

White Paper

Containerizing Key Business Workloads: Meeting End-to-End Kubernetes Data Services Needs for Enterprise Applications

Sponsored by: Pure Storage

Archana Venkatraman

May 2021

IDC OPINION

Picture this: By 2023, more than 500 million new logical applications will be created – equivalent to the number of applications created in the past 40 years.

This growth in applications predicted by IDC is a result of organizations' efforts to become "digital innovation factories" over the next five years. As such, organizations will develop digital products and services with digital-native speed and scale. The COVID-19 pandemic has further accelerated this digital journey as businesses prepare to survive in the post-pandemic new world order.

These future apps and services will be at the core of every industry's digital value proposition. A large number of these applications are developed and deployed in containers, use Kubernetes orchestration environments, and consume external API-fronted services. Container-enabled infrastructure is drawing tremendous interest from customers, and IDC expects the installed base of container instances to grow at a 62.1% CAGR from 2019 to 2023.

Why? Because containers allow for agile and automated deployment of modern applications at scale and at costs that are appealing. Their lightweight nature and provisioning efficiency mean that a single server can generally host many more containers than virtual machines (VMs) and achieve higher utilization.

Seeing the speed, efficiency, and practicality of containers in running cloud-native applications, enterprises are considering containers to address the larger need of transforming existing workloads. IDC predicts a wholesale shift to containers in three years where 80% of workloads will shift to or be created with containers and microservices, reducing per-app infrastructure needs by 60% while driving a 70% improvement in digital service resiliency. Already, 84% of respondents are using containers in production, according to a 2019 Cloud Native Computing Foundation (CNCF) survey.

Core application modernization via containers can deliver several benefits such as:

- Breathe new digital life into legacy applications
- Support cloud migration strategies
- Make more efficient use of computing resources
- Make data and applications portable to support multicloud and hybrid cloud environments

- Pivot to fully automated deployment and operations to eliminate slowdown and errors with manual processes

Pure Storage's acquisition of Portworx, the Kubernetes data services platform, helps the company bring complete data services through a single data control plane. Portworx by Pure Storage provides persistent storage, high availability, data protection, data security, and cloud mobility for containers deployed in any cloud (hybrid or multicloud) architectures. The platform helps Pure Storage further decouple its data services from hardware and take a software- and API-driven approach to deliver Kubernetes-native data services.

In IDC's opinion, this software-centric capability puts Pure Storage in the sweet spot for enterprises building software, modernizing apps, and accelerating their cloud-native strategies. IDC estimates that by 2025, up to a quarter of Fortune 500 companies will become software producers to digitally transform and maintain their F-500 status. The platform already supports dozens of Global 2000 companies run containerized data services in production.

Portworx's richer portfolio of hybrid cloud and cloud-native data services is ever more relevant in the new normal. In IDC's 2020 *COVID-19 Impact Survey*, organizations cited an aggressive push to the cloud as the top priority in their long-term IT strategy in the post-pandemic world. We predict that by the end of 2021, 80% of enterprises will put processes in place to shift to cloud-based digital infrastructure twice as fast as before the pandemic.

With the breadth and depth of data services and support for use cases including database as a service, container as a service (CaaS), SaaS, disaster recovery (DR), and backup for Kubernetes, Pure Storage is establishing itself as a single go-to source for data services needs for essential workloads such as analytics/AI, databases, and CI/CD pipelines and build systems, as well as hybrid cloud strategies.

This desire to containerize mainstream applications is overturning storage and data management paradigms as IT customers seek container-defined storage and data services to power their modern app environments. Keeping storage overheads under control as containers scale is critical as more persistent workloads enter the container world.

Enterprises cite storage and data persistence as one of the key barriers to broad adoption of containers besides security, performance overheads, and integration challenges. Storage solutions that deliver container-aware storage with simplicity, cloud integration, flexible costs, automation, elastic scaling, and smart management can truly help organizations modernize their IT core.

IN THIS WHITE PAPER

Containers were initially aimed at new cloud-native applications because containers are inherently stateless, but enterprises are now evaluating them for "lift and shift" applications such as classic three-tier web applications, Java applications, or any database-dependent workload. This widens the workloads that containers need to address, and it has become vital to respond to the storage, data integrity, data persistence, storage persistence, container patching, and security needs to accelerate container adoption across enterprises.

This white paper analyzes the multiple approaches to support data needs of containerized applications. It also assesses Portworx by Pure Storage and how it is differentiated in the market. In

addition, IDC examines the future opportunities and challenges within the context of the evolving container data services landscape.

SITUATION OVERVIEW

Containers are already in heavy production use in many hyperscale web and SaaS providers such as Google, which runs its entire infrastructure in containers. But container instances in non-hyperscale environments will increase at a CAGR of over 100% from 2019 to 2023, higher than the overall CAGR of 62.1%, as enterprises emulate hyperscaler-like architecture and software development methodology to compete in the digital economy.

Containers are appealing because they help enterprises:

- Develop software faster to improve time to market and compete with digital disruptors
- Scale applications to be highly responsive and to meet increasing application loads
- Modernize application design by using microservices and DevOps methodologies to suit the cloud world
- Address the constant need to make more efficient use of computing resources and overcome traditional virtualization licensing fees as application onslaught occurs
- Build in the cloud and deploy on premises by making applications portable as multicloud and hybrid cloud become the norm

Some organizations also feel that containers make it much easier to collaborate on development across different geographies. When IDC asked enterprises about IT strategies and investments this year, the top 3 initiatives cited were cloud infrastructure evolution (50%), IT infrastructure modernization (50%), and security, compliance, and privacy (49%). We see container adoption at the heart of this infrastructure modernization and cloud infrastructure-as-a-service (IaaS) evolution. Traditional workload modernization is necessary for digital innovation, and cloud and container platforms are enablers of this modernization.

According to IDC's 2020 *Multicloud Survey*, security, multicloud workload portability, automation, infrastructure modernization, and use of containers are among the key priorities to accelerate application delivery (see Figure 1).

FIGURE 1

Key IT Priorities for 2020-2021



Source: IDC, 2021

Kubernetes Becoming Mainstream

Most of the initial workloads deployed are "stateless," with less than 20% of applications today having a need for state. With developers now testing containers for lift and shift classic three-tier web applications and Java applications, a sizable chunk of container applications will have state.

IDC predicts a wholesale shift to containers in three years where 80% of workloads will shift to or be created with containers and microservices, reducing per-app infrastructure needs by 60% while driving a 70% improvement in digital service resiliency.

In addition, there is growing use of open source NoSQL databases such as MongoDB, MariaDB, and Cassandra so that developers can control data structures in their applications without reliance on DBAs or a formal database schema change process. These database nodes that integrate with Docker and Kubernetes are stateful, and IT architects need to ensure full resiliency and protection.

In IDC's opinion, as container technology matures further, there will be a balanced ratio of cloud-native and traditional workloads on containers and the arbitrary differentiation between stateless versus stateful will fade and the need for security, data services, and storage for containers will be critical.

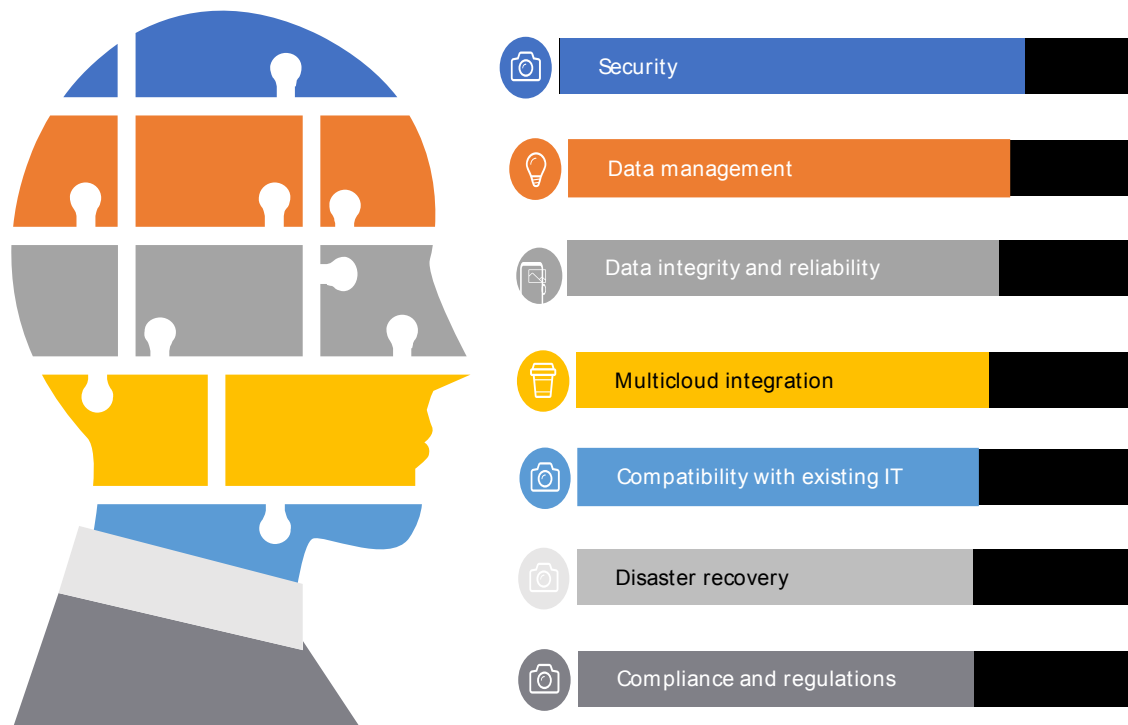
As more traditional workloads get containerized, complete data services such as enterprise-grade persistent storage, automated provisioning, data protection, DR, and security services become must-haves.

To scale their Kubernetes environments and accelerate cloud-native journeys, organizations need to assess and overcome container-specific challenges such as security, data management, and data mobility (see Figure 2).

FIGURE 2

Container Deployment Challenges

Top challenges when deploying containers



Source: IDC, 2021

Enterprises need the following capabilities in their container data services platform:

- **Hybrid cloud support** – Ability to adopt a "build once and deploy anywhere" strategy for application development
- **Day 2 operations** – Ability to address both "Dev" (speed and self-service) needs and "Ops" needs (security, reliability, DR, data protection, and high availability)

- **Kubernetes-native experience** – Integration with standardized container environment and automated data services on demand
- **Automation** – Automated container-aware data services for a variety of use cases such as persistence, data protection, scaling applications, provisioning multiple database services, and delivering a unified automated container data control plane that ingests, integrates, and correlates data for Kubernetes environments
- **Flexibility** – Support for heterogeneous container platforms, databases, and storage protocols
- **Single data control plane** – Support for CaaS, DBaaS, SaaS, and DR needs

We assess the six key data services needs and use cases in Kubernetes environments in the sections that follow.

Enterprise-Ready Container as a Service

The compelling benefits of containers are driving IT teams to architect a CaaS platform to support developers' infrastructure resource needs quickly. Moving beyond logical and composite applications to running enterprise-grade stateful workloads on CaaS platforms requires container-native storage and data management features to deliver self-service and everything as a code to developers without compromising on data security needs.

Enterprise-Grade Data Protection for Kubernetes Environments

Although containers represent a paradigm shift in speed and agility, basic business continuity, security, and data protection principles don't go away. Organizations require solutions to back up and protect containerized environments, including the data, metadata, application configuration, and Kubernetes objects – with a single click at the Kubernetes pod, namespace, or cluster level.

However, traditional backup environments aren't built for Kubernetes. Data management, security, and compliance and regulations are among the key concerns in using containers in production environments (refer back to Figure 2).

Companies building digital and data resiliency are racing to deliver zero data loss and zero downtime. Access to data protection services such as granular restores, flexible backup targets, application-aware backup, hybrid cloud data mobility and restores, and policy-based automation for data protection is critical to delivering resiliency.

Kubernetes Disaster Recovery

The pandemic has shone a spotlight on the value of DR and business continuity. Enterprises are looking to extend these capabilities to their modern cloud-native and container worlds. But without automation, DR for Kubernetes can be a complex process. Robust RPOs and RTOs for hybrid cloud and container applications are emerging as must-haves. Protecting data, Kubernetes objects, and configuration data can help recover applications faster and easier. Kubernetes environments are also home to heterogeneous databases including MongoDB, Kafka, PostgreSQL, Cassandra, and MySQL, requiring IT teams to develop skills in managing the recovery of each data service.

Database Services

Developers want speed and ease of use, with self-service, autoscaling, and automated deployments of databases. However, IT and governance teams have set out policies dictating strict rules around data protection, access control, encryption, and key management.

Provisioning application-driven database as a service with automated deployments, snapshots, and other data services can meet the needs of developers and operations teams. The IT teams need the flexibility to create secure guardrails centrally, as well as create resource limits when giving dev teams self-service database capabilities.

As one financial services customer told IDC, "Databases are our biggest bottleneck in accelerated software delivery strategies. It takes 3.5 hours a day for a database to deploy and 100% of our applications talk to databases."

A container data management platform that offers any database as service to users without requiring them to be expert in each data service, and provides automated deployments, autoscaling, and robust security, can accelerate time to market.

Scaling CI/CD

Large organizations, especially service providers and cloud providers, are building CI/CD pipelines on containers to improve their speed to market and agile software delivery objectives. The CI/CD systems (such as Jenkins or Bamboo) include all the images and artifacts developers need to run their tests. Developers need to recreate production environments rapidly for tests and roll back after each test to avoid failures.

However, many users are tempted to use the environment they have built, pulling all the relevant images and artifacts for a test run for multiple tests because of the time it takes to set the system up. But the slight configuration changes for each test can result in errors, delays, and inconsistencies, slowing down development cycles. The ability to snapshot the database configuration and use it as a "golden image" for multiple test runs is a key data service essential for container environments.

Scale and Efficiency in SaaS Environments

Many enterprises see the value in using Kubernetes as an orchestration platform to scale their SaaS applications. But this requires container-native storage and data management features, including end-to-end performance, reliability, and security, especially if the SaaS applications host critical data for analytics. Organizations that choose only a persistent storage layer without considerations such as automation or broader data services such as data protection, security, and data mobility will find it difficult to scale SaaS environments. Automation, particularly, is key as developers don't have to build deep expertise on every database and instead can ingest data at scale. This can ensure reliability of the application and also free up valuable time for DevOps and SREs to focus on more strategic tasks.

BENEFITS OF CONTAINER-DEFINED DATA MANAGEMENT

One important benefit of using container-defined data services is it meets the needs of all the use cases highlighted previously through a single, unified platform. It supports the data management needs of multiple workloads and facilitates cross-cloud data mobility. Enterprises can overcome the scalability and performance obstacles for their lift and shift applications in containers. This in turn brings portability and interoperability to go multicloud and accelerate cloud adoption in a risk-mitigated manner. As enterprises adopt Kubernetes environments for mainstream needs, they are evaluating moving beyond DAS and SAN/NAS storage plug-ins for containers to more software-defined, container-focused data services approach to meet their evolving needs. End-to-end container data management platforms bring highly scalable and configurable Kubernetes solutions with enterprise-grade features such as resiliency, availability, DR, and quality of service (QoS).

Containerized Applications: Designing Storage to Meet Modern Workloads and Users' Needs

Containers are needed for cloud-native applications (optimized for cloud and use automation, microservices, and DevOps processes) as well as for existing workloads such as mobile platforms that come with databases (such as MongoDB). Running existing workloads in containers is a game changer for IT. These workloads need storage persistence and enterprise-grade backup and security. As these persistent workloads move from being an exception to the norm, traditional storage infrastructure is not able to cope with the new container paradigms. When it comes to the top challenges around enterprise data management, 34% of organizations cite platform-as-a-service (PaaS) data protection and 24% cite containerized applications.

IDC believes that in a modern architecture, storage and data management need to be defined by the application and not by the infrastructure. Infrastructure and DevOps teams prefer solutions that deliver a full suite of data services to the application. The future belongs to data services platforms that are hybrid cloud enablers, API driven, automated, container native, and highly scalable and support multiple databases and infrastructures. This aligns with modern application development approaches and new personas such as SREs.

Container data services will be a preferred data control plane in the long term; 64% of organizations IDC surveyed in 2020 in its multicloud survey said they are investing in software-defined and container-aware storage to meet modern Kubernetes data needs (see Figure 3).

Each container-native environment shouldn't be a data management silo. Enterprises prefer data management platforms that integrate with existing data architectures to create a continuum of data services across their virtual, cloud, and container environments. Enterprises need deep integration with container orchestration frameworks, such as Kubernetes, to deliver persistence and data resilience in a self-service manner across hybrid cloud. They also want the platform to help them leverage data that may exist outside of containers, which can be useful for data pipelines as well as for exploiting the powerful capability of building apps once and deploying anywhere.

Ultimately, the performance of applications is dependent on the speed and quality of their interactions with data, and more often, this is where the bottleneck lies. Data needs to be closer to applications, but traditional storage cannot autoconfigure, autoscale, and follow applications closely, making performance unpredictable in the new scale-up and scale-out containerized IT environments.

IDC believes that whether a workload is in VMs or containers, there can be no compromise on performance, resiliency, security, and data services. In fact, containers are more dynamic, so there is a greater need for faster and automated provisioning of storage. Container workloads have greater management and scalability demands that require fast, automated provisioning. In addition, as containerized environments sprawl quickly, it is even more important for storage to be cloudlike.

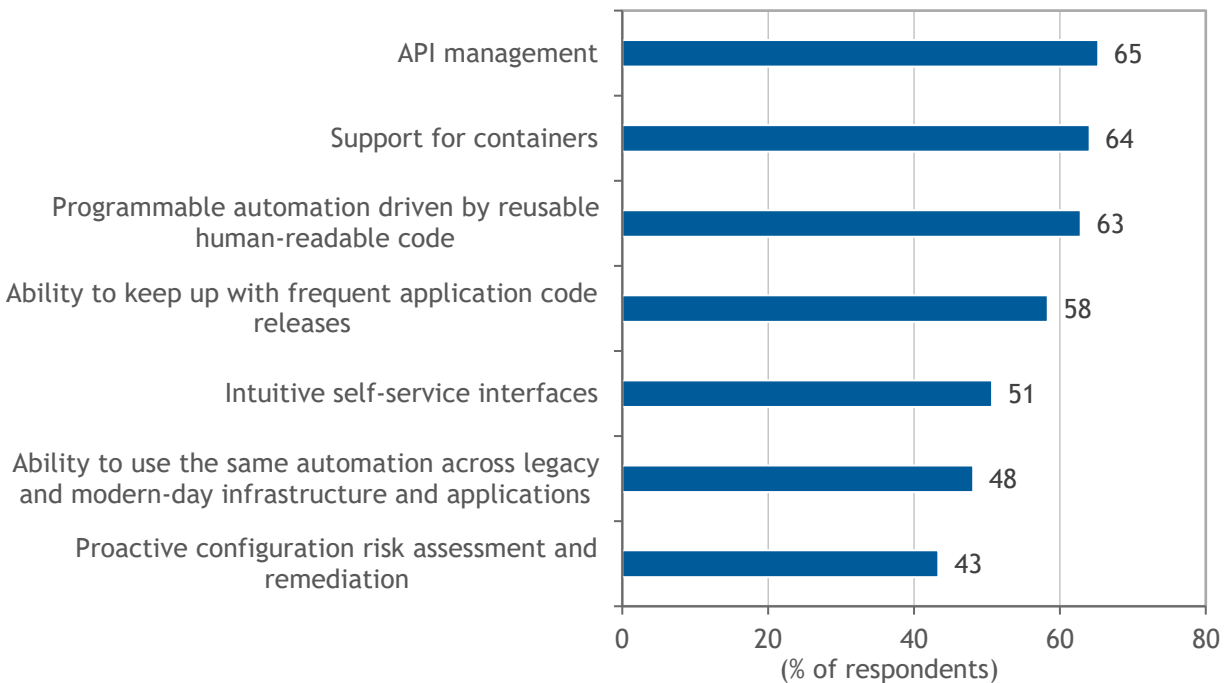
The emphasis will be on performance, data services, low costs, simplicity, and cloudlike experience in delivering storage to container workloads. Enterprises need to think about storage for modern infrastructure and for all key applications in a holistic manner. They need to consider whether the storage offering can bring the same tier 1 resilience, reliability, and protection that mission-critical applications depend upon when they are in container clusters.

Platforms that bring policy-driven storage with enterprise-grade data services will rise in popularity. The integration with the ecosystem – Docker, Kubernetes, Red Hat OpenShift, Google Anthos, AKS,

EKS, and VMware Enterprise PKS – can give developers easy access to storage volumes through these container platforms and allow them to manage storage via APIs, command line interfaces (CLIs), or graphical user interfaces (GUIs).

FIGURE 3

Key Infrastructure Priorities for Accelerated Application Delivery



Source: IDC, 2021

In many instances, developers are taking the lead in container adoption for cloud-native applications, but the infrastructure teams are evaluating containers too. In conversations with IDC, many telcos, OpenStack IT shops, and complex enterprise IT shops indicated that they are evaluating Kubernetes and looking at storage systems for scalability and high performance, making integration and functionalities such as scale out, application performance, and latency very critical.

Pure Storage ticks these boxes with its Kubernetes data management platform Portworx.

CONSIDERING PORTWORX – THE KUBERNETES DATA SERVICES PLATFORM

Pure Storage acquired Portworx, the Kubernetes data services platform to run business workloads in containers, in September 2020. This is significant in catapulting Pure's strategic role in delivering data services for modern hybrid cloud enterprises – giving customers the flexibility to run on any Kubernetes distribution, any cloud, or any on-prem infrastructure. The acquisition helps Pure Storage expand into the fast-growing market for multicloud data services to support Kubernetes and containers.

The Portworx Enterprise Storage Platform is an end-to-end storage and data management solution for multiple Kubernetes use cases, including persistent storage, data protection, DR, and security for Kubernetes, bursting to the cloud, cross-cloud data mobility, multidatabase services, consistency in CI/CD pipelines in containers, SaaS scaling, copy data management, and big data needs.

Designed for cloud-native environments, Portworx delivers enterprise-grade performance, reliability, and data management for the Kubernetes world, making these environments robust and resilient for critical workloads.

In IDC's opinion, what differentiates Portworx is the deep integration with the ecosystem, ranging from infrastructure to databases to container orchestrators, delivering a rich, agnostic, and hybrid cloud experience. IDC also believes that Portworx's features support data needs in the complete cloud-native journey from Kubernetes piloting to test and dev to production to multiregion environments and global container as a service. As a result, Portworx meets the needs of multiple personas including infrastructure teams, cloud architects, DevOps, application owners, SREs, Kubernetes architects, and storage professionals, delivering a unified experience.

Portworx also supports critical workloads in container environments with container-granular storage, metadata management, disaster recovery, data security, and multicloud migrations. It also brings modern data services experience for SREs, DevOps teams, DBAs, and application architects such as APIs, unification, golden images, scale, automation, and self-service.

Portworx includes the following features that help enterprises overcome key bottlenecks in their stateful container environments and database containers in production:

- Infrastructure as code
- Cross-cloud data mobility
- High availability, smart provisioning, and resiliency
- Scheduler-based automation and ability to meet data needs from familiar schedulers

These features support the following capabilities:

- Backup, DR, and security for stateful applications through encryption, access controls, and policy-based data management
- Run any data-intensive container (including MongoDB, Cassandra, PostgreSQL, or CI/CD pipes or process streams such as Kafka and Elasticsearch) on any hybrid cloud infrastructure (bare metal, VMware, Google, AWS, IBM, and Azure)

In IDC's 2019 *DevOps Survey*, 72% of organizations cited lack of integration into existing environments as the biggest technology roadblock in their modern application delivery pipeline. Portworx offers an opportunity to eliminate this bottleneck with unified, integrated, and automated data services for container environments used for app delivery.

Alongside container-defined data services, Portworx features such as storage tiering, intelligent placement of workload, smart provisioning, automated operations, and application-level replication bring business value such as cost savings, speed for DevOps, and freeing up professionals to focus on more strategic tasks.

"With Portworx, I can reduce Kafka brokers from five to three, a 40% savings, for the same level of reliability," said Jeffrey Zampieron, CTO, Beco.

What Portworx Means to Pure Storage and Its Customers

Pure Storage, with Portworx's container-native data services platform, can now position itself as a single go-to source for data services needs for all workloads – traditional, transforming, and cloud native.

The acquisition of Portworx comes at a time when organizations are seeking simple, unified data services such as consistent storage resources, data protection and security, data integration, and availability – for all applications hosted across their hybrid multicloud environments.

In six years, start-up Portworx has become a Kubernetes data services platform of choice for many global enterprises including GE, Comcast, Audi, Adobe, Roblox, DreamWorks Animation, T-Mobile, Lufthansa, Carrefour, ESRI, and Dailymotion, among others.

Pure Storage, which already has established customers, now has a wider set of personas to inform its engineering road map. The acquisition brings value far beyond just the Portworx platform IP. It gives Pure Storage access to a new breed of skills and talent, insights into cloud-native journeys of large complex organizations, and a production-ready software data services strategy.

Pure Storage started delivering data services and storage support for container environments back in 2017 with the launch of a container plug-in. It then evolved it into a full-blown container storage-as-a-service offering called Pure Service Orchestrator (PSO), which brings not just persistent storage for container workloads but also hybrid cloud connection, cloudlike pricing and consumption, deeper integrations, scale, automation, QoS, data services, and full compliance with CSI standards.

Portworx adds a full suite of data services for Kubernetes environments including persistent storage, high availability, data protection, data security, automation, and cloud mobility for containers deployed in hybrid cloud architectures. The platform helps Pure Storage further decouple its data services from hardware and take a software- and API-driven approach to delivering data services. This software-centric capability puts Pure Storage in the sweet spot for enterprises building software, modernizing apps, and accelerating their cloud-native strategies.

Developers and workload architects want a public cloudlike experience when provisioning storage for containerized applications in their own private clouds. In fact, IDC predicts that by 2022, over 80% of enterprises will prioritize the "public cloud experience" – such as access to new technologies and intuitive operations workflows – across on-prem, hosted, and public cloud environments. Companies need to evaluate important new technologies that provide this experience and leverage innovation in adjacent infrastructure areas to implement them.

The addition of a container-aware data platform further elevates the value of Pure Data Services, bringing storage and data services to all workloads, including virtualized, containerized, and cloud applications as well as monolithic traditional applications. The market for cloud application platforms is without a doubt red hot. Pure Storage is placing the right bets by acquiring Portworx to deliver an agnostic service for any PaaS and has added integrations to several PaaS platforms, including Azure Kubernetes Service, Red Hat OpenShift, Google Kubernetes Engine (GKE), Amazon EKS, Rancher Kubernetes Engine, and IBM Kubernetes Services. This adds to Pure Storage's shared storage service support for AWS, Microsoft Azure, and VMware Tanzu.

Enterprises continue to leverage PaaS and application platforms to help them build the next generation of highly scalable applications using newer tools such as microservices. The PaaS market

is estimated to grow at a strong double-digit CAGR. IDC expects the public cloud platform-as-a-service market to have a 26.8% CAGR from 2018 to 2021 and the hosted private PaaS market to grow at a 21.3% CAGR in the same period. Pure Storage needs to continue its focus on broad container and both public and private PaaS markets to remain relevant and widen its total addressable market (TAM).

As a storage industry disruptor, Pure Storage understands enterprise storage needs and infrastructure complexities, as well as the importance of offering features such as cloudlike scale and automated provisioning, to appeal to a new breed of personas such as developers, cloud architects, application owners, and DevOps users.

CHALLENGES/OPPORTUNITIES

Containers are proving to be highly useful for business agility, software development life-cycle (SDLC) acceleration, and cost savings in the era of the innovation economy. There is a fast-growing ecosystem of container schedulers and frameworks, and much of the innovation is focused on persistent storage challenges for stateful applications, as persistent storage is identified as one of the first significant barriers to faster adoption. In the past four years, Pure Storage has maintained the momentum in maturing its container storage offering to claim a strong stake in this burgeoning market.

Challenges

Lack of talent and skills is a major obstacle for many enterprises to use containers to production. Beyond extremely large enterprises, container use is still restricted to cloud-native and test and dev applications. Pure Storage will need to raise awareness and create innovation centers for tier 2 enterprises to experiment and understand how the company can overcome its storage obstacles to broaden its container adoption.

Limitations in the channel community will guide enterprises in their container adoption journeys. Pure Storage will also need to invest in training and upskilling its channel to demonstrate the benefits of Portworx data services to unleash more applications as container-ready candidates. The vendor needs to be mindful of competition from all sides. The container storage market is heating up, with traditional vendors as well as container storage start-ups battling for customer mindshare.

Pure Storage will also need to be mindful of the competition emerging in the container data services area, but its early entry to the market can help it be one of the leaders in the space.

Opportunities

The pandemic has accelerated digital transformation, and consequently, cloud and container adoption features aggressively in three quarters of enterprises' long-term IT strategies, according to IDC's research. Pure Storage appears to be ready to meet the data needs of mainstream container workloads.

IDC estimates that by 2024, net-new production-grade cloud-native apps will increase to 70% from 10% of all apps in 2020 because of adoption of technologies such as microservices, containers, dynamic orchestration, and DevOps.

Containers will coexist with virtual machines in the enterprise for the foreseeable future, and enterprises are likely to prefer storage solutions that help them build a continuum using familiar storage provisioning and management interfaces across multi-infrastructures.

IDC notes that many of Pure Storage's customers are running VMware environments and will welcome support for their container and cloud-native journeys through a container-native data services platform from their trusted vendor. IDC also sees AKS, EKS, GKE, Red Hat's OpenShift, and IBM Kubernetes Engine rising in popularity across all verticals, with CIOs seeing internal requests from developer teams for specific platforms. Investing in integration with these platforms are solid first steps for Pure Storage. But IDC research shows that on average, about 60% of the application estate is Windows based in a typical multinational enterprise, making only 40% of applications relevant for OpenShift, and the migration journey will be a multiyear one for enterprises.

Linux is increasingly becoming a preferred platform for modern infrastructure, including hosting traditional as well as next-generation applications. For instance, the Windows operating systems and subsystems (OSS) market fell by about a 5% CAGR from 2012 to 2017 compared with the Linux OSS market growing at a 14% CAGR during the same period. IDC expects more enterprises to choose Linux as they transform their IT environments, but the Windows installed base is well entrenched and will account for the bulk of the operating systems market for the foreseeable future.

Moving forward, Pure Storage will need to:

- Increase the use of Portworx platform in its core customer base.
- Continue enriching its container ecosystem through alliances with more PaaS, cloud, and data services vendors and broaden TAM beyond loyal customers.
- Continue expanding data services for edge container use cases beyond existing AWS Outposts and Equinix Platform integrations and adding object storage services to offer complete storage services per user demands.
- Focus on cloud providers and service provider customers that are building CI/CD pipelines on Kubernetes and communicate how they can find efficiencies, reliability, and cloud rebates.
- Expand data services messaging to benefits around automated container strategy, cloud optimization, autoscaling, and cloud migration.

From an engineering perspective, Pure Storage should now strengthen governance-related features around compliance, regulatory obligations, granular searches, rolling backups, and role-based access controls and security.

Containers are a new technology, and enterprises need help in accessing and migrating their existing application without compromising on cost, performance, and security. Vendors that can help customers navigate through this containerization and application estate modernization can remain a trusted partner for their long transformation journey.

ADVICE FOR TECHNOLOGY BUYERS

In the digital economy, an organization's advantage is not its size or strength, but its ability to adapt, be resilient, and quickly respond to crisis. Enterprises are required to deliver IT services at the speed of their business and, better still, the speed of their customers. Having a world-class technical foundation built on a hybrid cloud and container foundation is critical for competitive edge.

Enterprises need to:

- Make cloud migration and container adoption a key priority to accelerate digital transformation.
- Adopt containers for mainstream applications to modernize the core. Of course, container technology is not appropriate for every application, so take stock of the application estate and identify the traditional applications that can be optimized and containerized.
- Consider and address storage, data service, and security needs of containerized applications right at the onset and not as an afterthought when containerizing applications. To remain competitive and scale containerization, enterprise IT teams need to have container management teams and not storage management teams because infrastructure needs to be automated, API driven, and integrated with a container management layer.
- Assess whether the offering is container aware, scalable, automated, smart, open, and in sync with container standards when evaluating storage solutions.
- Keep up with innovation in container technology by training or hiring talent and skills and continuously improve on the processes and technologies incorporated in a containerized infrastructure.
- Choose end-to-end container-native data services platform to meet all data needs of containerized workloads.
- Choose data platforms that deliver cloudlike experience in technology, consumption, and innovation.
- Include elements that lead to the automation of the repetitive and onerous discovery, verification, and reconfiguration tasks associated with workload migrations, as with any technology, services, or organizational restructurings being considered today.
- Ensure that workloads and their accompanying data can quickly move across shared and dedicated cloud environments as well as critical edge locations without jeopardizing security or resiliency.

CONCLUSION

Interest in containers is clearly shifting to the next level, and IDC believes that innovation in adjacent areas such as storage and security will be key in converting this interest to reality. IDC also believes that the distinction between "stateless" and "stateful" will blur and the need for persistent storage, performance, and persistent data will be fundamental for mass containerization. Container-based microservices can be built, tested, and deployed more quickly, with improved agility and application resilience being the main business benefits. Portworx can help build microservice-based applications, container as a service, database as a service, SaaS, DR for Kubernetes, and Kubernetes backup. It has all the hallmarks of a modern container data management architecture. IDC believes that as a storage industry disruptor, Pure Storage is committed to addressing the fast-evolving Kubernetes data management needs of containerized applications.

As enterprises broaden their container adoption strategy, they may want to look at what Pure Storage offers.

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

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.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 2021 IDC. Reproduction without written permission is completely forbidden.

