

White Paper

Pure Storage's Evergreen Storage Service Delivers Storage-as-a-Service Benefits for On-Premises Infrastructure

Sponsored by: Pure Storage

Eric Burgener
May 2018

IDC OPINION

As enterprises undergo digital transformation, they are looking to streamline the efficiency of their information technology (IT) infrastructure. Cloud-based services provide a valuable option for certain workloads, allowing CIOs to not only offload IT infrastructure management tasks to a third party but also enable different purchase options. Cloud-based storage services provide a pay-per-use option that can be very attractive for short-term projects or to handle higher seasonal demands, but they also give enterprises the ability to move to an operational expenditure (opex) model for affected infrastructure. Even by itself, the ability to move to an opex model for IT infrastructure is attractive enough that many enterprise storage vendors are offering storage utility pricing programs for on-premises infrastructure. Many customers have taken advantage of these programs, often as a way to finance on-premises private cloud infrastructure. However, these offerings don't match the service level of a true cloud service and lack flexibility. In addition, given the way that these storage utility pricing programs are generally structured, upcoming changes to International Accounting Standards Board (IASB) standards for handling leases will reclassify those assets as capital equipment and move them back onto corporate balance sheets. These changes are slated to start taking effect in December 2018.

Pure Storage, recognized by IDC as a leader in the all-flash array (AFA) market over the past three years, has just introduced a new product offering – the Evergreen Storage Service (ES2) – that not only provides a true tier 1 cloud storage service experience on-premises but also preserves use of the opex model despite the IASB 2019 changes. ES2 provides a opex¹, always available, pay-per-use service for block-based storage infrastructure that meets enterprise performance, availability, and data integrity requirements – notably offering a lower total cost of ownership (TCO) than "enterprise" storage from existing public cloud storage providers. Customers looking for this class of storage with a pay-per-use model that is better than those offered by cloud providers such as Amazon, Google, and Microsoft will want to take a look at this offering.

IN THIS WHITE PAPER

Storage utility pricing programs available from many enterprise storage vendors provide an opex option for on-premises infrastructure that many customers appreciate. Because of the way most of

¹ *Opex treatment is subject to customer's auditor review*

these programs are structured, they fail to provide all of the key benefits of tier 1 cloud storage services, and the upcoming changes to IASB standards will, most likely, force the underlying assets back onto the balance sheet as capital equipment. These accounting changes are slated to become effective starting in December 2018. This white paper summarizes the changes, discusses the impact of the changes on legacy storage utility pricing programs in general, and takes a close look at a new storage pay-per-use private cloud offering from Pure Storage that not only will continue to provide an opex² purchase option for enterprise-class, on-premises infrastructure in the wake of the IASB changes but also can truly deliver the tier 1 storage-as-a-service experience.

SITUATION OVERVIEW

The use of cloud storage services has grown significantly in the past decade. Today, aggregate cloud storage infrastructure spend by IT organizations has already reached 60% of total storage spend. Cloud services offer attractive benefits, including significantly increased IT agility, simplified infrastructure management, improved economics, and the ability to move from a capital expenditure (capex) model to an opex model. Many enterprises have a "cloud first" strategy for new application deployment, but there are still reasons why certain workloads are hosted in on-premises infrastructure. Performance, security, and regulatory compliance tend to be leading reasons, although for workloads with certain access profiles, there can also be economic reasons to use on-premises storage infrastructure.

IDC draws a distinction between public and private clouds in the cloud storage market. In the public cloud model, a service provider (i.e., Amazon Web Services, Google Cloud Platform, or Microsoft Azure) makes IT resources, such as applications, virtual machines, or storage, available to the general public over the internet, typically using pay-per-use pricing that is considered an opex model. In this model, an enterprise uses the cloud provider's infrastructure but does not own it, and multiple different enterprises could be sharing the same physical IT infrastructure at the service provider (multitenancy).

A private cloud, on the other hand, is a particular model of cloud computing that involves a distinct and secure cloud-based environment that is dedicated to a single client. While public cloud providers can offer private cloud options running on their own infrastructure, many enterprises prefer to build their own on-premises private cloud infrastructure to address performance, security, or regulatory compliance requirements while enjoying the benefits of the cloud computing model such as IT agility and easy, self-service access for end users. Many enterprise storage vendors have introduced a "storage utility" pricing model that gives customers an opex option for on-premises infrastructure (although they do not necessarily provide a

Capex Versus Opex

Traditional capex models for the purchase of storage infrastructure typically required IT organizations to purchase much larger systems than they initially needed so that they could accommodate anticipated growth over the life of the asset. Cloud services, on the other hand, used a storage utility or "pay per use" pricing model that allowed enterprises to reduce their up-front expense for new purchases and pay for additional capacity only when they needed it. This opex model is one of the key benefits of the use of cloud infrastructure because it lowers the up-front costs to access needed IT resources. It also makes costs more predictable over the life of the asset.

² See footnote 1

true "services"-based experience). Many on-premises private clouds take advantage of this pricing option, although it doesn't truly give them a tier 1 cloud service experience, but it is also used by enterprises that want just the opex option and care less about the other cloud capabilities.

A high percentage of "storage utility" options provided by enterprise storage vendors and designed to solve for only opex pricing for on-premises infrastructure are in fact finance leases wrapped around a series of separate products and services. Under most of these programs today, customers have the right to control the use of the equipment, which is defined as "identified assets" (specific pieces of equipment with discrete serial numbers), and the terms tend to be long (generally at least 36 months or more). Despite the fact that these approaches offer "pay per use" pricing, they can be disruptive and inflexible over time — they typically do not truly offer a "public cloud" experience.

Under the impending IASB 2019 changes, lessees will be required to recognize assets and liabilities for leases with terms longer than 12 months on their balance sheets. Any leases that depend on the use of "identified assets" will now be accounted for in substantially the same manner as traditional capital leases, which again means that the covered equipment must appear on the balance sheet. This means that any "storage utility" pricing models that run afoul of these two changes (term length, use of identified assets) will no longer offer opex pricing that keeps capital equipment in use by an enterprise off the balance sheet. Public companies will be impacted by this change starting December 15, 2018, while privately held companies will not be affected by it until a year later (December 15, 2019).

The Definition of a True "Service" Model

Pay-per-use pricing and access to the infrastructure without up-front capital investment are clearly part of the definition of a true cloud service model, but there are other key aspects. The provider (e.g., Amazon AWS) is also responsible for the day-to-day management of the equipment used to deliver the service. Providers maintain and service that infrastructure to ensure that it delivers to predefined service levels, and they expand that infrastructure over time as demand for the service scales. Customer expectations around service levels dictate that any outages must be rapidly repaired, and the provider may even go so far as to provide availability guarantees at a certain level (e.g., four-nines). As new technology becomes available, a cloud service may want to refresh its infrastructure to make it better performing and more efficient, scalable, or cost effective, but technology refresh must be handled in a manner that is nondisruptive to its customer base. Technology refresh may also be undertaken to provide new services to help differentiate the provider in a competitive market or lower its costs.

Public cloud storage providers strive to operate in this manner. They provide a baseline of service; set customer expectations about performance, availability, and data integrity; and completely relieve customers of any responsibility around managing, upgrading, or refreshing the actual physical IT infrastructure that provides the service. From their customers' points of view, they have no control over the use of the infrastructure, there are no identified assets with a dedicated serial number that underlie the service, and cloud providers may move their "workload" around to different parts of the infrastructure without their customers' knowledge. This helps cloud providers meet their defined service-level agreements (SLAs).

Although many operating leases met the opex requirement in the past, it is interesting to note that the storage utility pricing models from most enterprise storage vendors never met this latter requirement to nondisruptively perform technology refreshes over time to ensure that the equipment enjoys performance, scalability, and efficiency improvements. With most SAN and NAS equipment, there is some limited ability to upgrade firmware and perhaps support multiple disk device sizes, but

cross-generational technology upgrades require a disruptive, risky, and expensive forklift upgrade as well as time-consuming data migration. Major upgrades that transition the array from SCSI to NVMe or to controllers built around next-generation Intel CPU technology just aren't possible without a frame swap. While it is true that these storage utility pricing programs would allow an enterprise to avoid putting capital equipment on the balance sheet, they didn't provide a true cloud storage service model for on-premises infrastructure. And with the impending IASB 2019 changes, many of them will no longer offer even the opex benefit.

Where Cloud Storage Providers Fall Short

As enterprises began their journey to hybrid cloud, they typically chose less critical secondary storage workloads as the candidates for outsourcing to public cloud providers. Many mission-critical workloads had much more stringent latency, recovery, security, and/or regulatory requirements and as a result were retained in on-premises infrastructure. Many of these critical workloads ran on very high-performance block-based storage platforms that offered extremely high availability and data integrity as well as a bevy of storage management features (compression, deduplication, RAID, thin provisioning, space-efficient snapshots, encryption, quality of service, replication, stretch clusters, etc.) that just were not available from block-based storage offerings in the cloud.

The benefits of public cloud environments — IT agility, ease of use and management, improved economics, and the transition to an opex model — could be quite interesting for many of these block-based storage environments if cloud-based offerings could meet their functional and cost structure requirements. Through the integration of flash-based media, higher-performance networking, and additional services, cloud providers are improving their ability to meet more of these types of functional requirements. While these improvements have increased enterprise adoption of these cloud storage offerings, for most block-based mission-critical workloads, enterprises still exhibit a strong preference for on-premises infrastructure.

Although the pay-per-use pricing of cloud-based storage services is more granular than what is available from owned, on-premises IT infrastructure, it is not ideal. Generally, cloud storage customers buy a certain amount of "provisioned capacity" for which they pay a fixed amount regardless of their actual usage, and their storage usage cannot exceed that amount unless they change the block storage configuration through the control plane. While some cloud storage services offer the ability to increase or decrease the amount of provisioned capacity dynamically, others are more restrictive with the changes allowed. As a result, most enterprise cloud customers overbuy on storage capacity to give room for growth and are paying for storage capacity they're not actually using, although they are paying less than they generally would if they were directly using owned, on-premises infrastructure (all of which they would have had to purchase and install up front). In addition, block storage pricing models do not typically offer "usage volume-based" or "up-front payment-based" discounts, limiting the cost flexibility available when using such services at a large scale. If cloud storage providers could offer block storage services in a manner that more closely matched the actual usage of block storage capacity with billing, while providing an opportunity to occasionally burst beyond the provisioned capacity, this would be viewed by customers as highly desirable. Customers would also benefit from access to volume usage discounts with cloud storage providers.

Pure Storage's Evergreen Storage: An Industry Leader

In 2015, Pure Storage introduced the Evergreen Storage Program. Evergreen Storage is a comprehensive set of programs and guarantees that significantly raised the bar throughout the entire sales and postsales customer experience (CX) for enterprise storage customers. As long as the

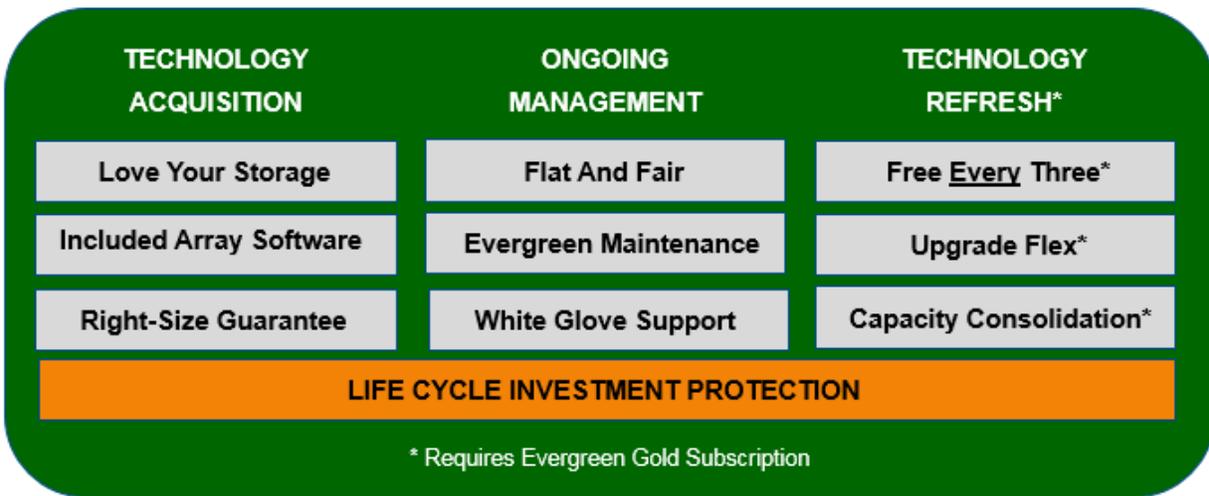
purchase is covered under the Evergreen Subscription, it covers everything from an unconditional money-back guarantee, bundled Pure software, and a storage efficiency guarantee to flat maintenance pricing, no-questions-asked replacement of any failed components for the life of the system, and controller upgrades every three years. Pure Storage's system architectures fully support nondisruptive technology refresh of any component in the system (including controllers and backplanes) and in-place data migration that establish a technology refresh model that is unmatched in the industry among products with similar enterprise-grade designs. This program, available on all Pure Storage enterprise storage systems, changed customer expectations around the value enterprise storage vendors should provide and has been a key contributor to Pure Storage's extremely high Net Promoter Score (NPS) rating of 83.7. The company's competitors have scrambled to respond, but Evergreen Storage continues to be the clear leader in the industry for programs of this type.

Evergreen Storage takes its name from the program's ability to support nondisruptive, multi-generational technology refresh on the customer's schedule, effectively enabling (but not requiring) the customer to upgrade the system each time a new controller, disk device, backplane, or host connection becomes available. Customers that want to upgrade controllers more frequently than every three years can use the Upgrade Flex program, included as part of Evergreen Storage, leveraging trade-in credits representing the full value of the controllers being replaced toward the purchase of new controllers. Pure Storage has delivered on this promise for over three years now, enabling customers to upgrade to new controllers with the latest Intel chipset as it is released and allowing backplane upgrades that moved systems from 6Gb SAS to 12Gb SAS to NVMe — all without any forklift upgrades or data migration and while preserving a customer's existing investment in hardware and storage capacity.

Customers have taken advantage of these upgrade opportunities in droves, enabling Pure Storage to have the highest percentage of its installed base on the latest firmware, hardware, and software of any vendor with storage systems in its class, a fact that improves system availability, functionality, and efficiency while lowering the cost of support. Figure 1 provides an overview of Evergreen Storage, but for more information on it, refer to *Pure Storage's Evergreen Storage Continues to Drive Differentiating Value for Enterprise External Array Customers* (IDC #US43750118, April 2018).

FIGURE 1

Evergreen Storage



Source: Pure Storage, 2018

Pure Storage's Evergreen Storage Service

With the introduction of the ES2, Pure Storage is delivering a storage pay-per-use service and pricing option for on-premises infrastructure that provides all the standard Evergreen Storage benefits in an opex³ model that is not impacted by IASB 2019. Because the storage service pricing model has often been used by enterprises to build on-premises private cloud infrastructure, this means that an opex⁴ model will still be available for Pure Storage customers that want to do that in a post-IASB 2019 world. As a potential customer evaluates what ES2 offers, it will be important to look at the features that match public cloud offerings and those that go beyond. The ES2 features that provide the benefits of public cloud include:

- **A true storage service that is "evergreen."** ES2 customers pay a set cost per GB for reserve capacity and for on-demand capacity (if used) during the term of the agreement. Reserve capacity is billed in equal 12 month installments over the course of the term, while costs associated with usage of on-demand capacity are billed monthly at a higher rate (and will vary based on usage). All pay-per-use billing is based on effective, not raw, capacity. Term lengths vary, just as they vary with public clouds, and can be as short as 12 months or as long as 36 months.

The customer is not responsible for any of the installation and maintenance activities normally associated with on-premises infrastructure. Pure Storage defines a "right sized" hardware configuration based on articulated customer requirements, deploys the hardware at the customer site, and upgrades the system over the period of the contract term as next-generation technologies become available without any disruption to application services.

Note that the "evergreen" capability includes the nondisruptive future migration to NVMe technology if and when it is needed. Pure Storage can deploy NVMe-based FlashArrays today

³ See footnote 1

⁴ See footnote 1

if the customer SLA requires it, but any SAS-based FlashArrays deployed under ES2 can be nondisruptively upgraded to an all NVMe-based array at no additional charge. Once NVMe over Fabric (NVMeF) becomes available later in 2018, customers who require that can also be nondisruptively upgraded to it. Software upgrades during the term of the contract will also bring major new functionality.

- **A true opex⁵ pricing model.** Because ES2 does not operate with "identified assets," and is delivered as a service, it is defined as a true opex model. There also is no buyout option, another factor that IASB calls out as something that forces the asset to appear on the balance sheet. So ES2 maintains its true opex⁶ status even as IASB 2019 requirements take effect starting in December 2018.

Both of these features differentiate ES2 significantly from most of the storage utility offerings provided by enterprise storage vendors today that will no longer meet opex requirements in the wake of IASB 2019. With respect to public cloud, though, ES2 goes beyond providing just an opex⁷ option for on-premises IT infrastructure — it actually significantly improves upon what leading cloud providers offer for block-based, mission-critical, primary storage applications. Features that differentiate ES2 from public cloud offerings include:

- **Built for tier 1 workloads.** ES2 customers specify a service level, the effective capacity, and the term, and Pure Storage will configure and deploy an AFA platform guaranteed to meet performance, capacity, and availability requirements. The SLA will determine specifically which FlashArray platform Pure Storage deploys to meet IOPS requirements, the effective capacity will determine how much raw capacity is deployed, and all storage services are proven to meet six-nines (99.9999%) availability. Because ES2 is based on Pure Storage's proven FlashArray platform, it offers true enterprise-class block storage. The FlashArray has delivered "six-nines" availability across Pure Storage's entire installed base since 2015 and includes a full set of enterprise-class data services, including inline compression and deduplication, dual parity RAID, thin provisioning, space-efficient snapshots and clones, encryption, quality of service, synchronous and asynchronous replication, support for VMware Virtual Volumes (VVs), and stretch clusters. This is reliable block storage that fully meets enterprise requirements for performance, availability, and data integrity.
- **Significantly lower pay-per-use pricing.** The ES2 starts at a fraction of the price of equivalent public cloud for "used" storage, with a 100TB minimum effective capacity for storage that provides consistent sub-millisecond response times under load. This is for effective capacity (which includes the combined impact of dual parity RAID protection, anticipated snapshot usage, and Pure Storage's inline storage efficiency technologies) and is priced on par with leading cloud providers' provisioned capacity (which is just raw storage capacity before data protection, snapshots, and other features that add to storage capacity consumption). Compared with the effective cost per GB of the most reliable provisioned storage capacity available from major cloud providers such as Amazon, Google, and Microsoft, ES2 is roughly half the cost and sometimes less.

ES2 customers set a "reserve" capacity for which they pay up front to cover the 12-month installment on each anniversary for the term of the agreement. Discounts are available for higher reserve capacity amounts. Pure Storage configures in an "on-demand" capacity buffer in the on-premises infrastructure that customers can burst into if needed, but at the cost of paying a slightly higher cost per GB on just that capacity "if it is used." Customers can adjust

⁵ See footnote 1

⁶ See footnote 1

⁷ See footnote 1

their reserve capacity amount upon a monthly basis, and Pure Storage will ensure that the on-demand capacity is adjusted as well so as to maintain a 25% buffer to accommodate burst capacity requirements. Downward adjustments are available at renewal only. The access to on-demand capacity that is not paid for until it is used provides an available buffer, but it also means that customers do not have to overprovision their "reserve" capacity in the same manner that public cloud customers overprovision (and pay for) their provisioned capacity. This goes above and beyond the straight lower price for reserve (versus provisioned public cloud capacity), reducing costs even more. What the customer pays on a monthly basis is the cost per GB of the blended usage during the preceding month (reserve plus on-demand capacity).

- **Changing data access patterns that do not impact cost.** While public cloud storage pricing can be very attractive relative to up-front capex purchases, the additional costs required to repeatedly access that data (not just store it) will vary based on frequency of access and can quickly get out of hand. Because of these "egress" charges, cloud customers carefully consider their data access patterns for particular workloads before they think about transitioning them to the cloud. ES2 on-premises customers pay the set prices for reserve and on-demand capacity, regardless of data access patterns; these set prices provide costs that are potentially both much lower than public cloud storage costs and much more predictable.
- **Nondisruptive transition between opex models and capex models.** ES2 supports another feature not available at all with public cloud offerings: the ability to nondisruptively migrate between an opex model and a capex model for the "storage service." Pure Storage does this using ActiveCluster, its stretched cluster solution. With this model, if an ES2 customer wants to transition back to the capex model, it purchases a net-new Pure Storage AFA, installs the AFA in its datacenter, establishes an ActiveCluster relationship with the ES2 FlashArray, migrates the data volumes, and then ships the ES2 FlashArray back to Pure Storage. All of this is done nondisruptively without impacting application services.

Hybrid clouds have become the mainstream deployment model. An IDC survey from mid-2017 found that 72.6% of datacenters are operating in a hybrid cloud environment. As a result, customers have become more focused on issues such as data mobility and the management semantics for data that may move between on-premises and off-premises infrastructure. The Purity1 storage operating environment that drives all FlashArrays supports features such as single-pane-of-glass management, multiple replication options, and REST. In addition, the environment supports a number of other APIs for broadly deployed third-party infrastructure, automation, true storage service pricing (with ES2) and out-of-the-box support for multiple public clouds that provide the capabilities enterprises need to deploy and manage an efficient, effective hybrid cloud environment.

CHALLENGES/OPPORTUNITIES

As with most programs, products, or services that break new ground, lack of market awareness can be a significant impediment to sales. This is the position Pure Storage's ES2 is in for the next year or two. Admittedly, leases are not the stuff of pulp fiction page-turners, but there is a clear preference among certain customer types to move to more of an opex model for IT infrastructure to achieve better IT agility, meet changing financial goals, and better ride technology curves. It is a smart move on the part of enterprise storage providers to offer storage utility pricing options to help them better compete with public cloud storage services. For customers that prefer the opex pricing model for on-premises infrastructure, IASB 2019 presents an obstacle. Most storage utility pricing models available from traditional enterprise storage vendors for their storage platforms will run afoul of IASB 2019 and lose their ability to keep on-premises IT infrastructure off the balance sheet.

While ES2 offers certain advantages relative to public cloud storage and other advantages relative to the storage utility pricing models for on-premises infrastructure of competitors, its true "evergreen" nature differentiates it from both. So far, only Pure Storage can offer a comprehensively nondisruptive technology refresh for storage solutions that meet true enterprise requirements for performance, availability, and data integrity.

There is no doubt that more enterprise customers are building their own private clouds in their own datacenters, and this trend is not likely to change going forward given the ability to do this while enjoying an opex⁸ pricing model. So far, no other enterprise storage vendors have made changes to their storage utility pricing programs to address the IASB 2019 issue, and public cloud storage providers are struggling to provide sufficient performance and availability along with reliable block-based storage for primary workloads. This provides a significant opportunity for Pure Storage to cater to customers that want to stay with and/or move to an opex pricing model for block-based workloads run in on-premises infrastructure in their own datacenters.

CONCLUSION

Customers using storage utility pricing programs that are really just finance leases will need to take action if they want to maintain a true opex model for on-premises infrastructure. If pay-per-use pricing is a component of datacenter financing strategy, then CIOs should familiarize themselves with the IASB 2019 changes and understand how the changes may impact them. This could also be an opportune time to consider a technology refresh to a different enterprise storage vendor for customers that are looking for a change.

IDC has identified Pure Storage as a consistent AFA market leader over the past three years. This standing has been based on the functional capabilities of the company's FlashArrays and ability to drive industry change around the use of storage efficiency technologies and cloud-based predictive analytics as well as the impact of the Evergreen Storage program in setting a new, higher bar for how enterprise storage customers expect to be treated by their vendors. Customers have responded very positively to this different treatment, rewarding Pure Storage with the highest NPS for a vendor of enterprise-class array solutions. With ES2, Pure Storage is now providing a true storage service pricing model for enterprise-class, block-based storage that customers want as part of on-premises (not public cloud) infrastructure. While other vendors offer storage utility pricing programs for on-premises equipment, ES2 is the first program that will preserve the opex⁹ model in the wake of the upcoming IASB 2019 changes. As enterprises ponder how to maintain the opex model for on-premises infrastructure for the future, they may want to look at what Pure Storage offers.

⁸ See footnote 1

⁹ See footnote 1

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-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 2018 IDC. Reproduction without written permission is completely forbidden.

