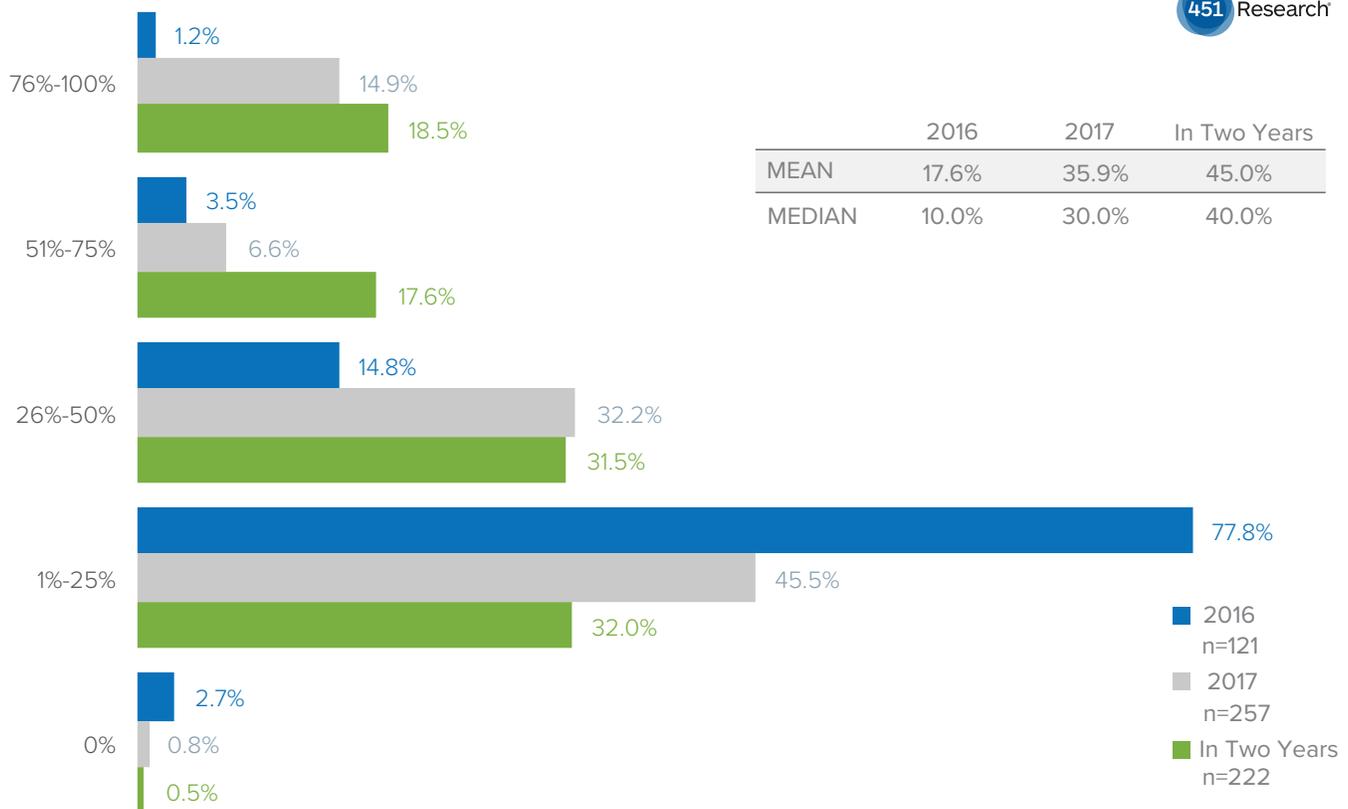
This transition has been happening quickly. Only four years ago, flash represented a very small percentage of the total storage capacity in typical datacenters, and enterprise storage systems were overwhelmingly based on disk. The majority of primary storage systems featured only a small amount of flash. Now, suppliers of those systems report that more than 50% of sales involve devices that are configured entirely as flash systems. That percentage will continue to increase as all-flash storage is applied to a wide range of applications including, but not restricted to, transactional databases, email, virtual desktop infrastructure, virtual server infrastructure and file storage, as well as analytics.

THE FALLING PRICE OF FLASH

The rate at which the price of flash has been tumbling over the last 15 years has been remarkable, and has resulted from exploding demand for flash memory chips in mobile phones and other devices. That soaring demand has driven up manufacturing capacity, and has driven major advances in flash technology itself. For the predominant NAND form of flash, prices have been consistently halving every two years for more than a decade. In 2005, NAND was more expensive than DRAM. Now, it is about 30 times cheaper.

Comparing the prices of different varieties of disk and flash storage systems is not straightforward because of the multiple variables involved, and the variations from environment to environment. However, one fact is certain: the tumbling price of flash has already made obsolete the fastest and most expensive enterprise disk drives – those running at 15,000 rpm, which were previously a storage mainstay for

AFA Storage Capacity Usage - Today and in Two Years



Q. Approximately how much of your organization's total primary (i.e., SAN-attached) storage capacity is deployed in an AFA configuration?
 Q. Approximately how much of your organization's total primary (i.e., SAN-attached) storage capacity is deployed in an AFA configuration in two years?

Source: 451 Research Voice of the Enterprise: Storage, Workloads and Key Projects, 2017

performance-sensitive primary or working data. On a capex basis alone, equivalent enterprise-grade flash drives are now cheaper per gigabyte of capacity, without even factoring in the opex savings that flash provides compared to disk (see below).

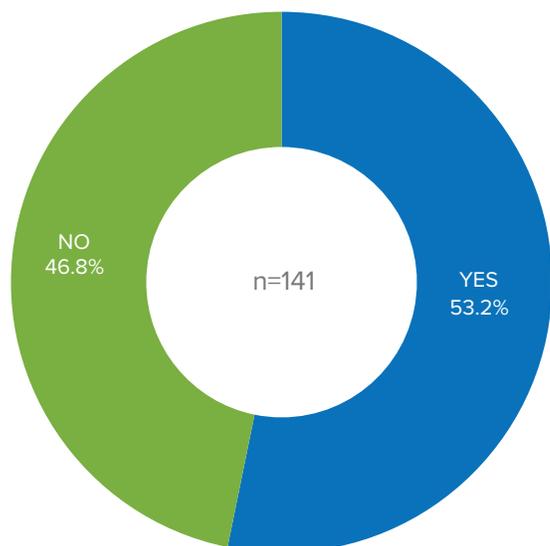
Next in line to become obsolete are enterprise disks running at 10,000 rpm. This is because flash prices per unit of data capacity are set to continue falling for several years yet. A major reason for this is that chipmakers are confident that they have refreshed the flash technology curve via a recent transition from two-dimensional to three-dimensional chips,

and by the ability to pack ever increasing amounts of data into individual flash memory cells.

FLASH REDUCES OPERATIONAL COSTS

Operating costs, or opex, often account for as much as half of IT budgets. Opex is lower for all-flash storage than for disk or hybrid disk-and-flash storage in multiple ways. The superior performance of all-flash storage reduces the number of systems needed to handle a given set of workloads, reducing administrative overhead. Four years ago, when usage of all-flash storage first began to increase, the huge majority of the systems were bought to handle

All-Flash Storage Default Implementation



451 Research

Q. Does your organization implement all-flash storage technologies by default for all new block-based (i.e., SAN-like) deployments?

Source: 451 Research Voice of the Enterprise: Storage, Workloads and Key Projects, 2017

data for single applications. Now the majority of all-flash systems are used to store data for multiple applications.

In large-scale studies, flash drives have proved to be more reliable than disk drives, reducing the week-to-week labor overhead of replacing failed and standby drives in storage systems. In 2016, the University of Toronto released a paper based on a study of flash drives in Google datacenters. It concluded that in terms of replacement rates, flash drives beat disk by a significant margin. During the entire four-year study, 4-10% of Google's flash drives were replaced, versus an annual replacement rate for disks of 2-9%, as reported in other studies.

Despite the relatively short working lives of the early generations of flash devices, enterprise flash drives

now enjoy significantly longer working lives than disk drives, and carry five-year warranties as standard, in comparison with three-year warranties for disk drives. The short life of disk drives has been one of the bigger reasons why over the last two decades, enterprises have adopted a policy of replacing storage systems every three to four years. Replacing those systems is a very labor-intensive process for IT administrators, and the cycle of purchase and replacement is often called the 'refresh treadmill.' Moving to flash has reduced the speed at which that treadmill spins, not just because of the longer working life of flash drives, but also because of the greater performance headroom that flash systems can grow into as workloads increase over time.

The need for analytics software to access data rapidly exists even for high-end, in-memory analytics implementations, despite the fact that they may appear to rely purely on DRAM to store data that is being analyzed. Because of the high cost of DRAM, in-memory systems only load subsets of data into memory and store the majority of data outside of server memory.

Initially, the need for fast, low-cost storage for analytics was met by storing data on local disk or flash drives inside servers. This direct-attached storage (DAS) model helped meet the need for analytics to access data very quickly by eliminating the performance overhead of the network between the analytics servers and the storage systems. However, this was a return to an architecture that had been largely replaced over the last two decades by the storage area network (SAN) model, in which data is stored in centralized, shared storage systems, thereby heavily reducing administrative overhead.

The declining cost and improving performance of all-flash storage systems have allowed centralized and shared storage to replace DAS or local server storage in analytics. This frees datacenters from the

administrative overhead of DAS and considerably increases the potential impact of analytics by allowing a broader range of data to be readily analyzed. This is done by eliminating the need to load what can be very large volumes of data onto the local storage within analytics servers.

Part III: Why IBM For Flash?

Real users of flash memory arrays describe their experiences and insights on IT Central Station, which provides data center professionals the means to share pragmatic reviews on a wide range of enterprise tech solutions. A diverse group of users have published reviews of IBM flash storage systems on IT Central Station. Despite coming from multiple countries, industries and organization sizes, they share a common positive view of IBM flash storage solutions. The reviews show how companies have applied IBM flash storage solutions to improve application performance and operate more economically through increased storage density and ease of use. The reviewers also praise the scalability of the solutions, which is essential for today's organization that is dealing with unpredictable storage growth needs as they look to rapidly exploit data and insights for competitive advantage.

IT ECONOMICS

A storage solution's impact on finance is certainly a high-level concern for IT departments. Money is seldom far from an IT manager's radar screen. A [Solution Architect](#) at a tech services company commented on his IBM flash storage solutions in this context by noting, "Benefits for us include lower overall cost, higher density, lower data center costs, and reduced total TCO [total cost of ownership], as well as moving to an all-flash solution versus a flash/

spinning disk solution." A [Senior UNIX Systems Administrator](#) at a financial services firm with over 1000 employees stated, "This product has the best quality, performance and cost."

FLASH ADVANTAGES

An [Enterprise Architect](#) at a tech services company commented, "The FlashSystem 900 is perfect for providing pure speed and low latency." Or, as the [Head of IT infrastructure](#) at a financial services firm with over 1000 employees put it, "Our automatic banking systems have become twice as quick with flash storage." A [Senior Analyst](#) at a company with over 1000 employees praised his IBM solution by calling it "one of the best flash storage systems."

“

It reduces storage costs by compressing it up to 80%, therefore saving rack space and power consumption.”

Reviewers note that IBM's flash storage solutions offer the advantage of flexibility. A [CTO](#) at an IT consultancy said, "That idea of starting small with their current flash rollout and then being able to grow as the budget allows, has been very beneficial to clients." It's also compact, which is helpful in space-constrained situations. A [Senior Technical Specialist](#) at a tech services company praised his IBM flash storage solution's real-time compression, saying, "It reduces storage costs by compressing it up to 80%, therefore saving rack space and power consumption."

EASE OF SETUP

How easy or difficult it is to set up a storage solution can affect the cost and time spent on implementation.

Ease of setup is therefore praised by IBM flash storage users like an [IT Specialist](#) at a European commercial bank who said, “The initial setup we did ourselves because it is very simple.” An [Enterprise Architect](#) at a tech services company concurred, saying, “The speed and the ease of installation are the most valuable features.”

IBM’s approach to system deployment is reflected in reviews of the flash storage solution. A [Senior Solutions Architect](#) at a tech services company noted, “The implementation was very easy. It’s almost a “set it and forget it” kind of product.” He added, “Follow the instructions that IBM sets forth in the Redbooks and make life simple.”

EASE OF USE

With IT departments struggling with a deficit of people, time and money, complex, hard-to-use solutions are an unwelcome drag on productivity. With IBM’s flash storage solution this is not an issue - according to the [Enterprise Architect](#) at the tech services company, it is “very simple and straightforward.”

The [CTO](#) at the IT consultancy further explained, “The [FlashSystem] V9000 and the investment that IBM has made to the user experience, makes it very easy for our customers to be able to self-maintain and self-administer their own.” A [Senior Technical Specialist](#) at a tech services company added, “Ease of use and the GUI are the most valuable features of this product. It’s the same platform all over across the same product type.”

SCALABILITY

Storage solutions have to be able to expand along with requirements. As organizations evolve and grow, demand for storage can increase rapidly. Users of IBM flash storage solutions on IT Central Station are pleased with their solutions’ scalability. The [Enterprise Architect](#) at the tech services company

said, “It’s a good, scalable product with which you can scale out. And you can add it to virtualization. Whether it’s a [FlashSystem] V9000 or you add it to an SVC, it becomes even more scalable.” A [Storage Admin](#) at a financial services firm commented, “You can scale out as much as you want. We use up to four in a cluster, four to six nodes, so it’s working fine for us.”

SPEED AND PERFORMANCE

Users like flash because it’s typically very fast, especially for data intensive workloads. This speed is very welcome. As noted, analytics workloads need as much speed as possible, but spinning disks have reached their peak performance levels. Flash delivers speed and performance needed for analytics and other data intensive workloads. The [IT Specialist](#) at the European commercial bank thought the most valuable feature of his IBM solution was its high IOPS rate. He said, “I think the performance of most of the application is improving,” indicating that flash storage was enabling his applications to run faster.

A [Storage & Backup Transition Engineer](#) at a company with over 1,000 employees expressed a similar sentiment, noting “It facilitates those applications that need low latency and high speed for transferring the data.” The [Senior UNIX Systems Administrator](#) at the financial services firm likes his IBM flash storage solution’s “latency and speed of the array,” as well as “its ease of management.”

CONCLUSION

Flash is already rapidly revolutionizing datacenter storage by reducing TCO while increasing flexibility and reliability, as well as dramatically improving performance. While all of those benefits apply almost universally to all workloads, the ability of flash to transform performance is of even greater value for analytics than it is for other applications. Indeed, without flash, the full benefits of analytics cannot be realized. For analytics, flash is key to realizing the competitive business benefits delivered by the new generation of analytics tools. Flash not only allows analytic insights to be delivered faster, but also allows those insights to be deeper. The reviews submitted to IT Central Station confirm the suitability of IBM's flash storage systems for this task, across the key parameters of cost, ease of deployment and use, scalability, and performance.

ABOUT IT CENTRAL STATION

User reviews, candid discussions, and more for enterprise technology professionals.

The Internet has completely changed the way we make buying decisions. We now use ratings and review sites to see what other real users think before we buy electronics, book a hotel, visit a doctor or choose a restaurant. But in the world of enterprise technology, most of the information online and in your inbox comes from vendors but what you really want is objective information from other users. IT Central Station provides technology professionals with a community platform to share information about enterprise solutions.

IT Central Station is committed to offering user-contributed information that is valuable, objective and relevant. We validate all reviewers with a triple authentication process, and protect your privacy by providing an environment where you can post anonymously and freely express your views. As a result, the community becomes a valuable resource, ensuring you get access to the right information and connect to the right people, whenever you need it.

www.itcentralstation.com

ABOUT 451 RESEARCH

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

www.451research.com