

# BUILD AN ON-DEMAND AND AUTOMATED INFRASTRUCTURE

Seizing opportunities to advance your data-centric business strategy.

You might have heard that “data is the new currency,” referring to its dominant role in today’s economy. Data-centric organizations have the power to disrupt traditional ways of working and create new economies. However, data alone is inherently powerless. It doesn’t actually do anything unless you know how to use it — or at least have a guide to start the process.<sup>1</sup> Like currency, the potential value of data remains untapped by organizations that don’t take purposeful action to invest it.

Change is hard, but inaction has consequences. The main challenges in tapping that potential value lie in siloed data, clouds, and organizations.

- **Siloed data.** Data is often spread across the enterprise in discrete silos to meet the needs of specific applications. However, this inhibits modern analytics-driven workflows that demand data be available and delivered quickly as a consistent whole, not fractured in data silos. Additionally, cold data, collected as the economics of storing it improved, had its potential frozen in time when it was captured and forgotten, unintegrated with the objectives of your business.
- **Siloed clouds.** Cloud has earned its place in enterprise IT, but that place has so far been separate and distinct. Separate isn't shared, isn't invested, and isn't valuable. Those data silos withhold data's value from your business, constraining innovation and degrading your users' experience.
- **Siloed organizations.** At the same time, organizational silos, inherited from traditional separation of IT builders and technology operators, prevent the agility of DevOps where the builders are the operators and the service is their objective.

Forward-thinking [data-centric organizations](#) constantly reinvent themselves around mobility and data — positioning themselves to ride opportunity into the stream of future revenues — rather than staying in their comfort zone by doing more of the same and watching competitors take the lead. These organizations understand the mere accumulation of data does not create value, but that it must be processed and used. That, in fact, the value of data appreciates the more it’s shared, and it returns multiples when shared the right way, in the right place.

A data-centric architecture positions organizations to make sense of information and turn an avalanche of data into invaluable business intelligence, which can be the deciding factor in separating data economy winners from losers. A data-centric architecture is an approach to designing an end-to-end environment — across compute, network, storage, and cloud — that is optimized for ubiquitous and fast consumption of data. A data-centric architecture is characterized by five key pillars:

- Fast, shared data
- On-demand and automated
- Globally reliable and secure
- Hybrid cloud by design
- Constantly on and improving

This paper focuses on the on-demand and automated pillar of a data-centric architecture.

# Data-Centric Architecture

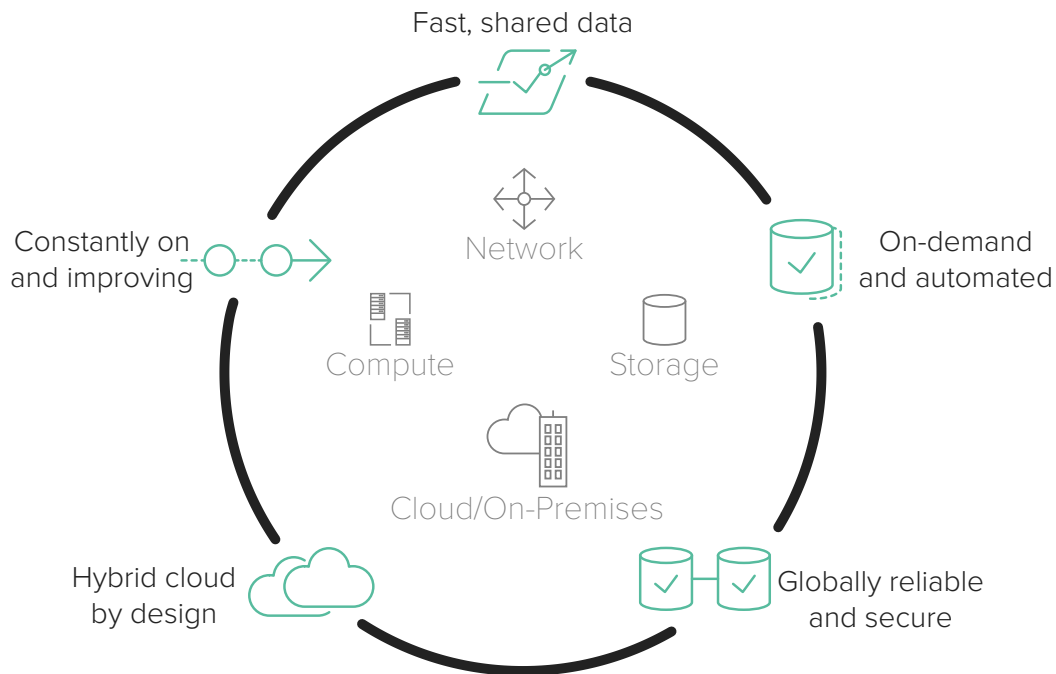


Figure 1. The five defining traits of a *data-centric architecture*.

## Better Service Offerings and Faster User Interactions

Data-centric businesses must be able to develop, scale, and iterate services on-demand, while at the same time reducing the need for IT intervention. Implementing an on-demand and automated design provides two key business benefits. First, it enables IT teams to efficiently develop and quickly roll out repeatable services to internal employees. For example, DevOps teams are tasked with creating and updating a catalog of on-demand services for the internal user community to provision resources, grant access to tools, or provide assistance. Also, product development teams could be creating more robust catalogs of services that are the business' revenue-generating offering to external customers, such as those provided by software-as-a-service (SaaS) providers. The second key benefit is in the consumption of services: internal IT customers can efficiently point-and-click their way to the resources they need, while external customers of paid services benefit from the frictionless ability to sign up, monitor, manage, and use the services they need without a lot of overhead.

In both cases, the benefits of an on-demand and automated infrastructure are reduced cost to deliver services through the efficiency of automation and the upside opportunity that comes with the agility to innovate at speed and maintain happy customers. A data-centric design will work across data silos, allowing automation across data types, owners, and purposes. At the same time, a design that enables DevOps will inherently break down at least one traditional organizational silo that prevents value generation from data: the divide between developer and operator in the technology team. In an *on-demand and automated* services delivery model, the developers *are* the operators, allowing for rapid iteration of service delivery as they experience any obstacles.

## Orchestrating Data Services

A data-centric architecture delivering the ability to be on-demand and automated enables technology teams to plan, implement, and orchestrate workloads across their infrastructures. This is part of knowing what to do in order to release the value of data. The simple state of holding data isn't what creates value. Acting on that data, sharing it, and enabling workflows to use it are part of value creation. As the scale of data and its potential grows, so too does the need to scale IT services with automation of both delivery and consumption.

There's a saying in the DevOps world: "Anything that you do more than twice should be automated." APIs are one of the key elements in building a **data-centric architecture** that allows organizations to optimally store, share, and ultimately analyze data. They enable you to build an on-demand and automated infrastructure that lets a storage team deliver data as a service to development teams.<sup>2</sup>

## Pure Storage® Supports On-Demand and Automated Services

Pure solutions feed an on-demand and automated data-centric architecture with building blocks of standardized file, block, object, and [virtual machine \(VM\) services](#), orchestrated with an API-first portfolio and policy-based automation.

A consistent and extensible REST API enables one authentication that works across all open-source array frameworks, such as Swagger®. Developers can easily discover existing arrays across the organization and access historical performance and capacity data. For FlashArray™ storage, this detailed information includes load, version, model, volume, snapshots, pods, file systems, network interfaces, and other metrics.

Telemetry collection and machine learning (ML) pattern recognition enable automated artificial intelligence (AI) tuning, [predictive workload planning](#), and [proactive support](#). VM analytics deliver full-stack performance troubleshooting and topology discovery.

Pure supports on-demand and automated architecture:

- **FlashArray all-flash block-storage solutions:** Pure all-flash block-storage solutions that support REST APIs, for simple, HTTP-based integration into leading private-cloud platforms.
- **FlashBlade™ all-flash data hub for file and object data:** FlashBlade [file and object data storage](#) enables a truly scalable architecture, from hardware to software, for powerful performance to consolidate various data silos.
- **Pure1® cloud-based storage management:** Powered by Pure1 Meta™, Artificial Intelligence (AI) and Machine Learning (ML) engine, delivering end-to-end visibility of [hybrid data-storage infrastructure](#) across multiple private- and public-cloud environments. The Pure1 platform provides full stack monitoring across VMs and arrays, consistent APIs and SDK, intelligent recommendations, hardware simulations, predictive support, and an optimized workload planner for IT to maximize their storage capacity, cost, and performance.
- **VMware® integrations:** Pure's deep integration with VMware simplifies ongoing storage operations, freeing up time so you can drive more innovation for the business. With the [VMware vSphere® Web Client plug-in](#), you can easily monitor and manage FlashArray from VMware vCenter®, alongside the rest of your virtual infrastructure. Integration with VMware vRealize® Suite enables comprehensive operations management along with automation and self-serviceable storage services. For instance, pre-built Pure workflows with [VMware vRealize® Orchestrator™](#) provide the most common tasks involved in storage orchestration, including setup of per-VM-level storage provisioning, with policy-based control through support of VMware vSphere® Virtual Volumes™.
- **Cloud Block Store:** Move data and applications more freely. [Cloud Block Store](#) delivers storage services, resiliency, and APIs consistent with Pure's on-premises infrastructure, so you can seamlessly run your existing apps in the cloud. Connect to storage and automate deployments in the cloud in the same way as you would on-premises. Develop your apps once and deploy seamlessly across your hybrid cloud.

## Build a Data Foundation That Is On-Demand and Automated with Pure

Reaping the benefits of an on-demand and automated strategy requires a programmable, **data-centric architecture** that makes service delivery as simple as DevOps offering services and users consuming them.

Pure can help you build an architecture that is on-demand and automated, allowing you to provide scalable, elastic services to applications and developers on-demand, with policy-based automation that minimizes, and even eliminates, the need for human intervention.

For more information about building a data-centric architecture that supports your needs for on-demand and automated, contact your Pure representative today.

© 2019 Pure Storage, Inc. All rights reserved. Pure Storage, Pure1, Pure1 Meta, the P Logo, FlashArray, and FlashBlade are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. All other trademarks are registered marks of their respective owners.

<sup>1</sup> Digitalist. "Data: The New Currency." December 2017  
[www.digitalistmag.com/cio-knowledge/2017/12/11/data-new-currency-05592449](http://www.digitalistmag.com/cio-knowledge/2017/12/11/data-new-currency-05592449).

<sup>2</sup> Pure Storage. "Get a Sneak Peek at Self-Driving Infrastructure Management from Cisco and Pure Storage." June 2018.  
<https://blog.purestorage.com/get-a-sneak-peek-at-self-driving-infrastructure-management-from-cisco-and-pure-storage/>.