PURESTORAGE®

WHITE PAPER

Enterprise Data Cloud Architecture

Reclaiming Control in a World of Data Chaos

Contents

Today's enterprise data environments are increasingly fragmented, operationally complex, and difficult to scale. The Enterprise Data Cloud (EDC) introduces a new model. One that virtualizes all storage into a unified data layer, governed through automation, delivered as a service, and powered by intelligence. With built-in resilience, integrated services, and SLA-backed flexibility, EDC enables IT to shift from managing infrastructure to delivering business outcomes. This paper explores the architecture behind EDC, how it's built with the Pure Storage® Platform, and why it fundamentally changes the way IT operates.

Introduction

Enterprise data environments are at a breaking point. Across industries, organizations are grappling with a perfect storm of operational complexity, fragmented infrastructure, and rising security risks.

Traditional storage architectures have fractured. Block, file, and object systems are now scattered across on-premises, hosted, and public cloud environments—each with its own tools, policies, and lifecycle requirements. This sprawl has created disconnected silos that are difficult to scale, impossible to standardize, and slow to respond to business needs.

At the same time, information security is under pressure. With data spread across so many platforms, visibility is limited, and enforcement is inconsistent. Policies slip through the cracks. Gaps appear between systems. Attackers exploit the chaos. And as ransomware threats rise and compliance expectations tighten, the margin for error is shrinking fast.

Patching these problems with more tools or manual processes doesn't work. The traditional model of isolated teams, overlapping systems, and reactive governance simply can't keep up.

What's needed is a new model. One where storage arrays no longer operate as isolated systems but instead function as a global pool of capacity, orchestrated by software and delivered as a service via API. Where the entire storage footprint is automated, aligned to business policy, and managed as a unified environment. And where IT shifts from executing manual tasks to defining policies that govern global data operations automatically.

That foundation is the Enterprise Data Cloud.

What Is the Enterprise Data Cloud?

The Enterprise Data Cloud is a cloud operating model for data and storage. It brings the simplicity, elasticity, and service-based experience of public cloud to any environment: on premises, in the cloud, or across both. It gives IT full control over performance, protection, and cost, while delivering the speed and flexibility the business needs.

It replaces fragmented data infrastructure with a unified platform that virtualizes block, file, and object storage into a single, policy-driven pool of resources. Instead of operating in isolated silos, storage arrays function as part of a global pool, dynamically orchestrated by software and delivered as services via API.

Data becomes mobile, accessible, and centrally governed, not confined to specific systems or tied to physical locations. Provisioning, scaling, and protecting workloads happen automatically, powered by built-in intelligence and consistent policies.

This is what cloud should feel like: elastic, programmable, outcomes-driven, and always under your control.

What is an Enterprise Data Cloud?

FIGURE 1 Enterprise Data Cloud outcomes

At its core, the Enterprise Data Cloud creates a virtualized cloud of data and storage that spans every environment. Block, file, and object data are abstracted from the underlying infrastructure and managed as a single, integrated data fabric. That eliminates the silos, duplication, and complexity that slow teams down and create risk. Instead of managing storage system by system, IT operates with a unified view. Data is accessible, secure, and governed consistently.

Automation is key. The Enterprise Data Cloud is built around an API-first design that enables integration into DevOps pipelines, compliance frameworks, backup platforms, and cloud automation tools. A customer's global storage footprint can be fully automated and governed according to corporate policy, enabling IT to focus less on tasks and more on intent.

Provisioning, protection, optimization, and policy enforcement can all be automated based on business needs. Workloads are dynamically matched with the right resources and service levels, without manual effort.

A defining characteristic of the Enterprise Data Cloud is how services are delivered. Core capabilities like data protection, compliance, encryption, replication, snapshots, and performance management are integrated directly into the platform. These are not bolted on or managed through separate systems. They are built in. This simplifies governance and recovery, strengthens security posture, and ensures consistency across the board.

Critically, the environment supports a broad range of applications, each with the service level it needs. Traditional enterprise workloads and modern, cloud-native services can all operate in the same environment, while getting the performance, availability, and protection they require. You don't need different infrastructure for every workload. You just need a platform that can adapt.

And it's not just flexible in architecture. It's flexible in how it's consumed. The Enterprise Data Cloud is delivered as a service, aligned to measurable outcomes through guaranteed performance, uptime, and efficiency SLAs. IT teams no longer manage capacity manually or chase infrastructure refresh cycles. Instead, they consume what they need and when needed, and focus on delivering business value instead of managing systems.

At its core, the Enterprise Data Cloud puts IT back in control. It enables organizations to shift from manual, reactive data management to a software-driven model governed by policy. That's how teams regain visibility, consistency, and agility across the data landscape and move from infrastructure operations to business outcomes.

How to Build Your Own Enterprise Data Cloud with Pure Storage

You need more than automation tools or cloud connectors to build a true Enterprise Data Cloud. You need a platform that's architected for it from the start. The Pure Storage Platform is built on that foundation, bringing together always-modern infrastructure, a virtualized data layer, intelligent control, and integrated services.

This isn't a collection of products. Our platform has a unified architecture designed to support consistent data operations across every environment. It's how you eliminate silos, reduce risk, and move faster without sacrificing visibility or control.

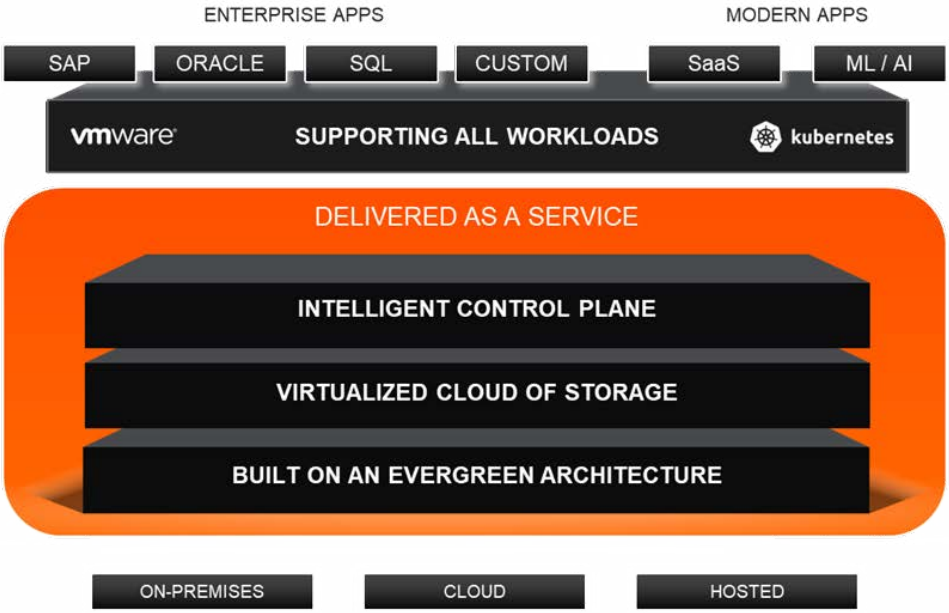

FIGURE 2 The Enterprise Data Cloud architecture.

Here's what makes it work:

- **Evergreen architecture.** At the heart of the Enterprise Data Cloud is Evergreen Architecture. This is Pure Storage's foundation for continuous innovation without disruption. It eliminates traditional lifecycle bottlenecks by allowing systems, controllers, software, and capacity to be upgraded in place without downtime or data migration.

  There are no forklift refreshes. No rearchitecture every three to five years. With Evergreen, always running on modern technology, without the risk and cost of legacy refresh cycles.

- **Virtualized cloud of storage.** Built on our Evergreen® architecture is a virtualized cloud of storage that unifies block, file, and object into a seamless environment across on-premises, hosted, and public cloud infrastructure. This is delivered through FlashArray™, FlashBlade®, and Pure Storage Cloud Services, with orchestration and pooling handled by Pure Fusion™. Pure Fusion abstracts infrastructure into a single, software-defined environment. Storage is no longer tied to individual systems. It becomes part of a global pool that can be provisioned, consumed, and scaled through policy. This is where the EDC model comes to life. Storage behaves like a cloud service: elastic, automated, and delivered on demand, but with complete enterprise-grade control

  This also changes how organizations manage data access. Teams can access the original source directly instead of creating copies for every workflow—AI, analytics, dev/test, or compliance. Data remains consistent, governed, and under control. You reduce duplication, reclaim capacity, and eliminate the sprawl that clutters legacy environments. With a virtualized cloud of storage, data becomes shareable, not siloed, opening the door to faster insights, cleaner operations, and far less overhead.

- **Intelligent control plane.** Managing at this scale takes more than visibility. It takes intelligence, automation, and deep integration across your environment. Pure1® gives teams real-time insight into what matters most: performance, risk, capacity, and cost. But it's the AI behind the scenes that makes the difference. It analyzes telemetry across your entire footprint, spots patterns humans would miss, and identifies issues before they disrupt workloads. Instead of chasing alerts, you get recommendations with context—clear actions you can take.

  That intelligence doesn't just live in a dashboard. It drives automation. With Pure Fusion, teams can orchestrate provisioning, resource allocation, and policy enforcement through a unified control layer. And because it's built API-first, storage becomes programmable. You can integrate it directly into DevOps pipelines, compliance workflows, and broader automation stacks. This shifts operations from reactive to proactive. Routine tasks are automated, resources are used more efficiently, and teams can focus on delivering outcomes rather than managing infrastructure.

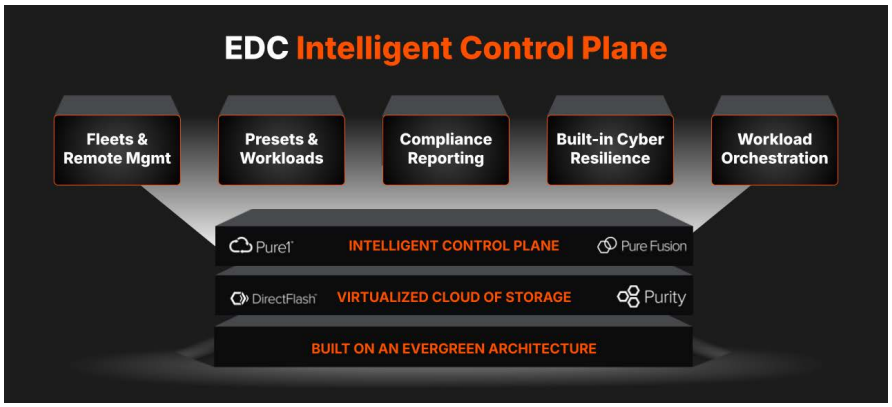

FIGURE 3 EDC Intelligent Control Plane

- **Integrated services.** Integrated services turn infrastructure into a complete data platform. Pure Storage provides enterprise data storage and protection services as part of the core architecture. This includes thin provisioning, deduplication, replication, snapshots, QoS, and multi-protocol support. All essential for performance, efficiency, and flexibility. These services are available across all environments and workloads, without manual tuning or extra tools.

  Protection and security are also built in. SafeMode™ Snapshots, policy-based retention, encryption, and role-based access are always on. Compliance and recovery are simplified because the architecture enforces policy and resilience by design.

- **Workload consistency.** Whether you're supporting traditional enterprise applications or modern, cloud-native services, the experience stays consistent. Databases, virtual machines, containerized apps, and AI workloads all run on the same foundation. With VMware, Kubernetes, and Portworx® support, Pure Storage enables a unified operational model across all workloads.

  That means teams don't have to manage separate infrastructure for different application types. Everything shares the same protections, automation, and visibility.

- **Flexible delivery.** Evergreen//One™ delivers the Enterprise Data Cloud as a fully managed service, backed by guaranteed performance, availability, and efficiency SLAs. It's outcome-based by design. You don't just deploy infrastructure—you consume a platform aligned to business needs, with built-in resilience, predictability, and scale. You get the full power of the Pure Storage Platform without owning or operating the hardware. Capacity scales elastically. You pay only for what you use. And the experience mirrors public cloud, delivered wherever you need it—on premises, or across hybrid environments. Infrastructure fades into the background, and the platform just works, delivering consistent outcomes with enterprise-grade control.

Building your Enterprise Data Cloud with Pure Storage means more than modernizing storage. It means implementing a platform that simplifies your operations, strengthens your protection, and accelerates data agility, and you can do it all without compromise.

Why the Enterprise Data Cloud Changes the Game

The architecture matters, but outcomes are what count. The Enterprise Data Cloud doesn't just modernize infrastructure. It fundamentally changes how IT teams operate and how the business grows.

With traditional workflows, every change carries hidden delays, manual effort, and operational risk. Provisioning can take days. Upgrades require planning windows. Backup and recovery feel disconnected. Any attempt to move data across environments adds cost, complexity, or both. The model may hold together, but it doesn't scale.

The Enterprise Data Cloud breaks that cycle. When isolated systems are replaced with a unified platform, each workflow becomes faster, more secure, and fully automated.

Here's what that looks like in practice:

- **Provisioning** becomes a self-service experience. Volumes are deployed in under an hour using policy-based templates. DevOps teams can move without waiting. Security is built in, not retrofitted.
- **Performance management** shifts from reactive tuning to real-time insight. ML-driven telemetry helps detect anomalies, prevents performance slowdowns, and even catches ransomware in its tracks, without needing to bolt on siloed tools.
- **Upgrades** stop being disruptive events. With Evergreen architecture, systems are refreshed in place, without downtime or data migration. Modernize continuously, without breaking anything.
- **Data protection** is fully integrated. SafeMode Snapshots, automated retention policies, and clean-room recovery are built into the storage platform. You don't just back up. You protect by design.
- **Cloud mobility** is seamless. With the same operating environment on premises and in the cloud, workloads move without rewrites or re-platforming. You avoid egress fees and gain the flexibility to optimize by cost or performance.

These are not just operational improvements. They are business accelerators.

The Enterprise Data Cloud reduces time to value for every project by simplifying the stack and eliminating manual workflows. It strengthens cyber resilience, lowers operational risk, and frees staff to focus on transformation instead of tickets. It turns storage from a cost center into a force multiplier.

When the entire platform is always current, you're never stuck waiting for the next upgrade window or trying to keep up with the business's needs.

This isn't just a better way to run storage. It's a better way to run IT.

The Bottom Line

Enterprise data has outgrown traditional storage strategies. Fragmented infrastructure, inconsistent governance, and operational sprawl have made it harder to protect data, control costs, and move at the speed business demands.

The Enterprise Data Cloud changes that.

By consolidating data and storage into a unified, intelligent platform, IT teams regain control across every location, workload, and service level. Integrated services and built-in automation replace manual work and tool sprawl. SLAs ensure performance, resilience, and cost predictability. And a single control plane makes it all visible, manageable, and programmable.

This isn't just simplification. It's a shift to a model that’s scalable, secure, and built for outcomes.

With the Pure Storage Platform, the Enterprise Data Cloud isn’t a concept. It’s infrastructure that’s ready to run right now.

To learn more, visit https://www.purestorage.com/pure-advantage/platform.html

CM1461-01-en 06/25