



Sponsored by:
Pure Storage

Authors:
Harsh Singh
Eric Burgener

April 2020

Business Value Highlights

290%
three-year ROI

73%
reduced total cost of operations

4 months
to payback

57%
more efficient IT storage staff

4x
more staff time spent on innovation-related activities

156%
more usable terabytes of flash storage

54%
improved storage array deployment agility

43%
more efficient DBAs

Pure Storage Provides High-Performing Storage Platform for Optimizing Oracle Database Workloads

EXECUTIVE SUMMARY

Over the past eight years, all-flash arrays (AFAs) have come to dominate primary storage spend in the enterprise. Given their ability to deliver consistent performance at scale, they are commonly used to host latency-sensitive, online transaction processing workloads that rely on relational databases (DBs) like Oracle. Pure Storage is a leader in the AFA market, and the vendor's influence over the past decade has helped mold an enterprise storage experience that is very different from that which has historically been provided by established enterprise array vendors. Pure Storage targeted the relational database market (among others) early on with their FlashArray and has built up a considerable installed base of Oracle customers. The vendor's current flagship array, the FlashArray//X, offers a number of features that optimize the performance, availability, and manageability of Oracle environments.

IDC interviewed a cross-section of enterprise customers to understand the impact of running Oracle database applications on the Pure Storage platform. These customers reported that Pure Storage serves as an efficient and cost-effective platform while delivering the performance, agility, and reliability required to optimize Oracle applications. IDC calculates that study participants will realize value worth an annual average of \$2.23 million per organization (\$60,700 per database) by:

- Fostering more efficient storage and database administrator (DBA) teams and allowing them to spend more time on innovation and business-related projects
- Improving the agility required for deploying and managing storage arrays and providing more usable flash storage availability

71%

improved agility to deploy
new databases

85%

faster database query speed

- Improving the efficiency of DBA operations, leading to quicker and more agile database deployment and faster query speeds
- Reducing the total cost of operations

These improvements are driven not only by the fact that the FlashArray uses only solid state storage but also by the vendor's deep integration with Oracle-specific tools to enable better automation, a control plane (Pure1) that leverages artificial intelligence and machine learning (AI/ML) to self-manage storage to operator-defined policies and simplify administration, a highly available design that offers flexible "defense in depth" data protection, and efficiency features like inline compression and deduplication that deliver high data reduction ratios to maximize flash capacity utilization.

SITUATION OVERVIEW

Oracle is the undisputed market share leader in the relational database market. In 2018, the company grew its number 1 market share by 5.4%, generating \$14.77 billion in revenue. Most enterprises, particularly if they have been in business for a while, have Oracle databases supporting mission-critical workloads. The performance and availability of these databases can have a critical impact on not only an enterprise's customers but also on businesses' ability to provide responsive customer service, rapidly deliver new products, and perform a variety of internal processes in a timely manner. Managing relational databases and their underlying storage infrastructure to meet these business requirements largely falls to two constituencies: database administrators and storage administrators. These two constituencies do not always see eye to eye on issues that impact storage infrastructure decisions.

DBAs are responsible for a number of tasks that are done at the database (rather than the storage) level. These include database product selection, database design, deployment, configuration, and ongoing management with a focus on the database and the applications it supports. What they need from the underlying storage infrastructure includes consistently predictable performance, which can be delivered to meet defined service-level agreements, availability (along with resilience and simple, rapid recovery capabilities), ease of provisioning (to easily enable expansion of existing as well as new projects), and the ability to rapidly clone databases for analytics, test and development, and other operational needs. Traditionally, DBAs have had to interact with storage administrators to ensure they get the things they need from the underlying storage infrastructure through formal requests, and this interaction has often been viewed by DBAs as time consuming, frustrating, and inefficient.

Storage administrators, on the other hand, have tended to be one step removed from the various applications (including relational databases) and have been focused on making the storage infrastructure itself capable of high performance, able to be configured to meet necessary availability levels, creating resilient infrastructure, recovering data in the event of data loss and/or corruption, creating copies of underlying data when it is needed, and managing other lower-level functions, like data reduction, snapshot creation, and encryption, as needed to meet business and/or regulatory requirements. While the cost of the storage infrastructure can potentially be a concern of DBAs when charge-back may be in place, creating a cost-effective storage infrastructure that meets all business requirements has historically been much more of a storage administration (rather than a DBA) concern.

In the past decade, there have been two major evolutions in enterprise storage product strategies among vendors that have simultaneously improved the lives of both DBAs and storage administrators significantly. The first of these was the introduction of all-flash arrays (AFAs) — enterprise storage systems that used solid state storage as a persistent tier. Until that time, most primary storage arrays servicing block-based, latency-, and availability-sensitive workloads were using hard disk drives (HDDs).

AFAs were built using flash memory, a storage medium that delivered much lower latencies, much higher throughput and bandwidth, and a much more agile ability — compared with HDDs — to handle I/O spikes in real time while continuing to provide predictably consistent low latencies, particularly under load. While HDD-based arrays were delivering latencies that could range from 5 to 20ms, AFAs could consistently deliver sub-millisecond latencies under load. As nonmechanical devices, solid state disks (SSDs) could be more reliable than HDDs and were not subject to the same shock and vibration issues that HDDs were. Flash storage was also more efficient than HDD-based storage, enabling the creation of higher density devices, which consumed less power than HDDs on a watts per terabyte (TB) basis.

Although in the 2000s several vendors were experimenting with solid state storage as caches, there were two issues with the use of flash memory as an enterprise-class persistent storage medium — cost and endurance. The performance advantages of flash relative to HDDs were clear for latency-sensitive workloads, but few in the industry felt that the cost of such a system would not be prohibitive for mainstream usage. Starting in the 2010s, two new technologies were employed in enterprise array design to address these concerns — lower cost consumer-class flash memory and inline data reduction technology.

Because it had originally been developed for HDDs, the SCSI storage protocol that was in broad usage in enterprise arrays was not very efficient at writing to flash memory, and the endurance of the medium in write-intensive enterprise environments was a significant concern. Algorithms were developed in the Flash Translation Layer (FTL; which sat between

SCSI and the actual flash devices) to modify how data was actually written to flash, significantly improving the endurance (as well as the performance) of the underlying medium. And while data reduction technologies like compression and deduplication were already in use, HDDs were too slow to allow them to be implemented “inline” without significantly impacting application performance. For that reason, these technologies were not widely used against latency-sensitive primary workloads. The much lower latencies of flash enabled vendors to develop inline data reduction that could be used with these types of workloads, a feature which significantly cut the \$/GB cost of SSDs.

Customers that purchased AFAs noticed a number of pleasant surprises when using them in performance-sensitive database environments. DBAs and end users noticed the order of magnitude improvement in latencies. The lower latencies translated to higher CPU utilization in application servers, which translated to needing fewer servers to meet a given level of performance. This in turn translated into needing fewer Oracle licenses, driving potentially significant cost savings. Running on flash, virtual machines (VMs) could now offer sufficient performance to enable database servers to be virtualized, improving IT infrastructure agility, improving efficiency, and lowering costs. The ability of flash to handle I/O spikes without sacrificing latency meant that laying databases out on storage was faster and easier and hitting even stringent latency requirements required little if any database performance tuning. Higher flash throughput also meant that databases could be moved significantly faster, whether that was for recovery or other DBA workflows.

In addition, with flash performance, storage administrators could deploy all sorts of useful storage management services at the array level without undue impact on database performance — features like thin provisioning, compression, deduplication, RAID, snapshots, encryption, and replication that protected data better made copy creation almost immediate without giving up performance and improved availability and security. Almost overnight, application performance problems due to storage disappeared (although there were still other issues outside of storage which could impact it). AFAs now account for almost 80% of all primary enterprise storage revenue, deliver far better performance for latency-sensitive workloads like Oracle databases than HDD-based arrays, enable the use of array-based data services in performance-sensitive environments, have streamlined storage infrastructure by increasing storage densities and enabling performance and capacity requirements to be hit using far fewer devices, and have increased the efficiency of IT infrastructure by lowering energy and floorspace consumption. AFAs enable much higher infrastructure density and lower total cost of ownership (TCO) relative to HDD-based arrays for performance-sensitive workloads.

The second major evolution was a shift away from very human-intensive storage administration toward self-managed storage. Newer vendors developed their storage operating systems and management tools under the assumption that storage administration would over time be migrating away from dedicated, highly trained storage administrators to IT generalists (like virtual or Linux administrators). This is in fact happening. Vendors have simplified traditional storage workflows, including initial array deployment, provisioning, snapshot creation and usage, data protection and recovery, and system expansion. Vendors have added policies that ensure storage administration remains in compliance with established governance as these workflows are performed while they also streamline and improve the reliability of operations. Vendors have added automation APIs to their management consoles and have qualified their systems with various orchestration tools that enable automated workflows to span different vendors' products. As application deployment strategies move from virtual machines to containers, integration with Kubernetes-based orchestration frameworks will become increasingly important.

Recently, vendors have begun to leverage artificial intelligence and machine learning to dynamically optimize systems to more cost effectively meet performance and availability requirements. Cloud-based predictive analytics tools leverage these same technologies to monitor system health and performance, proactively remove failing devices before they can cause a problem, troubleshoot problems faster, pre-validate upgrade options, aid in performance and capacity planning, optimize workflows, and tier data between available storage media — from storage-class memory to flash memory, HDDs, and cloud — to help lower costs and streamline management requirements. With these technologies in use, manual performance tuning is largely a thing of the past as administrators establish performance targets using quality of service (QoS) tools at the application level and systems self-manage to hit them even as workloads evolve over time.

As the control planes for IT infrastructure become more business objective oriented, they clearly become easier and more intuitive to use. Combined with role-based access control, these updated management interfaces can in fact be used in many cases to allow DBAs to directly perform certain recurring tasks like provisioning or database cloning safely, reliably, and in accordance with IT governance. Other enterprises that establish private clouds can offer internet-based service portals that can achieve the same thing. These developments have been welcomed by DBAs and storage administrators alike — DBAs get resources they need from storage much faster (and in some cases without even having to interact with IT) to keep their projects moving along, and storage administrators are freed up from these types of mundane chores without having to worry about risk or IT governance issues, allowing them to focus on more strategic tasks.

PURE STORAGE OVERVIEW

Pure Storage is a \$1.6 billion enterprise AFA vendor with a broad portfolio of all-flash offerings that cover primary and secondary as well as structured and unstructured workloads. The FlashArray//X line covers primary storage, the FlashArray//C addresses tier 2 and other secondary workloads, and FlashBlade covers unstructured (file/object) workloads. Over the past eight years, Pure Storage has introduced many features and programs that have differentiated its offerings from those of the legacy enterprise storage providers and driven real changes in the industry, including pioneering the use of AFAs for enterprise workloads, launching the industry's first purpose-built AFAs for unstructured data workloads, and introducing broad-ranging guarantee programs that have significantly changed industry expectations around customer experience (CX).

With respect to the first major evolution discussed previously (the move to AFAs for primary storage), Pure Storage spearheaded that trend. As a start-up founded in 2009, Pure Storage was formed with the idea of creating a true enterprise-class storage array that used only flash as persistent storage. The vendor harnessed two strategic choices — the use of less expensive consumer-grade flash (whose volume production would allow it to ride an aggressively downward-sloping cost curve) and the use of inline data reduction technology — to narrow the effective dollar-per-gigabyte cost difference between SSD- and HDD-based arrays. The lower endurance of the consumer-grade flash was addressed with an optimization of the flash translation layer in the vendor's arrays so that even write-intensive enterprise workloads could be reliably handled by flash media. And through the use of optimized data reduction, which included compression, deduplication, and other more proprietary algorithms developed by the vendor, they were able to provide data reduction ratios of up to 5:1 or more for latency-sensitive workloads, thereby cutting the effective cost of an SSD by 5x relative to an HDD. Once this issue was addressed, the performance of SSDs (rather than HDDs) meant that storage performance was no longer the “usual suspect” it had been when end users complained about application performance. Many administrators, using AFAs for the first time, noted that time spent dealing with storage performance tuning tasks dropped to almost zero.

Ease of use had always been a concern with enterprise storage, and many of the start-ups developing enterprise-class arrays in the 2010s were building their storage operating systems in a very different manner from the legacy systems of old. Higher performance CPUs and storage devices, the use of policy-driven administration rather than task-driven administration, and broader use of orchestration and automation tools, when coupled with newer technologies like AI/ML, all allowed new storage vendors like Pure Storage to develop enterprise-class arrays that just required less manual involvement on a daily basis to meet business requirements, even as workloads evolved. Based on this design approach, the

Source: IDC, 2020

vendor's arrays are not only much easier to manage than legacy systems but the combination of all flash with self-managed storage allowed many operations, such as provisioning, snapshot creation, failed device replacement, and capacity expansion, to occur in just minutes.

Readers will note from the primary research results that the interviewed users of Pure Storage FlashArrays in Oracle database environments enjoyed a 43% improvement in DBA productivity and a 57% improvement in storage management productivity (for storage administrators). For DBAs, database agility improved by 71%, whereas for storage administrators it improved by 54%. These improvements are a result of the combination of all flash with an array built around a self-managed storage design.

The Pure Storage FlashArray//X

The FlashArray//X array portfolio spans from the entry-level X10 up to the X90. Based on a dual controller array design, the systems use all NVMe storage devices internally. They support both NAND flash-based storage devices and storage-class memory-based options, which allow the system to deliver latencies as low as 150µs under load.

The storage operating system, called Purity, features a comprehensive set of enterprise-class data services including inline storage efficiency technologies (compression, deduplication, and thin provisioning), dynamic multipathing, dual parity RAID, space-efficient snapshots, always-on quality of service (Purity Optimize), encryption (FIPS 140-2 certified), and various replication (synchronous, asynchronous) options including support for stretch cluster configurations (ActiveCluster). The highly available design includes RAID-protected hot-pluggable storage devices; redundant, hot-pluggable power and cooling; transparent recovery from component failures; and can support "six-nines plus" availability requirements (when deployed in ActiveCluster configurations). The systems support both SCSI and NVMe over Fabric host connections using either Fibre Channel or Ethernet. The FlashArray's highly efficient design can house up to 3.3PB (assuming a 5:1 data reduction ratio) in only 3U and requires only 6 cable connections when a system is initially deployed.

Pure Storage does not use commodity off-the-shelf (COTS) SSDs. Instead, it has built custom memory modules, which it calls DirectFlash Modules (DFMs). Available at multiple capacity points (including a 57.8TB DFM), these storage devices provide higher performance and endurance with higher density, better efficiency, and lower cost than COTS SSDs. The raw flash on the DFMs is managed by Purity at the system level, dispensing with the device-level approaches in COTS SSDs that, for enterprise workloads, are just less efficient.

When new flash media becomes available (such as the quad-level cell flash that became available in 2019), Pure Storage can offer storage devices built from that medium faster than

its competitors that have to wait for component suppliers to develop and make available storage devices based on it. The DFMs give Pure Storage systems based on it (FlashArray and FlashBlade) significant relevant differentiation in the areas of performance, endurance, density, efficiency, and cost.

Pure Storage also features extensive cloud integration capabilities, including CloudSnaps (snapshots that are easily portable between on- and off-premises locations), support for cloud-based tiering, data reduction for the cloud as well as data in-flight (Purity Reduce), RESTful APIs for integration with VMware, Amazon Web Services, Microsoft Azure, and OpenStack, Container Storage Integration (CSI) support, the ability to run Purity directly in the public cloud (Cloud Block Store), and a single control plane for Pure Storage environments in both on- and off-premises (i.e., cloud) locations (which includes the vendor's AI/ML-driven cloud-based predictive analytics platform Pure1).

All Pure Storage AFAs can be purchased with a subscription to the vendor's Evergreen Storage program. A key pillar in the highly regarded CX the vendor delivers (turning in an industry-high Net Promoter Score [NPS] of 82), the Evergreen program includes not only 24 x 7 support but also a money-back guarantee, all-inclusive array software bundling, a right-size guarantee (that ensures your FlashArray//X will have the promised effective capacity), fixed maintenance pricing over the life of the system, free controller upgrades, and guaranteed trade-in credits to defray the cost of controller upgrades or capacity consolidation. Free controller upgrades and guaranteed trade-in credits require an active, eligible Evergreen subscription. Controller upgrades occur in an entirely nondisruptive manner with these systems.

Evergreen is an industry-leading CX program that has forced Pure Storage's competitors to respond, and it drives real value for customers. For more detailed information on the Evergreen Storage program, see *Pure Storage's Evergreen Storage Continues to Be the All-Flash Array Market's Customer Experience Program to Beat* (IDC #US45761920, January 2020). For more information on just what NPS means to customers, see *Net Promoter Score Becoming an Important Metric for Enterprise Storage Managers to Understand* (IDC #US43896818, June 2018).

THE BUSINESS VALUE OF PURE STORAGE

Study Demographics

IDC conducted research that explored the value and benefits of using the Pure Storage platform for Oracle database workloads. The project included seven interviews with organizations that were using this solution and had experience with or knowledge about its benefits and costs. During the interviews, companies were asked a variety of quantitative and

qualitative questions about its impact on their IT infrastructure, DBA and storage operations, core businesses, and cost profiles.

Table 1 presents study demographics and profiles. The organizations interviewed had a base of 7,968 employees. This workforce was supported by an IT staff of 376, managing the use of 197 business applications. The total overall storage capacity was 2,475TB. Most of the companies surveyed were based in the United States with one based in the Dominican Republic. These organizations hailed from a wide range of vertical markets, including the financial services, education, government, insurance, real estate, and travel/hospitality sectors. Additional metrics are presented in Table 1. (Note: All numbers cited represent averages.)

TABLE 1 Firmographics of Interviewed Organizations

Firmographics	Average	Median	Range
Number of employees	7,968	2,000	1,200 to 31,000
Number of IT staff	376	360	75 to 700
Number of IT users	7,968	2,000	1,200 to 31,000
Number of external customers	927,500	850,000	10,000 to 2 million
Number of business applications	197	175	40 to 500
Total overall storage (terabytes)	2,475	1,000	25 to 10,000
Revenue per year	\$3.51 billion	\$3.35 billion	\$350 million to \$6.9 billion
Countries	United States (6) and Dominican Republic		
Industries	Financial services (2), education, government, insurance, real estate, and travel/hospitality		

Source: IDC, 2020

Choice and Use of Pure Storage

The companies that IDC surveyed described usage patterns for the Pure Storage platform as well as providing relevant data about their overall IT and business environments. They also discussed the rationale behind their choice of Pure Storage, citing a number of factors for their choice such as the need for better performance latency, the desire to move to a more automated environment as part of their DX efforts, and improvements in deduplication capability. Study participants commented on these benefits:

- Improved ROI and ease of use:** *“The main reason to deploy Pure Storage was to achieve economies of scale. As an organization, we needed to know we had good value for what we purchased. We found that the ROI was not always good on the most expensive product. We did a lot of homework and what we liked about Pure Storage was ease of use. It’s also really fast. The UI is great, and the service has been really good.”*

- Better support for Oracle workloads:** *“Oracle workloads take up a lot of resources. We did not want to have noisy neighbor syndrome or have the Oracle system stopping everything else on the array. Provisioning is faster with Pure Storage. The expansion capabilities allow us to expand what is there and not have to spin up new storage. Snapshotting is vastly superior to spinning disk. Deduplication capabilities for storage efficiencies are much better than what we can get from spinning disk. What Pure Storage offers is what we needed. We did proof of concept with two other vendors and Pure came out on top.”*
- Helped with digital transformation (DX):** *“Our digital transformation process began last year to improve efficiencies and customer service. As part of the transformation, we wanted to move from traditional storage. We were looking for flash storage and Pure Storage was chosen. Pure has a lot of features for improving automation and supports our company’s move to more technical and digital transformation. We wanted to replace the consultant resources and do it ourselves. Pure Storage supports that effort with its tools for proactive monitoring.”*
- Needed better latency performance:** *“We needed higher performance for latency-sensitive business-critical applications, mostly databases and virtual servers. We also needed the additional flexibility that all flash provided for our overall automation project. System administration in the virtualized environment worked much better with all flash.”*

Table 2 describes typical organizational usage associated with the use of Pure Storage for these companies. As shown in Table 2, there was a substantial Pure Storage footprint consisting of 125TB, including 3.3 in-memory terabytes. In terms of Oracle database usage, interviewed companies had 37 databases running 196 applications. This Pure Storage infrastructure played a strong role in the business operations of these companies by supporting 68% of all revenue. Additional metrics are presented in Table 2. (Note: All numbers cited represent averages.)

TABLE 2 Organizational Usage of Pure Storage

Pure Storage Use	Average	Median
Number of terabytes	125	30
Number of in-memory terabytes	3.3	1
Number of arrays	2.3	2
Number of databases	37	25
Number of applications	19	15
Number of internal users supported	5,357	1,250
Percentage of revenue being supported	68	100

Source: IDC, 2020

Business Value and Quantified Benefits

IDC's Business Value model expresses the benefits for organizations using the Pure Storage solution to support their business-critical Oracle database applications and workloads. Survey data obtained from Pure Storage customers was applied to this model to arrive at an array of quantified post-deployment benefits. Using this methodology, IDC found that these customers realized significant value for their storage, database, and business operations.

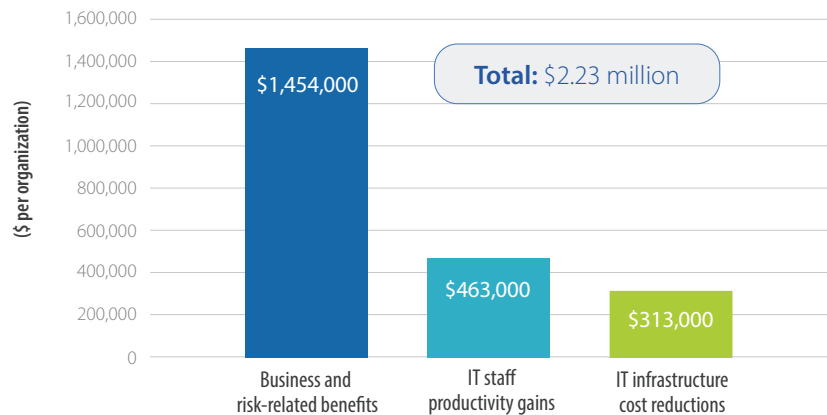
Interviewed organizations reported that use of the Pure Storage platform fostered more efficient storage and DBA teams, thereby allowing them to spend more time on innovation and business-related projects. It also improved operations that involved deploying and managing storage arrays and increased flash storage availability while reducing the total cost of operations. The platform also improved the efficiency of DBA team operations, leading to quicker and more agile database deployment and faster query time. Study participants described these benefits in greater detail:

- **Spinning up new database environments quicker:** *"I get to sleep at night because of the ease of setup, training, and onboarding of new administrators. Instead of a week of training, it's done in a single day or less. Provisioning is simple with Pure Storage. For the people doing disaster recovery, we can revert back to snapshots as needed. There is no restore necessary from backup; the snapshots in time handle it. Time to resolution of data loss is dramatically improved ... it was two to three hours before, but with Pure it takes less than five minutes ... In addition, no paperwork is needed ... now it is done by email. A DBA can make a request, get an email approving the request, and do it themselves."*
- **Easier management:** *"The really big win with Pure Storage was to automate a lot more of our processes. We were able to expand virtualization and maintain the performance we needed. System and DBA administration was more effective. Pure Storage was also just faster for IT activities. Pure Storage, and all flash in general, is much simpler to manage. Now that our organization has moved from the previous technology, things are a lot easier to manage."*
- **Can deploy new databases more cost effectively:** *"One of the greatest benefits ... is flexibility in terms of needing new landscapes to do new development. The flexibility of instantly creating a new database landscape is there. Of course, cost-wise, if I were to buy a new database subsystem, the cost is approximately one-fourth overall."*
- **Strong performance improvements:** *"Pure Storage is rock solid. It just works. We have put some other non-Oracle things on the array, and it races along. When our business side with its own application monitoring comes to us with an issue, we are always certain it's not a storage issue. In the past, we would have to prove it was not the storage. Any larger-scale downtime for this application will not be related to storage. And it translates to less time troubleshooting."*

Based on interviews with seven Pure Storage customers, IDC quantified the value they will receive over five years at an annual average of \$2.23 million per organization (\$60,700 per database) in the following areas as shown in Figure 1:

- **Business and risk-related benefits:** The improved agility and performance of Pure Storage environments enabled interviewed organizations to run more effective business operations through database optimization. In addition, reducing unplanned downtime helped these organizations minimize their business risk. IDC quantifies these benefits as worth an average of \$1.45 million per organization (\$39,600 per database).
- **IT staff productivity benefits:** IT storage and database infrastructure teams benefited from increased efficiencies in their routine operations with the automated self-management functionality of the Pure Storage solution. IDC puts the value of these benefits at an annual average of \$463,000 per organization (\$12,600 per database).
- **IT infrastructure cost reductions:** Operating a more efficient storage and database infrastructure reduced ongoing operational costs for Pure Storage customers. IDC projects that study participants will save an annual average of \$313,000 per organization (\$8,500 per database).

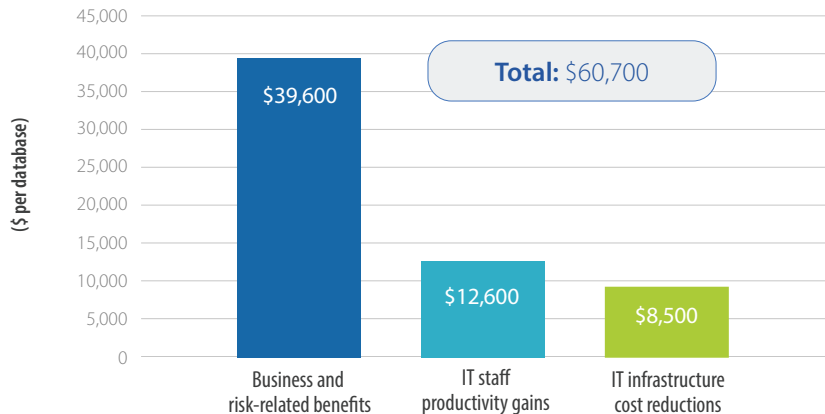
FIGURE 1 Annual Average Benefits per Organization



Source: IDC, 2020

These same quantified benefits are presented in Figure 2 and are calculated per database.

FIGURE 2 Annual Average Benefits per Database



Source: IDC, 2020

Improvements in Storage Operations and Management

In today's IT environments, DBAs have major responsibilities for protection, availability, and performance of company data — all of which are foundations for business growth. However, fulfilling the required tasks to maintain storage infrastructure that supports database operations falls outside the parameters of DBA management and is the direct responsibility of storage infrastructure teams. For these teams, traditional storage infrastructure poses its own set of challenges for managing storage tiers, RAID groups, LUNs, and other key elements.

The Pure Storage platform is designed to alleviate this real-world conundrum. Study participants spoke to IDC about how Pure Storage served to cost effectively improve the efficiency of their Oracle database operations and cited valuable features and capabilities such as better ability to support business needs, improved data compression functionality, and easier replication and disaster recovery. They commented on these and other benefits:

- More efficient IT operations:** *"We are seeing IT benefits across all of those areas: savings, efficiencies, agility, and high availability. We also see that resourcing and processing time is better. For example, a job that would run for 4 hours now takes 30 minutes. In terms of IT efficiencies and flexibility, with Pure Storage, replication is easy, disaster recovery is easy, and migration between controllers is easy."*
- IT is more cost effective:** *"What the IT organization is seeing is that cost is a factor. In IT, whatever we do, efficiency matters. The cost savings are incredible using Pure Storage, but at the same time, the flexibility and the amount of time spent by DBAs and the system administrator team to provision has gone down significantly."*

- **IT can better support business needs:** *“When one of the applications owners says: ‘I need this new feature now,’ guess what happens? You have a parallel build to do in terms of the application. Guess what happens when you talk about parallel builds. You need a new database landscape to be created. Now, with snapshots and zero-copy clone, that’s almost instantaneous ...”*

Interviewed companies reported that Pure Storage makes it easier for storage teams to manage their environments. The platform’s self-management orchestration tools enable automated workflows that simplify traditional tasks such as array deployment, provisioning, and system expansion. Table 3 quantifies these impacts on storage team efficiency. As shown in Table 3, IT storage management productivity, measured in FTEs per organization, decreased from 6.2 to 2.7, representing a substantial improvement (57%). This translated into an annual salary savings of \$357,000.

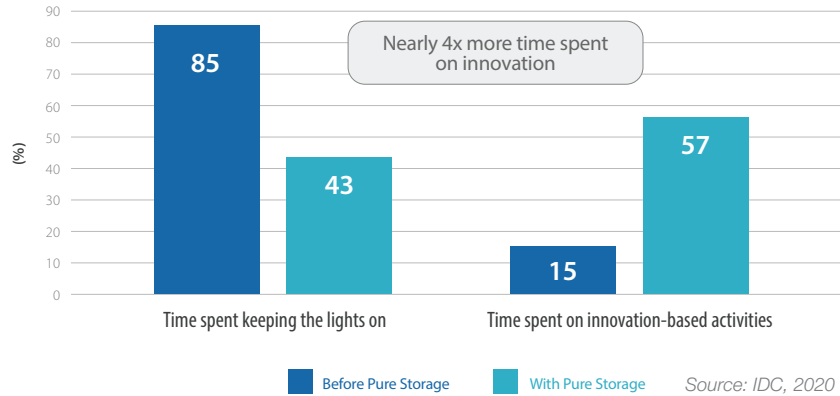
TABLE 3 IT Storage Management Impact

	Before Pure Storage	With Pure Storage	Difference	Benefit (%)
Management of IT storage productivity impact (equivalent FTEs)	6.2	2.7	3.6	57
Salary cost per year per organization	\$625,000	\$267,000	\$357,000	57

Source: IDC, 2020


Interviewed companies found that their storage teams were able to focus less on keeping the lights on and more on strategic- and business-related activities. As one study participant commented: “We are seeing time savings because of snapshot technology. We can replicate a database without taking up space. We also see benefits from compression and deduplication. Pure Storage screens provide visibility and monitoring from one location. Our IT teams can focus on learning the Pure Storage system and supporting the transformation of processes across the IT organization and the company.” Figure 3 shows further post-deployment impacts on IT storage team activities and indicates that 50% less time was spent on routine maintenance. As a result, storage teams were freed up to spend nearly four times more time on business-oriented innovation.

FIGURE 3 Impact on IT Storage Team Activities




Implementation of the Pure Storage platform has improved the agility of deploying new storage resources. AFAs are easier to manage than legacy systems, and the combination of all-flash with self-managed and automated functionality allowed operations such as provisioning, snapshot creation, and capacity expansion to be carried out more quickly and efficiently. As one study participant noted: “On the IT side, we have built automation around the Pure arrays. If more space is needed, zero touch is required from the storage team.” Figure 4 quantifies these agility impacts. As shown in Figure 4, the number of days required to deploy a new storage array showed significant reduction (54%).

FIGURE 4 IT Agility Impact



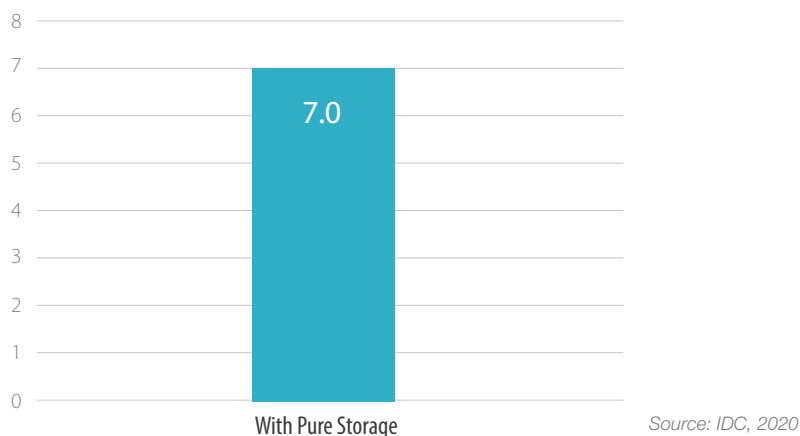
Study participants reported that Pure Storage increased the availability of usable flash storage for their organizations. Figure 5 quantifies this benefit. As shown in Figure 5, after deployment, the amount of available flash-based storage was significantly increased from 18.6TB to 47.7TB, a capacity improvement of 156%.

FIGURE 5 Flash Memory Usage Impact




Another key area affecting interviewed companies is data compression. Survey results showed that Pure Storage customers are seeing higher data compression rates in part because the platform has been designed to maximize capacity while minimizing the number of writes required to store unique data, resulting in increased storage efficiency and reduced costs. Inline compression uses less space than original formats, while append-only write layout and variable addressing removes wasted space introduced by fixed-block architectures. As one study participant commented: “We’re getting a lot more bang for the buck. Because of compression, we don’t need to expand capacity on an annual basis. Every year I put in for a budget to purchase more arrays if necessary and we have not had to do that.” As shown in Figure 6, after deployment of the platform, data compression was improved by a factor of 7:1. It should be noted that Pure Storage advises customers that compression rates will vary based on each customer’s usage and workloads.

FIGURE 6 Data Compression Impact



Survey results confirmed that, in addition to the operational benefits described, Pure Storage represented a cost-effective solution with solid price/performance. The AFA-based platform increased IT infrastructure efficiency by lowering energy and floorspace consumption, thereby lowering TCO relative to HDD-based arrays used for performance-sensitive workloads. Figure 7 presents data on the infrastructure savings realized on a per database basis over a five-year period. As shown in Figure 7, the cost of storage required to support a single database is projected to decrease from \$72,100 to \$39,100 over five years, representing a 46% savings.

FIGURE 7 IT Infrastructure Savings per Database (Five Years)



Optimizing Database Performance and Operations

As described previously, AFA solutions such as the Pure Storage platform deliver better performance for latency-sensitive workloads like Oracle databases than HDD arrays. In part, this is because they can streamline storage infrastructure by increasing densities and meet performance and capacity requirements using fewer devices. Study participants spoke about how, with Pure Storage, DBAs could more easily and quickly get the resources they needed from storage teams and how both storage administrators and DBAs were freed up from routine tasks to work on strategic projects. Study participants cited valuable features and capabilities such as improved I/O and query performance, significant reductions in latency, and the ability to virtualize application server resources. Study participants commented on these and other benefits:

- Oracle-specific tools enable better automation:** *“Pure Storage allowed us to virtualize all Oracle DB servers, including production. It created large new potential efficiencies through automation, because of the virtualization. Pure Storage’s Oracle-specific automation tools (OCAT) and organizational expertise with Oracle hugely supported and accelerated the benefits we could achieve by implementing that automation.”*

- **Improved database server virtualization and performance:** *“The big win was virtualizing all of our Oracle database and CRM/app servers, including production. This allowed the retirement of legacy spinning-disk arrays. From our first benchmarks, our Oracle database/apps environment responded on average twice as quickly once moved onto Pure Storage. All end-user request latency was halved on average. The value to our organization was that everything could be done twice as fast.”*
- **Stronger performance:** *“Performance was a key upgrade ... in making the transition to Pure Storage. Performance, particularly query speed, has been impressive. We have applications that require a large amount of I/O to help the departments we support. This product does very well with that.”*
- **DBAs have time to work on more strategic projects:** *“For DBAs that handle backup and recovery, if we needed to recover and create a clone, it used to take hours. Now it’s minutes. With the time saved, I’m not cutting down my DBA staff but I’m using the remaining time to do innovative things like automation in the database world. For example, they can work on effective code reviews for our application people.”*

The increased productivity of database administrators is shown in Table 4. Annual DBA FTE values were reduced from 2.7 to 1.5, representing a 43% improvement. This translated into an annual staff time cost savings per year of \$116,000.

TABLE 4 Database Administrator Impact

	Before Pure Storage	With Pure Storage	Difference	Benefit (%)
Database administrator (FTE equivalent per organization per year)	2.7	1.5	1.2	43
Staff time cost per year	\$269,000	\$153,000	\$116,000	43

Source: IDC, 2020

Study participants reported that the Pure Storage platform with its Oracle-specific orchestration tools and automated workflows increased agility by simplifying traditional storage workflows such as initial array deployment, provisioning, snapshot creation, and system expansion. Table 5 quantifies these benefits. As shown in Table 5, the time to deploy new databases was significantly reduced from 2 days to 0.6 days, representing a 71% improvement.


TABLE 5 Database Agility Impact

	Before Pure Storage	With Pure Storage	Difference	Benefit (%)
Time to deploy new database (days)	2	0.6	1.4	71
Staff time to deploy new database (hours)	17	6.4	10.6	63

Source: IDC, 2020

The Pure Storage platform helped improve the performance of Oracle database operations in a number of ways when measured by various KPIs such as query speed and transaction rates. Companies reported significant levels of improvement in the four KPIs that were evaluated. As shown in Figure 8, the greatest improvements were realized with query speed (87%) and backup processes (67%).

FIGURE 8 Database Performance Impact



Source: IDC, 2020 (% improvement)

Better overall performance for storage systems and the database applications they supported has had positive implications for reducing the incidence of unplanned downtime. Interviewed companies reported that, after deployment, the amount of Oracle database downtime was reduced by 95%, significantly mitigating episodes of business interruption. In addition, the annual frequency of events was reduced from 5.3 to 1.3, a 75% improvement. In the aggregate, these improvements yielded an annual business value of \$1.48 million (see Table 6).


TABLE 6 Unplanned Downtime Impact

	Before Pure Storage	With Pure Storage	Difference	Difference (%)
Frequency per year	5.3	1.3	3.9	75
Time to resolve (hours)	12.5	2.7	9.8	78
FTE impact (lost productivity due to unplanned outages)	22.4	1.2	21.2	95
Hours lost per user	7.9	0.4	7.4	95
Value of lost productivity per year	\$1.57 million	\$85,300	\$1.48 million	95

Source: IDC, 2020

Study participants reported that Pure Storage served as a cost-effective storage platform for their Oracle workloads. Having the benefit of operating a converged and efficient IT platform reduced ongoing IT infrastructure-related costs for maintenance, power, and floor space and compared favorably with legacy or alternative solutions while, at the same time, delivering greater performance. In addition, Pure Storage systems accommodated growth in a cost-effective way because additional resources could be quickly and efficiently deployed. Figure 9 shows the cost of operations per database over a five-year period including the cost of lost productivity, unplanned downtime, and IT staff management. As shown in Figure 9, aggregate costs of the Pure Storage solution were \$107,600 compared against the cost of alternative or legacy solutions, representing a 73% reduction.

FIGURE 9 Cost of Operations per Database (Five Years)



Source: IDC, 2020

ROI Summary

IDC's analysis of the financial and investment benefits related to study participants' use of the Pure Storage solution is presented in Table 7. IDC calculates that, on a per organization basis, interviewed organizations will achieve total discounted three-year benefits of \$5.61 million based on IT staff efficiencies, increased business results, and lower costs as described. These benefits compare with projected total discounted investment costs over three years of \$1.44 million on a per organization basis. At these levels of benefits and investment costs, IDC calculates that these organizations will achieve a three-year ROI of 290% and break even on their investment in four months.

TABLE 7 Three-Year ROI Analysis

Three-Year ROI Analysis	Per Organization	Per Database
Benefit (discounted)	\$5.61 million	\$152,800
Investment (discounted)	\$1.44 million	\$39,200
Net present value (NPV)	\$4.17 million	\$113,700
ROI (NPV/investment)	290%	290%
Payback (months)	4	4
Discount factor (%)	12	12

Source: IDC, 2020

CHALLENGES/OPPORTUNITIES

Although Pure Storage established an early technology lead in the AFA market, its competitors have learned from its success and now offer many of the same "flash optimized" array features as the FlashArray. Over the past several years, the AFA market dynamic has evolved to the point where many functional capabilities have become checklist items, which are offered by a number of the established enterprise storage vendors. Over the years, Pure Storage has taken note of these developments and has pioneered innovations in other relevant areas, like the efficiency (not just the existence) of their data reduction technology, rapid integration with new Oracle database tools and features, the value it drives with its AI/ML-driven cloud-based predictive analytics platform (Pure1), multigenerational nondisruptive technology upgrades, its ability to support multiple public cloud-based environments, and the overall quality of the customer experience it provides. These innovations continue to

differentiate Pure Storage from other enterprise storage vendors among both IT infrastructure managers and DBAs even as they lower the cost and administrative involvement over time to support high-performance, mission-critical Oracle database environments.

Pure Storage has been able to capitalize on its differentiators to bring new customers on at a higher rate (as a percentage of the installed base) than its competitors throughout most of its history. It is interesting to note that while 2019 was a difficult year for the enterprise storage market overall (which exhibited flat to slowly declining revenue over that period), Pure Storage turned in 20% year-over-year growth. The vendor has been successful in generating new points of differentiation even as its market has evolved — not only will it need to continue to do so but it will also need to effectively capitalize on the opportunities that differentiation presents. The high Net Promoter Score of the vendor is indicative of how happy its customers are today, but Pure Storage will need to stay focused to keep customer satisfaction high as its installed base grows into tens of thousands of systems. A high NPS score translates into strongly positive word-of-mouth support from existing customers, and the vendor has been effective at leveraging it as it brings new prospects on board. Strong business value metrics like those derived from our customer interviews for this white paper have also helped make the vendor a compelling choice in enterprise storage.

CONCLUSION

Enterprise storage customers have made the transition to AFAs very rapidly over the past six years, and these types of systems drive over 80% of all external primary storage revenue today. As this white paper shows, the total cost of ownership justifications for a move to all-flash storage are very strong. AFAs deliver much higher performance and throughput, they need far fewer devices to drive a given level of performance and capacity resulting in lower energy and floor space consumption, and their much lower latencies drive higher CPU utilization on the server side, reducing the number of CPU cores and associated software licensing costs. Flash performance results in noticeably lower storage provisioning times and enables the use of a wide variety of enterprise-class data services (inline data reduction, thin provisioning, RAID, snapshots, encryption, replication, etc.) while still delivering sub-millisecond latencies at scale. In Oracle environments, AFAs deliver much faster database provisioning and query times, minimize performance tuning efforts, and enable the high availability needed for mission-critical databases.

Pure Storage was one of the earliest entrants in the AFA market in 2012, and it is a leader in that space, not only in revenue growth but also in other areas that have become much more important to Oracle enterprise users in the past six years: NVMe technology, the use

of artificial intelligence and machine learning to deliver self-managing systems, the extent of integration with the Oracle database and other tools, and the nature and quality of the customer experience throughout the storage life cycle.

For this white paper, IDC interviewed seven users of Pure Storage FlashArrays in Oracle database environments. This white paper indicates that those customers that have moved to AFA vendor Pure Storage's AFAs for their Oracle database environments are achieving significant benefits in terms of higher performance, faster response, improved productivity, and lower costs. DBAs have enjoyed 71% faster new database deployments, 85% faster database query speed, and 43% higher DBA productivity, while storage administrators have seen 54% faster new storage deployments, 156% more usable flash capacity, and 57% higher productivity. All in all, interviewees for this white paper reduced their total cost of operations by 73% and achieved a 290% three-year return on investment. Over a three-year period, these Pure Storage/Oracle customers interviewed indicated total discounted three-year benefits of \$5.61 million based on increased business results, staff efficiencies, and lower overall costs.

APPENDIX

Methodology

IDC's standard ROI methodology was utilized for this project. This methodology is based on gathering data from current users of Pure Storage using Oracle environments, applications, and workloads as the foundation for the model. Based on interviews with these organizations, IDC performed a three-step process to calculate the ROI and payback period:

- **Gathered quantitative benefit information during the interviews using a before-and-after assessment of the impact of Pure Storage.** In this white paper, the benefits included IT cost reductions and avoidances, staff time savings and productivity benefits, and revenue gains.
- **Created a complete investment (three-year total cost analysis) profile based on the interviews.** Investments go beyond the initial and annual costs of using Pure Storage and can include additional costs related to migrations, planning, consulting, and staff or user training.
- **Calculated the ROI and payback period.** IDC conducted a depreciated cash flow analysis of the benefits and investments for the organizations' use of Pure Storage over a three-year period. ROI is the ratio of the net present value (NPV) and the discounted investment. The payback period is the point at which cumulative benefits equal the initial investment.

IDC bases the payback period and ROI calculations on a number of assumptions, which are summarized as follows:

- Time values are multiplied by burdened salary (salary + 28% for benefits and overhead) to quantify efficiency and productivity savings. For purposes of this analysis, IDC has used assumptions of an average fully loaded \$100,000 per year salary for IT staff members and an average fully loaded salary of \$70,000 for non-IT staff members. IDC assumes that employees work 1,880 hours per year (47 weeks x 40 hours).
- The net present value of the three-year savings is calculated by subtracting the amount that would have been realized by investing the original sum in an instrument yielding a 12% return to allow for the missed opportunity cost. This accounts for both the assumed cost of money and the assumed rate of return.
- Further, because Pure Storage requires a deployment period, the full benefits of the solution are not available during deployment. To capture this reality, IDC prorates the benefits on a monthly basis and then subtracts the deployment time from the first-year savings.

Note: All numbers in this document may not be exact due to rounding.

IDC Research, Inc.

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 2020 IDC.
Reproduction without written
permission is completely forbidden.*

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.