



RETHINKING HYBRID

HOW DATA SERVICES BRIDGE CLOUD WITH ON-PREMISES

Introducing Astraeus, OpenDrives'
cloud-native data services platform



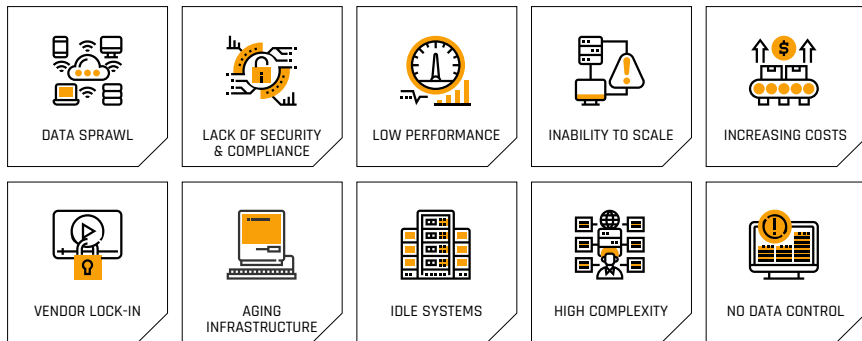
Today, IT directors in enterprise organizations realize that the expectations for traditional cloud or on-premises IT infrastructure remain unfulfilled. That's because IT infrastructure has not kept up with their business needs. As a result, workflows continue to fracture, data siloes continue to spread, and innovation continues to stall. Instead of completing projects and proving value, you and your teams end up spending more time hunting for data, using disparate, monolithic tools, and troubleshooting complicated workflows. So if IT infrastructure can't solve your data and workflow challenges, what can?

IT Infrastructure Expectations

Traditional infrastructure often bears a heavy load. Does yours measure up with the following:

- Cost-predictable resource optimization and scalability
- Easy data access and availability
- Strengthened security
- Reduced operational and compliance risks
- Increased productivity and innovation

DO ANY OF THESE DATA AND WORKFLOW CHALLENGES SOUND FAMILIAR TO YOU?



IT Infrastructure Doesn't Solve Workflow Troubles

Unrealistic expectations have been placed on IT infrastructures. Most of these expectations were unknowingly or casually determined. However, the truth is, IT infrastructure cannot solve workflow challenges because it is not the right tool for the job.

Let's face facts. IT infrastructure is an organization's structural lifeblood for building workflows. But expecting infrastructure to transform data into

business outcomes is not only unrealistic; it's unlikely. After all, infrastructure is just a passive vessel that holds data rather than an active enabler of workflows. Modern workflows demand cloud-native flexibility and scale, real-time data access, and the ability to adapt instantly to changing bandwidth and performance needs; on-premises, cloud, and hybrid infrastructures can't do any of this. Standard infrastructure can't keep pace with evolving, resource-intensive workflows like edge computing, Artificial Intelligence (AI), or Machine Learning (ML) that consume unprecedented

amounts of processing power. Without re-architecting or starting all over again, this reality is what prevents infrastructure from repairing broken workflows.

You can use workarounds to force infrastructure to perform certain data services, but the end result is often a cumbersome environment that cultivates data sprawl and tool fragmentation. Since standard infrastructure offerings fall short, the best path forward requires an unexpected shift—one that rethinks where and how your data and workflows coexist.

BACK TO INFRASTRUCTURE BASICS: IT INFRASTRUCTURES AT A GLANCE

PHYSICAL DRIVES

HARD DISK DRIVES
(HDD)SOLID-STATE DRIVES
(SSDs)

Hardware components that store digital media

CLOUD TYPES

**Public Cloud**

Resources hosted by third-party cloud providers

Private Cloud

Resources dedicated to a single organization

Hybrid Cloud

Combines public and private cloud environments

SAN VS NAS

STORAGE-AREA NETWORK



Block-level access

Fibre-channel
dedicated network

Excellent scalability

Expensive to implement

Complex to manage

NETWORK-ATTACHED STORAGE



File-level access

Easy to use and
implement

Ethernet network
connection

Economical scalability
for shared projects

So how did we get here? The global pandemic drove organizations to take calculated risks and adapt to a rapidly changing world, including lifting and shifting out of on-premises and exploring cloud-first or cloud-only strategies. There was no choice. Surviving chip shortages, supply chain constraints, and inflation, and adjusting to new remote workflows meant doing things differently. Eventually, the public cloud became the primary data home, but this came at a high cost.

60%

In response to the global pandemic, nearly 60% of organizations increased cloud usage in 2020. This rapid shift aligned with strategic goals, but fueled a new set of operational and governance challenges. Today, organizations seek a new path forward.

YOUR IT INFRASTRUCTURE CHECKLIST

Is it time to rethink your infrastructure? If you find yourself agreeing to the following, it's time to make another change.



You want an easy way to orchestrate and manage your Kubernetes containers and monolithic applications all on one central platform.



Cloud costs keep rising, continued cloud spending is no longer practical or affordable, and you find that the cloud isn't always the right home for all your workflows.



You want the convenience of the cloud, but you also need the security, performance, and control over your data that on-premises infrastructure provides.



You're looking for a solution that integrates into your existing workflows and environment, without having to re-architect the applications you've come to rely on.



You need to future-proof your environment and prefer a flexible solution that can evolve with emerging technologies and won't lock you into one vendor or application.



Studies show organizations are overspending by 25 to 40% on unnecessary cloud costs.

Get Your Data Out of the Clouds

Cloud reliance comes at a steep price. Some costs, like paying for more resources, are quantifiable. And don't forget egress fees—moving data in and out of the cloud is not an inexpensive proposition. Other costs are less tangible, but equally as concerning, such as: more complexity, high security risks, bottlenecks that slow down productivity, and strained resources. The cloud also requires upfront investments in terms of time, upskilling, and the hiring of dedicated staff. No wonder many feel stuck in the cloud infrastructure they spent so much time building.

CLOUD GLOSSARY



Cloud Native

Building and running scalable, resilient applications in dynamic cloud environments using containers, microservices, and declarative Application Programming Interface (API) technologies.

Cloud First

New workloads are deployed in the public cloud initially, while legacy systems remain on-premises. This hybrid approach adds flexibility, but makes it challenging to find, place, and maintain data.

Cloud Only

All workloads run in the public cloud, eliminating on-premises infrastructure. This boosts scalability and agility, but increases reliance on cloud providers, raising risks around cost, security, reliability, and performance.

Data Home

Authoritative location where data is stored, managed, and maintained.

▲ VERSUS ▼

Data Store

Temporary location for active data before it's moved elsewhere.

You can have multiple data stores, but only one primary data home.

Cloud or On-Premises: Where's the Right Data Home for Your Workflows?

The answer is you shouldn't have to choose. You should access your data from wherever it lives with one data services platform.




Choose on-premises if you:







- Want cost predictability
- Need customization, interoperability, enhanced security
- Require always-on performance for modern workflows
- Expect high growth and need to scale out
- Looking for intelligent storage and data management services

Choose cloud-based if you:

- Prefer a pay-per-use model
- Need flexible scalability
- Look for easy accessibility
- Require fast and affordable implementation
- Don't want to worry about operational overhead and maintenance costs

Modern Cloud-Native Development



	Cluster Group of connected machines that operate as a unified system to run and manage applications in a modern, distributed environment.
	Kubernetes Open-source orchestration platform that automates deploying, scaling, and managing containerized applications.
	Containers Portable, lightweight units that package applications and their dependencies for fast, consistent deployment in clustered environments.
	Microservices Building applications as collections of small, loosely coupled services that work together—often deployed in containers for flexibility and scale.
	Elasticity Ability to automatically scale resources based on workload demand, optimizing performance and reducing costs.
	Virtualization Creating virtual versions of physical resources so that multiple virtual machines (VMs) run on a single physical system.



"We're starting to see a trend toward repatriation, people moving back from public cloud to deploying in private cloud and on-premises. This is resulting in a different kind of deployment process where people are spending more time evaluating where to best put their workflow, as opposed to defaulting to the public cloud, because broad data accessibility is now less critical than the overall ROI and the total cost of that deployment. On-premises, you can manage and control your data more cost effectively and securely."

OpenDrives' Senior Director of Integrated Strategy,
Michael Wilsker

IABM ChainTracker "Store" Report



The Path Forward: From Infrastructure to Data Services

Knowing that IT infrastructure alone is not entirely to blame for your data and workflow challenges, what comes next? It's time to start rethinking where your data lives and how it's accessed.

Data repatriation is the process of moving data and workflow applications from a public cloud environment back to on-premises or hybrid cloud. This approach provides your organization with the flexibility you need to better position your data to deliver the greatest value. The optimal data home helps to streamline cost and performance, enhance productivity, and ensure resources are aligned with strategic priorities—something the public cloud nor on-premises can always deliver. IT infrastructure alone is not the solution; what you actually need is a new kind of hybrid infrastructure driven by data services.

Data services enable organizations to move beyond static infrastructure and into modern workflows that enable the generation of real value. Data services help you use and control your data. They provide the functional ability to move, store, process, manage, protect, organize, and so much more. In fact, you can consider the very workflow applications you and your team rely on to be data services. When data services and your workflow applications work together in one data home, your work just flows.

The problem is legacy infrastructures don't provide you everything you need in one place. What you need is a single platform that combines the data services you need with the workflow applications you depend on, to get the job done—your way.

A data service is a software application that securely delivers, processes, or manages data on request using an API. It's really an application that does something significant with data to produce a business outcome. In cloud-native environments, these data services are abstracted and activated only when needed, maximizing efficiency and scalability. With a data services platform, you get next-generation data services that work alongside your critical third-party applications for the ideal workflows you've always wanted.

Let Your Data Go & Let Your Work Just Flow

Introducing Astraeus, Your New Data Home

Astraeus is a dynamically scalable data services platform designed around cloud-native principles. This comprehensive platform enables you to configure, store, orchestrate, manage, secure, and deploy your mission-critical applications across dynamic environments, whether it's on-premises or in public, private, or hybrid clouds.

With Astraeus, you get the following:

- a unified namespace to seamlessly manage data sprawl across various data stores
- performant storage management and intelligent compute and resource optimization to ensure your critical workflow applications are accessible, available, and extensible
- dynamic scalability, and reliable data protection and resilience

To solve evolving data and workflow challenges, you need a single platform that combines cloud-like resiliency, familiarity, and experience with on-premises benefits. The result? You get everything you need: predictable costs, meaningful productivity gains, and more available and durable data. You get a workflow that just works.

astraeus

cloud-native data services platform

Scalable Services Manager



Cloud-Native
Customer Applications



Analytics



Management
Tools

Astraeus Data Services Controller (ADSC)



Data Stores
Management



Security



Networking



Monitoring

Astraeus Application Core

Astraeus OS

Infrastructures that Astraeus runs on



One-node



Multi-node



Private
Cloud



Public
Cloud



Scale out and
manage as many
clusters as needed

One
cluster



Multiple
clusters





astraeus
cloud-native data services platform

Astraeus is much more than IT infrastructure, storage management, or a Kubernetes orchestration product. It's a data services platform that modernizes and extends your workflows so that you can intelligently and automatically transform data into real business outcomes. It helps you find, use, and monetize data wherever it lives (using a combination of OpenDrives' built-in data services and your applications). It's an open platform with a cluster-first design that streamlines and future-proofs your ever-evolving workflows. It provides the cost control, security, and data governance you need in one easy-to-use, flexible platform. No re-architecting. No DIY. No lock-in.

With Astraeus, you can let your data go where it needs to go so that your work can just flow.

Astraeus provides data services that provide clear business outcomes with your data:

High-Performance Data Storage Management: Boost productivity and streamline operations with optimized and performant storage management, one of Astraeus' core data services. Unify critical workflow applications from various data stores onto a single platform, and optimize resource utilization, processing, provisioning, capacity, and performance across teams with precision. Quickly aggregate, access, and share resources with always-on availability and data integrity, eliminating concerns around resource allocation, overspending, idle systems, or wasting time hunting for data.

Dynamic Resource Allocation and Orchestration: With Astraeus' Scalable Services Manager, you gain explicit control over scheduling and allocating the right amount of storage and compute resources your business needs for smoother operation. Choose to spin up or down specific data services and when to turn them on or off. Grant more processing power as you build clusters, launch applications, and extend your workflows. With dynamic resource allocation and orchestration, you're in the driver's seat of the environment you want to create.

Flexible Deployment and Scalability: Leveraging cloud-native and software-defined principles, Astraeus connects all your data stores, whether on-premises or in the cloud, enabling you to scale out and manage as many clusters as you need. Its flexible, scalable design simplifies integration into a single environment that will evolve with your business and support emerging technologies and workflows. No data sprawl. No vendor lock-in. Just the freedom to architect your ideal workflow.

Autonomous Application and Cluster Management: Aggregate your cloud-native applications and clusters into Astraeus and manage them effortlessly with the Astraeus Data Services Controller. Its Kubernetes-based framework preserves everything you've already built, so that there's no need to re-architect what already works. Deploy entire workflows seamlessly, while autonomous application and cluster management services handle the monotonous tasks and heavy lifting. Now, you can focus on letting your data work for you—not the other way around.

Unified Data Plane: Break down silos and reduce operational friction by repatriating cloud-native applications from distributed data stores into a single, unified system. Consolidate unstructured data stores, object stores, and structured data sources in one data plane that adheres to your organization's access control, governance, and sovereignty policies.

Data Optimization & Tiering: A growing list of data optimization and tiering services help you find and access needed data across all connected data stores. A policy-based Hierarchical Storage Management (HSM) layer automates data movement for greater efficiency and cost control. Monitor key usage insights to improve performance, accelerate project completion, and reduce risks.

Data Governance & Protection: Safeguard your most valuable assets with protection and governance services that ensure data is shielded, secure, and available throughout its lifecycle. Minimize disruption with proactive business continuity practices, and in the event of a disaster, quickly recover data and resume critical operations.



LET YOUR DATA GO & LET
YOUR WORK JUST FLOW.

**READY TO MAKE ASTRAEUS
YOUR NEW DATA HOME?**

hello@opendrives.com
opendrives.com
310-659-8999