



Accelerating Safely:

Linking Traditional Enterprise Systems to the Modern Data Stack. Without Migrations.



Executive Summary

Across industries, enterprise organizations are under increasing pressure to extract more value from data—faster—while contending with fragmented systems, legacy infrastructure, and compliance demands.

Yet, data integration remains a persistent bottleneck. Traditional approaches are slow, rigid, and often incompatible with complex enterprise environments.

This paper argues that enterprises can modernize their data strategy without costly or disruptive migrations. By adopting Dataddo—a platform purpose-built for hybrid and regulated environments—organizations can unlock legacy systems, accelerate data integration by over 90%, and enhance agility while maintaining strict operational control.

Dataddo enables a shift from control-centric infrastructure to insight-driven—operations without compromising compliance, uptime, or performance.

Agility vs. Control: A False Trade-Off

Enterprise data architecture is under dual pressure: to accelerate time to insight while preserving stability, security, and compliance. One of the biggest blockers is the delay between when data is created and when it becomes usable—whether for dashboards, ML pipelines, or regulatory reporting. This bottleneck almost always appears in the integration layer.

Foundational databases like Informix, DB2, and Sybase still run critical workloads in sectors like insurance, finance, and healthcare. But these systems **are notoriously difficult to integrate with modern tools**. Traditional solutions require invasive migrations, on-host agents, or full platform rearchitecture—slowing proof-of-concepts and increasing operational risk.

So how can enterprises achieve both agility and control?

By connecting valuable data from foundational systems to modern applications rapidly and securely—without migrations or platform disruption—Dataddo helps enterprises move fast without sacrificing the reliability, compliance or stability these environments demand.

This paper describes how Dataddo supports modernization initiatives in organizations heavy with foundational systems **by combining direct extraction from older systems with the reliability, observability, and compliance tooling expected of a modern data stack**. Through its architecture and control-plane approach, Dataddo enables repeatable data movement while maintaining strict operational boundaries and auditability.

Works anywhere, without disruption

For complex enterprises, **the architecture priority is not to replace systems, but to connect them.**

Despite their age, foundational platforms cannot be dismissed—they are still the operational core for thousands of business processes. What the platforms lack is not relevance, but agility.

Dataddo's role is to sit next to these systems, extract what matters, and deliver it—continuously, reliably. It enables modern teams to access valuable data and respond without delay.

The platform's containerized design and Kubernetes compatibility allow it to run anywhere:

- **Fully on-premise**—even in air-gapped networks
- **As a hybrid setup**—with data agents behind the firewall and a cloud-based control plane
- **In a private VPC** under client governance




Figure 1. Dataddo hybrid on-prem architecture at a glance

This flexibility ensures that Dataddo fits the client—not the other way around. And critically, it does so without middleware, polling agents, or proprietary SDKs. There's no lock-in, no host-level installation, and **no disruption to existing systems.**

Embedded Compliance and Security.

- Every execution and change is logged, auditable, and attributable.
- Role-based access control (RBAC), single sign-on (SSO) environment separation.
- Certified to SOC 2 Type II, ISO/IEC 27001
- Complies with global standards: GDPR, HIPAA, CCPA, LGPD, and POPIA.
- All data is encrypted in transit
- Key management via KMS
- Sensitive data hashing at the source



Accelerated time to insight, without complexity

One of the most persistent challenges in enterprise architecture today is the delay between when data is created and when it becomes useful. Whether the goal is a business dashboard, a machine learning pipeline, or a regulatory report, **the bottleneck almost always appears in the integration layer.**

Dataddo removes that bottleneck by delivering speed at every layer of the pipeline: fast POC delivery, rapid backfilling, short incremental sync cycles, and predictable recovery timeframes. Even full dataset reloads—often needed for test/dev or after schema changes—can be completed in hours, not days.

This performance is enabled by Dataddo's **Mesh Ingestion model**: a parallel extraction method that replaces slow, serial ETL jobs with concurrent data readers. Unlike traditional tools that run one job per source, Dataddo distributes the load—reducing extraction time dramatically without risking system stability.

And because it requires no agents, custom runtimes, or host-side software, Dataddo delivers this performance **without adding complexity to your architecture.**

In real-world benchmarks, **Dataddo replicated over 1TB of data in just 2.5 hours**—12x faster than agent-based tools, and without infrastructure changes. And because the platform is database-agnostic, this speed holds across cloud platforms, foundational databases, and modern warehouses.

By connecting legacy systems to modern tools at high speed and low friction, Dataddo ensures that insight keeps pace with business.

Faster deployments, Faster Results

In large organizations, even modest integration projects can stall under the weight of legacy constraints, security reviews, or vendor timelines. But **Dataddo is engineered to bypass those roadblocks.** It deploys quickly, without needing agents, SDKs, or rearchitecting your infrastructure.

Pipelines can be configured in hours—even in highly segmented, compliance-heavy environments. When custom connectors are needed for internal tools or private APIs, Dataddo delivers them within two weeks—without requiring a typical six-month roadmap cycle or professional services contract.

As a result, proof-of-concepts launch in days, not quarters.

This isn't theoretical. A case study later in this paper describes how a large European insurance company—with foundational systems and tight change controls—successfully deployed Dataddo, validated performance, and began ingesting billion-row tables within days of initial setup, despite significant procedural overhead.

For enterprise teams under pressure to show results fast, **Dataddo reduces the gap between idea and implementation** to its minimum.

CASE STUDY:

How a European Insurer Deployed Dataddo in Days

A major European insurance company initiated a proof-of-concept to evaluate Dataddo's ability to operate in a highly controlled enterprise environment. Their infrastructure included foundational systems like **MSSQL and Informix**, **strict network segmentation**, and **deeply entrenched change management protocols**—conditions that had historically slowed down technical implementation.

Despite these constraints, **Dataddo was deployed without host-side installation or architectural disruption**. Within days of initial setup, the platform began ingesting billion-row tables via full and incremental loads. **Metadata-driven orchestration**, **recovery workflows**, and **schema evolution** were all **validated** against internal protocols. No core systems were restructured. No manual workarounds were required.

The project moved from POC to production rollout faster than any previous integration effort—without compromising compliance, stability, or control.

Use Case



Goal: Full replication of ~600M+ rows from an Informix database into Azure



Constraints: Legacy systems, segmented networks, strict change control



Result: Billion-row ingestion live within days

Performance Comparison

Full data replication runtime (~1TB equivalent load)

Solution	Runtime	Notes
In house developed solution	30+ hours	Requires full platform contract and setup overhead
Traditional Agent-based tool	12–18 hours	Requires on-host software, challenging to scale or adapt
Dataddo	2.5 Hours	12× faster. Parallelized on 8 CPUs / 32GB RAM using 4 parallel workers

Agility, without losing control

In modern enterprise environments, speed is only part of the equation. True modernization requires data infrastructure that is **observable, auditable, and compliant** with strict regulatory frameworks. Dataddo delivers on all fronts.

Every pipeline is fully transparent—monitored through logs, ddashboards, and real-time alerts. Built-in integrations with enterprise observability tools like Prometheus, Grafana, Kibana, OpenSearch, Slack, JIRA, PagerDuty, and ServiceNow ensure that **data operations are visible to the right teams, at the right time.**


Compliance is embedded in the platform's design. Every flow execution and configuration change is logged, auditable, and attributable. Role-based access control (RBAC), single sign-on (SSO), and environment separation are native features. Dataddo is certified to SOC 2 Type II, ISO/IEC 27001, and complies with global standards including GDPR, HIPAA, CCPA, LGPD, and POPIA. All data is encrypted in transit and at rest, with key management supported via KMS integration for AWS, Azure, and GCP.

For added data protection, **sensitive fields** can be hashed or masked at the source—ensuring privacy even before data leaves controlled environments.

Even disaster recovery is covered—Dataddo guarantees an RPO under 15 minutes and an RTO under one hour.

In short: nothing about Dataddo's agility comes at the cost of control, reliability, or compliance. It's **built for organizations that need to move fast—without breaking policy.**

Key Factors Behind Dataddo Performance



This approach allows enterprise clients to reduce replication time by up to 90% while staying compliant with existing infrastructure, licensing, and data locality constraints.

Built for the Future of Your Architecture

Enterprise modernization isn't a one-time migration—it's an ongoing process that demands adaptability, continuity, and minimal disruption. In organizations where stability is as important as speed, **Dataddo becomes the connective layer that moves data forward without forcing premature change.**

The platform works alongside your existing systems. It enables a gradual shift from legacy to modern infrastructure by integrating with what's already in place—foundational databases, cloud warehouses, and modern SaaS tools—without replatforming or system retirement.

It also supports the tools and practices of modern engineering. Configurations can be managed in YAML for version control and rollback. The platform integrates with CI/CD pipelines, supports geo-distributed workloads, and plugs into orchestration frameworks natively. This reduces operational overhead, lowers FTE requirements, and improves data availability across teams.

Just as importantly, **Dataddo is ready for what comes next.** When your organization adopts a new tool or platform, chances are it is already supported. A new custom connector can be delivered in weeks, even for private APIs or edge cases. This ability to adapt ensures that **Dataddo continues to add value as your architecture evolves.**

The result is a platform that supports modernization not just today, but at every step of the journey—enabling both stability and scale, on your terms.

Summary

Dataddo is an enterprise-grade data integration platform built for hybrid, multi-system environments.

It connects foundational systems to modern tools—quickly, securely, and without disrupting existing architecture.

In organizations where legacy databases still power mission-critical operations, **Dataddo enables high-speed, low-friction data movement to the cloud.**

Its architecture supports direct extraction, full observability, and compliance tooling without agents, SDKs, or replatforming. Proof-of-concepts launch in days, pipelines are configured in hours, and custom connectors are delivered in weeks.

The platform adapts to your infrastructure—not the other way around—making it uniquely suited for complex, compliance-sensitive environments.

In one real-world deployment, Dataddo ingested over 1TB of data in under three hours—**12× faster than traditional agent-based tools.**

This paper outlined how Dataddo helps enterprises modernize without compromise: **accelerating time to insight**, maintaining **strict operational boundaries**, and enabling **long-term** architectural flexibility.

For architecture leaders and CTOs planning their next phase of infrastructure investment, Dataddo is not just another tool to evaluate—it's a way to unlock the value of existing systems and accelerate modernization on your own terms. The result is a platform that supports modernization not just today, but at every step of the journey—enabling both stability and scale, on your terms.

Ready to modernize on your terms?

To explore deployment options or schedule a technical session, reach out to:

✉ juho.antikainen@dataddo.com

Why Dataddo Is a Strategic Fit



Stability vs. Agility

- Infra teams value security, compliance, uptime
- BI teams need speed, flexibility, adaptability
- Dataddo satisfies both without compromise



Time to Insight Shouldn't Take a Quarter

- POCs in days
- Production flows deploy with zero disruption
- Backfills complete in hours
- Custom connectors in 2-4 weeks



Integration ≠ Migration

- Works alongside existing infrastructure
- No forced replatforming
- Supports hybrid and fragmented environments



Solving the Ingestion Challenge

- Fragmented systems block analytics & automation
- Dataddo's connector-first model removes the bottleneck
- Built for long-term adaptability



Dataddo has been recognized as one of the Honorable Mentions in the 2024 Gartner® Magic Quadrant™ for Data Integration Tools



Dataddo is an enterprise-grade data integration platform designed for hybrid multi-system environments.

We connect, transform, and deliver data across modern and foundational infrastructures with minimal engineering overhead. Our platform adapts to your systems, not the other way around.

We are the only EU-based data integration platform built for complex enterprise needs, offering region-compliant scalability and future-proof architecture that supports digital sovereignty without relying on US cloud ecosystems.