LTM | datamotive
eliminating cloud boundaries

WHITEPAPER

Cross-Cloud Disaster Recovery with Datamotive

Nitin Kumar
Senior Principal – Architecture, Cloud & Infrastructure, LTM

Tushar Tarkas
Director of Technology, Datamotive

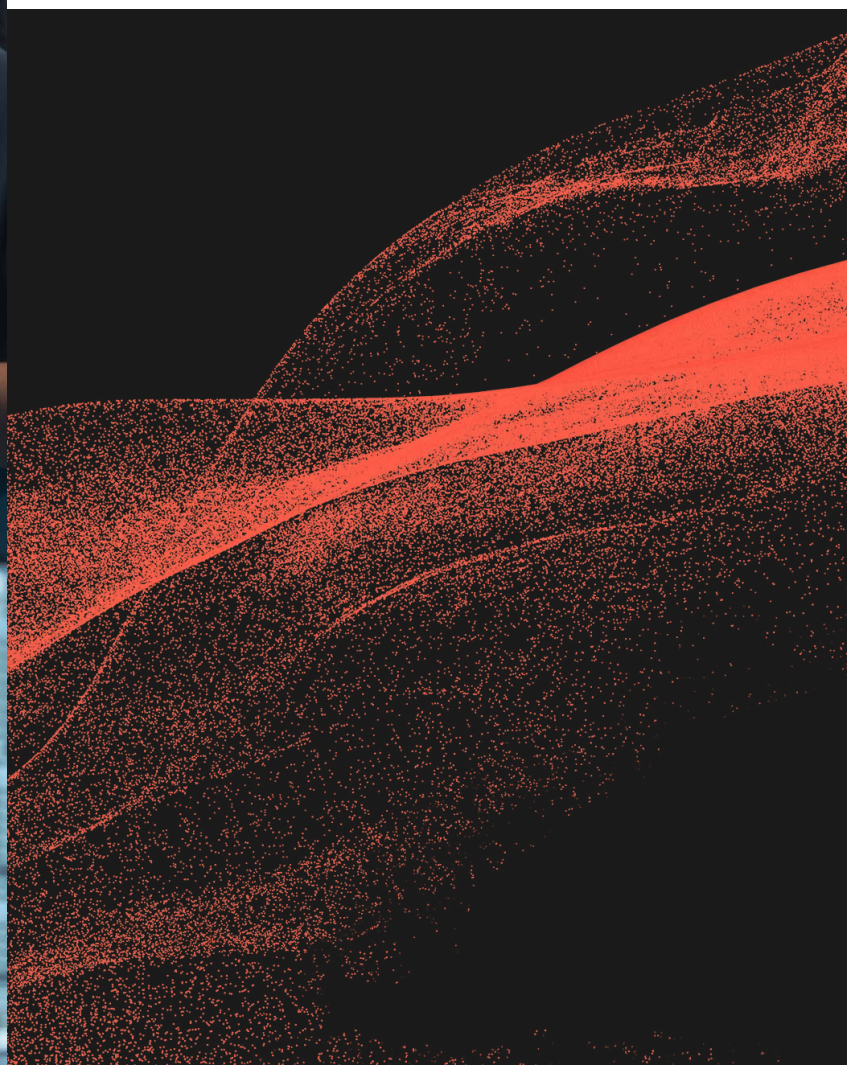

Table of Contents

Executive Summary 3
Introduction 3
Legacy DR Challenges 4
Overcoming DR Challenges with Datamotive 5
Integration with Nutanix, VMware 9
How the Datamotive DR Platform Leads the DR Landscape 12
Risks and Mitigation 13
Use Cases 13
Industry Aligned Methodology and Strategy for Datamotive 15
Conclusion 17
References 17
Authors 18

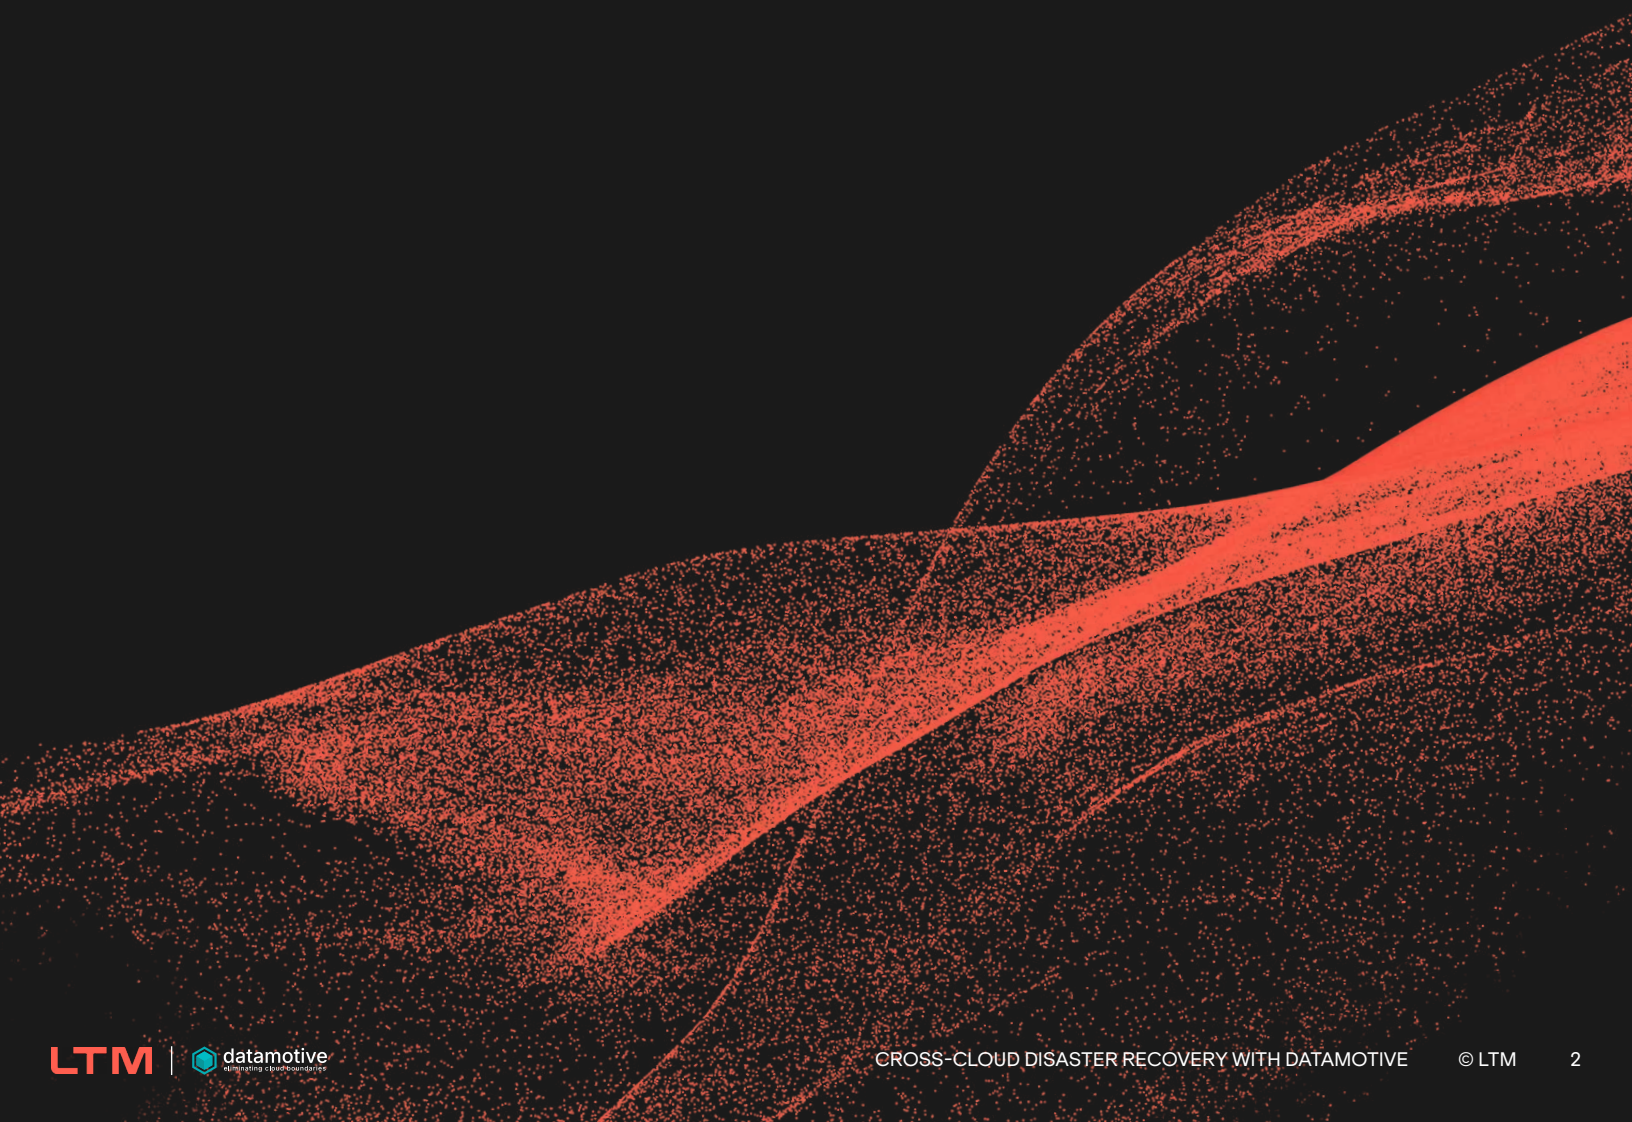

Executive Summary

Business disruption caused by system outages entails significant financial consequences, including lost revenue, regulatory exposure, and increased recovery costs. Moreover, traditional disaster recovery solutions often require substantial capital investment and ongoing operational expenses, limiting financial flexibility. Datamotive offers a cost-optimized, cross-cloud disaster recovery solution that protects business operations while reducing the total cost of ownership. By eliminating the need for physical disaster recovery infrastructure and enabling recovery across existing cloud platforms, Datamotive transforms disaster recovery from a fixed cost burden into a more efficient operating model. Organizations using Datamotive have realized up to 60% reduction in disaster recovery costs, reduced failback effort by up to 70%, and achieved up to 65% lower overall cost of ownership. Rapid recovery capabilities minimize downtime-related revenue loss, while standardized DR processes reduce operational overhead and reliance on manual intervention. Proven in global deployments, Datamotive delivers a clear return on investment and strengthens resilience, reduces operational complexity, and scales disaster recovery alongside evolving cloud strategies, without compromising security or performance.

Introduction

Business continuity is mandatory in the digital-first world that we live in today. Unplanned downtime, cyber threats, and high regulatory demands are mounting pressure on organizations, and a well-developed disaster recovery plan is an essential component of IT resilience. Any downturn in availability can lead to massive financial losses, reputational damage, and even operational disruption.

This has seen disaster recovery becoming a board-level issue instead of merely being a technical issue. In addition, traditional disaster recovery solutions are becoming less effective. They are also expensive to maintain, complicated to use, and not well adapted to modern, heterogeneous, and multi-cloud environments. These constraints deny organizations the agility, scalability, and cost efficiency needed to meet changing business needs.

Datamotive solves these issues by offering an cost-effective cross-cloud disaster recovery solution that will work across various workloads and cloud platforms. The solution is based on high-performance replication technology that Datamotive has tested and proven to guarantee data protection, fast recovery, and reliable failback, without undermining enterprise security, governance, and compliance requirements.

Datamotive helps organizations to protect business continuity with ease by providing predictable recovery times, lowering the cost of infrastructure, and simplifying disaster recovery operations. Datamotive offers flexibility, reliability, and control to be successful in a multi-cloud environment, whether planning to use the cloud or to react to an outage that was not planned.

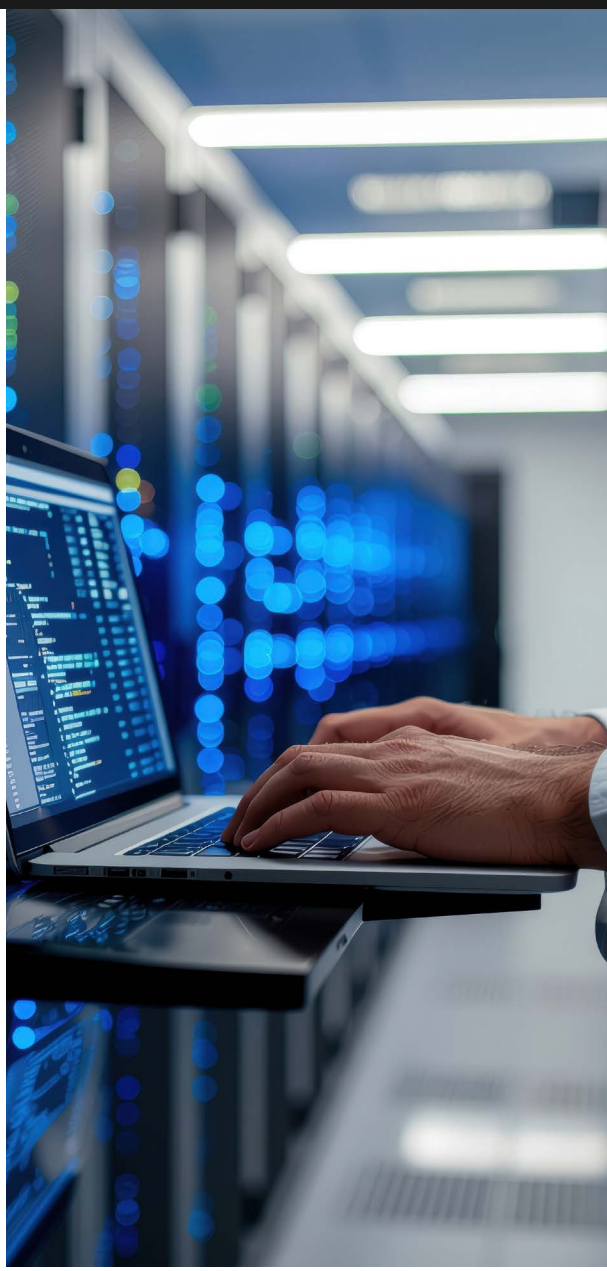

Legacy DR Challenges

Conventional Disaster recovery (DR) solutions were designed for closed, static environments and are no longer applicable to dynamic IT environments. Contemporary businesses are running on multi-cloud and hybrid environments that require flexibility, scalability, and adaptability; functions that legacy DR models are not adept at.

As a result, organizations are encountering increased difficulties in ensuring continuity of business processes across different platforms. The incomplete infrastructure, growing operational complexity, and lack of interoperability render the process of guaranteeing reliable recovery results in case of an outage or disruption quite challenging. Such restrictions are subjecting businesses to longer periods of downtime and escalated costs as well as operational risks.

The reasons why the continuity of business at various platforms has emerged as a major challenge to organizations in the modern era include:

01 **Cross-Cloud Limitations**
Legacy DR tools lack seamless data mobility and orchestration across multiple clouds, resulting in slow recovery times and operational inefficiencies.

02 **Complexity and Management**
Managing fragmented environments—on-premises, private cloud, and public clouds (AWS, Azure, GCP)—with disparate APIs and tools creates a highly complex DR landscape.

03 **Security and Compliance Risks**
Maintaining consistent security policies, encryption standards, and compliance (e.g., GDPR) across heterogeneous environments is difficult, increasing exposure to threats.

04 **Integration with Legacy Systems**
Older infrastructure often requires extensive rework to integrate with modern cloud platforms, delaying DR readiness.

05 **Data Gravity and Bandwidth Constraints**
Transferring large datasets between clouds or back on-premises strains network bandwidth, causing latency and slow recovery.

06 **Cost Management Challenges**
Unpredictable costs for data egress, storage, and multi-provider management make DR financially burdensome.

07 **Lack of Visibility and Control**
Achieving unified oversight across diverse environments is challenging, creating operational blind spots.

08 **Testing and Maintenance Gaps**
Regular, realistic DR testing is often skipped due to complexity, making recovery plans unreliable.

09 **Vendor Lock-In**
Dependence on proprietary tools limits multi-cloud portability and inflate costs.

10 **Skills Gap**
Finding IT staff skilled in both traditional systems and multiple cloud platforms is increasingly complex.

11 **Automation Deficiencies**
Traditional DR lacks automation for modern, cloud-native workloads, slowing recovery and increasing manual effort. [5][6][7]

Why Traditional DR Fails?

Traditional DR was designed to support static infrastructure and, as such, does not support the distributed environment of hybrid and multi-cloud services. This results in sluggish Recovery Time Objective (RTO)/ Recovery Point Objective (RPO), increased chances of data loss, and high management overheads, which is not suitable for contemporary business continuity requirements.

Overcoming DR Challenges with Datamotive

Cross-cloud disaster recovery entails copying workloads to another cloud computer system in case of an outage or disruption. This has become a significant requirement as organizations are seeking to de-risk their operations to enhance business continuity and guard brand reputation as well as customer trust. Due to the growing dependency of enterprises on heterogeneous cloud environments and geographically based limitations posed by providers, slippery workload portability between platforms has become the cornerstone of present-day DR strategies. Datamotive's EasyHybridDR solution plays a notable role in this landscape by facilitating the experience of predictable, secure, and operationally seamless cross-cloud migrations; addressing issues that cannot be solved in traditional and cloud-linked DR models.

Datamotive's DR solution provides cross-cloud DR via an architecture whose 10-minute predictable RTO per workload delivers high performance and certainty in terms of source cloud ability, workload size, and operating system. This uniformity forms the basis of providing recoveries that are consistent, reproducible, and geared towards continuity adherence levels within the enterprise. [1][2]

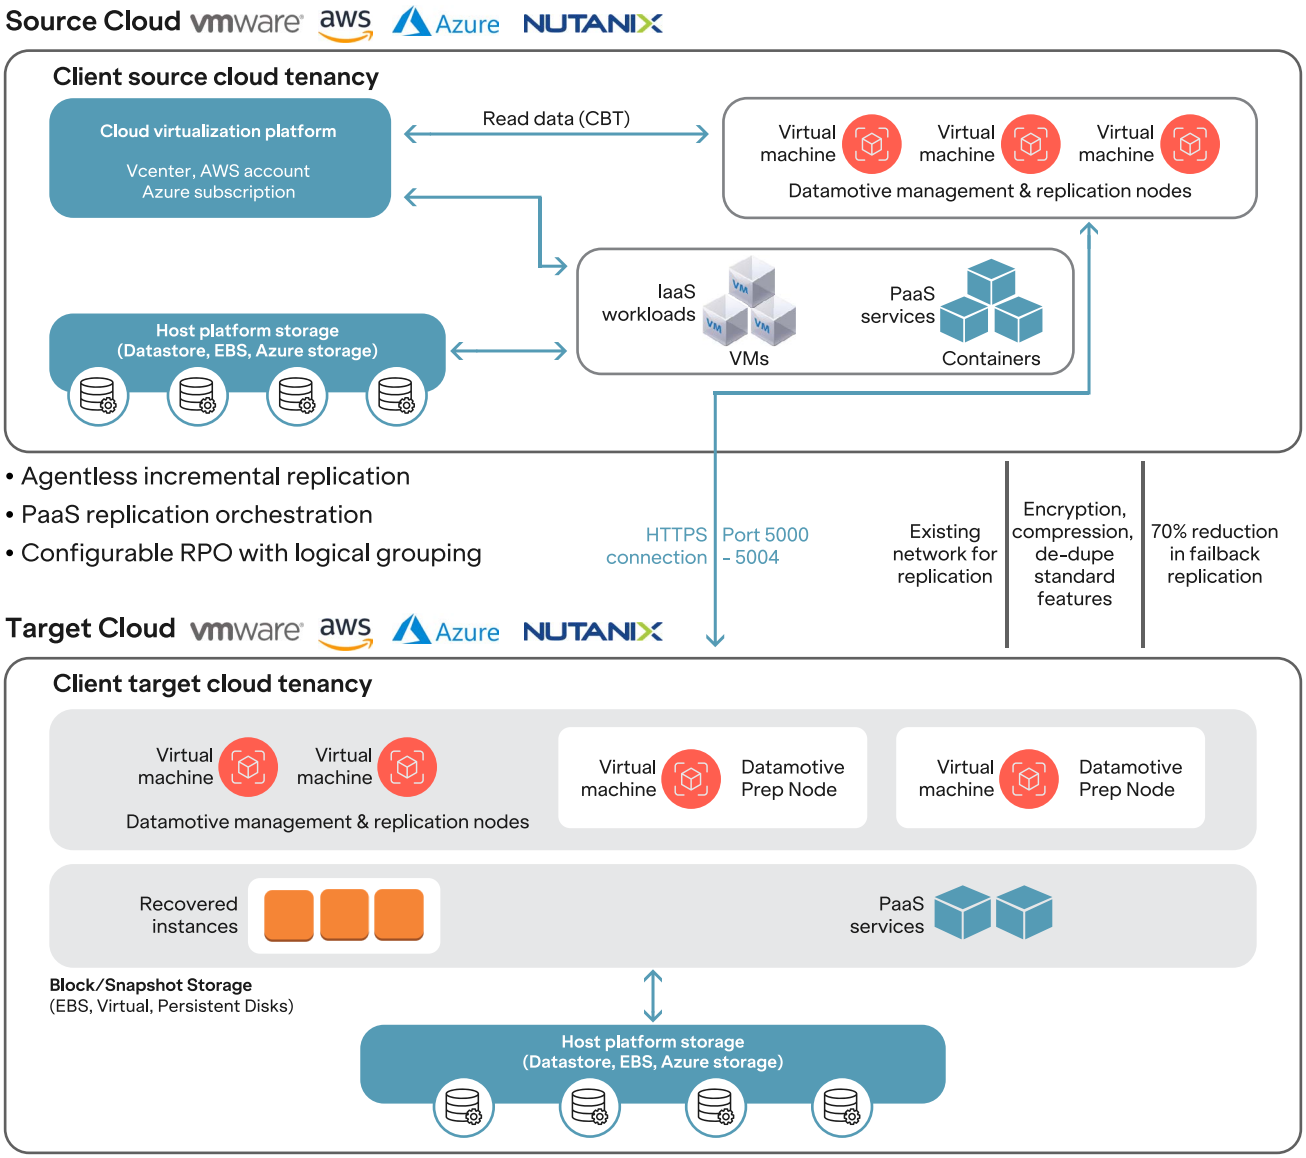

Figure 1: Datamotive Technical Architecture

The DR solution to be implemented by Datamotive is a hosted appliance installed in the customer's source and target clouds. It works with the customer using the existing compute, network, and security controls, making sure data replication is within established security guardrails. This strategy does not add risk to the enterprise security, governance, and data residency needs.

The agentless incremental replication engine incorporates encryption and compression as normal features, where the transfer of the information is safe and efficient to the network. This removes the operational overhead of the agent and provides it with a scalable low impact replication strategy that is appropriate in a wide range of application landscapes.

In the case of data retention, Datamotive makes use of the built-in block storage tools of the target cloud, which maintain native encryption and integrity provisions of the provider. This model will provide clients with full ownership of their data and will cut steady-state DR expenses by up to 60%. Long-term resiliency is strategy-wise and cost-efficient.

To ensure zero-recovery cost, Datamotive pre-creates all the workload compute components and parallels them before shutting them down reducing recovery cost. These elements can come online instantly during an event, which ensures 100% **recovery** ready without any performance degradation or performance provisioning lag.

The RTO of Datamotive, set at 10 minutes per workload (predictable), combined with parallel activity during failover, provides an unmatched level of clarity during DR operations. The integrated orchestration engine enables one-click recovery, reducing operator dependency and minimizing errors in stressful situations moments. Datamotive also boosts the level of failback preparedness, making it a pillar rather than an ancillary one. Using incremental failback replication and the ability to use the available assets of the production site, Datamotive saves up to 70% of the effort, egress cost, and time to restore the organization to a steadystate and in a way that is reliable and with the least amount of disruption.

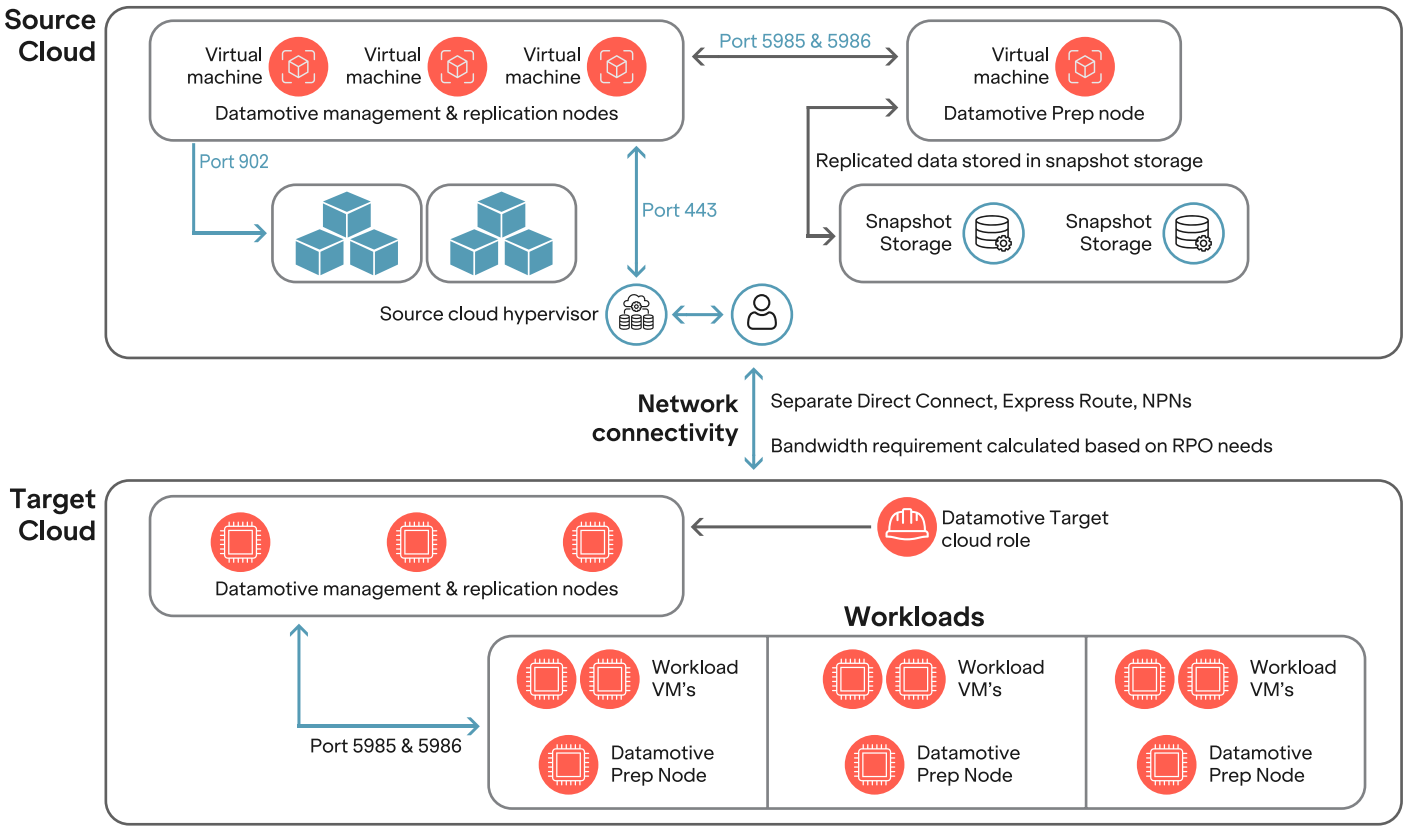

Figure 2: Product Components-wise Architecture Show

Table 1: Source and Target Cloud details

Source Cloud	Target Cloud
Supported platforms: VMWare, Nutanix, AWS, and Azure	**Supported platforms:** VMWare, Nutanix, AWS, and Azure
Deployment Model: Hosted	**Deployment Model:** Hosted
Appliance deployed: Datamotive management nodes, replication nodes, and prep-nodes (one per AZ) deployment in the source cloud tenancy.	**Appliance deployed:** Datamotive management nodes, replication nodes, and prep-nodes (one per AZ) deployment in the target cloud tenancy.
Connectivity with workload compute: Agentless via a service account, an Identity and Access Management (IAM) role, or an app registration.	**Connectivity with workload compute[TT1.1]:** Agentless via a service account, an Identity and Access Management (IAM) role, or an app registration.

Replication engine

The Datamotive replication engine exhibits the following characteristics and values to customers.

Table 2: Replication engine characteristics and their value for the customer

Characteristics	Target Cloud
Agentless replication between the source and target cloud	• Zero performance impact on production workload • Zero change to the workload security profile
Forever incremental	Less data to replicate, leading to reduced time and cost
Compression and encryption for the replication cycle	Compression reduces data size before transfer, lowering replication volume and cloud egress costs during steady-state operations. Encryption secures data in transit and at rest, reducing breach impact and ensuring compliance.
Differential reverse replication	• Reduces failback replication data transfer by 70% • Reduces failback data transfer cost by up to 60%

Datamotive's agentless, incremental, differential reverse replication ensures customers' data is replicated securely, leveraging the most cost-effective mechanism, without impacting production workloads' performance or security profile.

Orchestration Layer

The Datamotive platform includes an in-built orchestration layer designed to deliver a seamless, singleclick disaster recovery experience. This orchestration layer is enabled through a combination of Python, Go, and Shell scripts, which can incorporate cloudspecific CLIs, API integrations with network services, certificate generation, and interapplication connectivity workflows. To support flexible and controlled recovery operations, these orchestration scripts can be defined and applied at four distinct levels during both replication and recovery, enabling granular customization aligned with application and infrastructure requirements.

Datamotive's DR solution has a flexible, multilevel orchestration framework for both replication and recovery workflows, enabling precise control over actions at the instance and group levels.

During replication, orchestration can be executed at four levels: prereplication and postreplication actions for every individual instance, as well as prereplication and postreplication actions applied at the group level. This allows customers to perform preparatory checks, environment configuration, dependency validation, or postreplication tasks either per instance or collectively across an application group.

Similarly, during recovery, Datamotive supports four orchestration levels: prerecovery and postrecovery actions for every instance, along with prerecovery and postrecovery actions at the group level. These levels enable fine-grained control over recovery sequencing, including infrastructure readiness, network and security configuration, application startup ordering, validation, and post-recovery verification.

This tiered orchestration model ensures consistent, repeatable, and automated disaster recovery operations while allowing customization based on application complexity and business requirements.

Examples of such scripts include:

Replication

Level 1: Cloud-specific PaaS service replication before the instance replication cycle.

Recovery

Level 1: Updates to the cloud-specific cloud balancer for each instance

Level 2: Dollar test for testing the entire service

Execution

Datamotive automates the execution of these scripts, as well as the replication or recovery workflows. For this, Datamotive may require the necessary username and password to log in to the instance and execute these scripts. The credentials can be uploaded to Datamotive as an encrypted file. These credentials are held by Datamotive only during execution

Network Considerations

Datamotive is designed to support the existing network profile established between source and target Data Centers (DCs). The only requirement is that these source and target Data Centers can communicate on the ports defined in the deploy-ment guide. The network size is defined based on the customer's Disaster Recovery requirements. The Recovery Point Objective (RPO), the expected change rate, and the number of workloads provide an estimate of the network bandwidth requirements.

Datamotive's architecture is designed to use the existing network connectivity between target and source datacentres, with below defined components and bandwidth sizing aligned to DR objectives -

- Source and Target DCs
- Data replication engine
- Orchestration layer
- Network considerations (10G uplinks, SFPs)

Integration with Nutanix, VMware

Most enterprises already run mission critical workloads on Nutanix AHV or VMware, with established processes for provisioning, networking, security, and operations. The Datamotive DR solution enables orchestration of replication and recovery by using existing APIs, constructs, and management workflows in the manner that it is built and executed in modern application development and operation.

From a business perspective, it is an integration that reduces recovery time and operational risk by removing custom tooling and manual intervention during a crisis. It enables predictable, repeatable recovery outcomes, maintains dependencies among applications, and aligns DR operations with platform governance and compliance criteria. As a result, organizations can achieve faster RTOs, lower operational overhead, and greater confidence that recovery actions will behave as expected across hybrid and multicloud environments

Datamotive is an agentless replication engine that integrates with the on-prem hypervisors Nutanix and VMware via Prism Central and vCenter respectively.

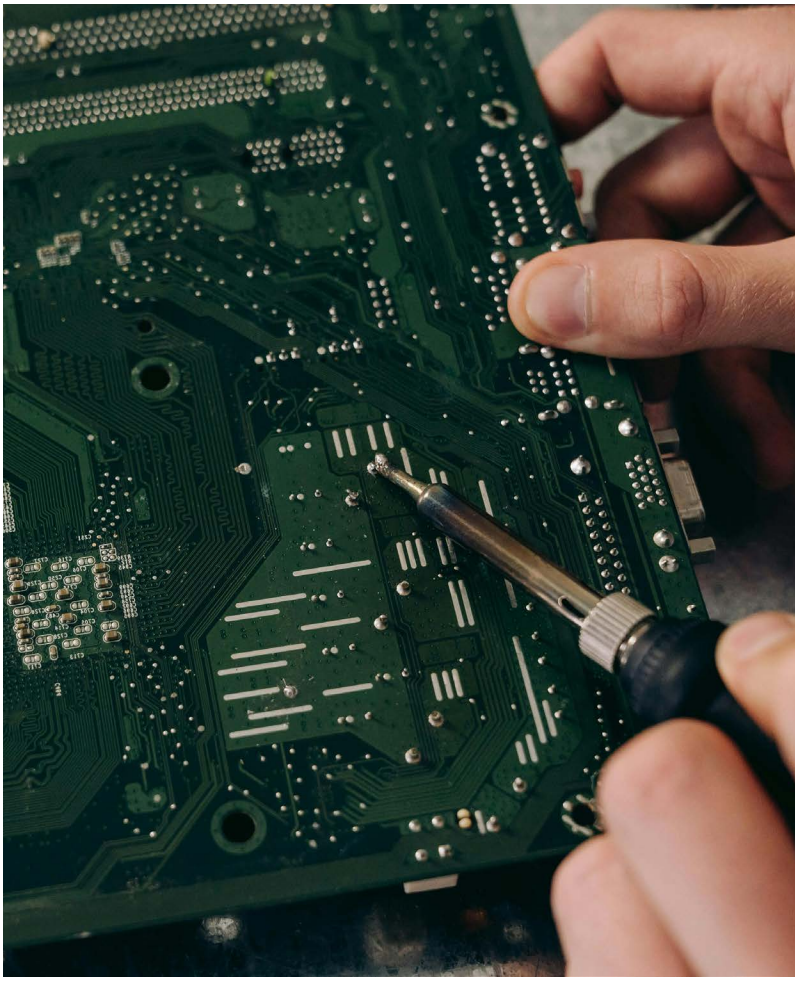

VMWare Integration with Datamotive

The figure below demonstrates the integration with these hypervisors.

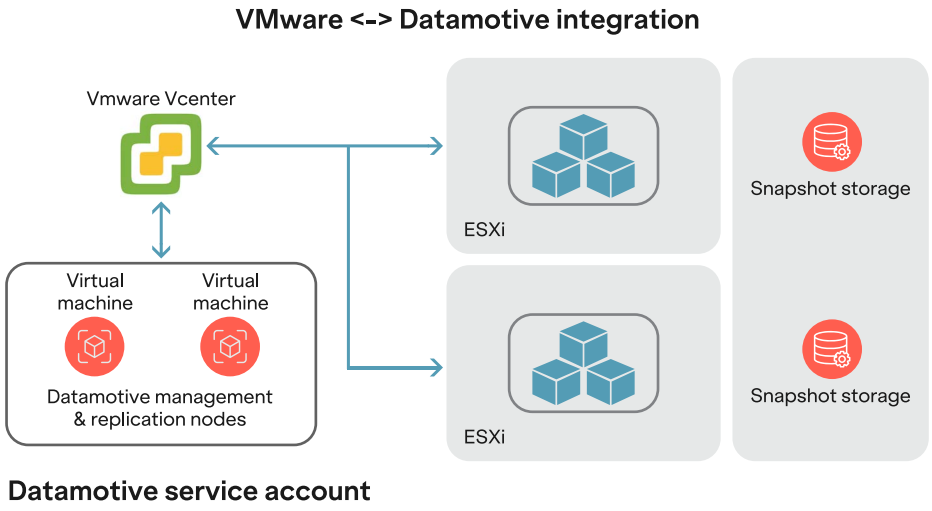

Figure 3: VMware integration with Datamotive

Datamotive requires a service account within the VCenter environment. The service account consists of access related to:

Request, create, delete, and manage snapshots.

Ability to create, update, delete virtual instances.

Read and write from Datastores.

None of these accesses involves deleting any datastores or data from the customer's environment. Datamotive does not install any agents on any workload, ensuring zero performance or security impact.

Nutanix integration with Datamotive

The figure below demonstrates that Nutanix integrates with Datamotive.

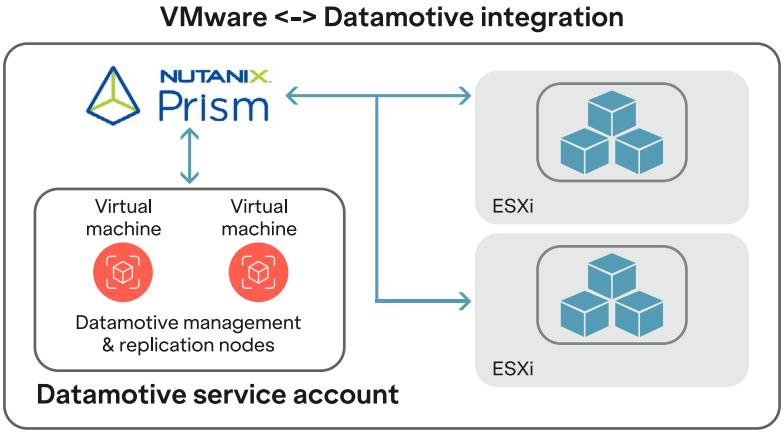

Figure 4: Nutanix integration with Datamotive

Datamotive requires a service account role to integrate with Nutanix Prism Central for access to the CBT APIs. The service account consists of access related to:

Request, create, delete, and manage snapshots.

Ability to create, update, delete virtual instances.

Read and write from Datastores.

None of these accesses involves deleting any datastores or data from the customer's environment. Datamotive does not install any agents on work-loads, ensuring zero performance or security impact.

Key Features

Designed for enterprise resilience and operational simplicity, Datamotive DR solution has the following key features:

- Cross-cloud failover and failback for seamless disaster recovery.
- Automated workload migration to reduce manual intervention.
- SLA-driven recovery objectives ensuring predictable RPO/RTO.
- Built-in security and compliance for enterprise grade protection.

Key Benefits

The key technical and business benefits delivered by **Datamotive** in a cross-cloud DR scenario:

Category	Benefits
Technical	• Secure, agentless replication with encryption and compression, ensuring compliance and data integrity across clouds. • Predictable recovery with consistent 10-minute RTO per workload and single-click orchestration for simplified failover. • Cost-efficient architecture leveraging native cloud block storage, reducing DR costs by up to 60%. • Incremental failback replication minimizes egress costs and effort by up to 70%. • Seamless support for diverse workloads and multi-cloud environments without vendor lock-in. • Automated orchestration and parallel recovery capabilities for faster, more reliable DR operations.
Business	• Enhanced business continuity with minimal downtime and operational disruption. • Significant reduction in total cost of ownership by eliminating physical DR. • Improved compliance posture and security governance across hybrid and multi-cloud setups. • Greater agility and scalability to adapt to evolving business and technology needs. • Proven outcomes with measurable savings and operational efficiency for global enterprises.
Operational and Strategic Benefits	• Simplifies workload migration across environments with minimal complexity. • If it is supported on a specific major hypervisor and cloud region, it is 100% flexible and portable. • Data security is also inherent to the design, and it is built into the process with encryption and compliance. • Ensures congruence with organizational migration strategies, with guaranteed results. • Allows migration and recovery with less than zero downtime, maintaining continuity. • Can be used to enable efficient failover and failback, ensuring quick recovery. • Provides centralized, secure management across environments to provide complete visibility and control.

So, the Datamotive DR solution provides a secure, high-performance replication engine with predictable 10-minute recovery times per workload, agentless replication, and seamless failover/failback processes, enhancing automation, orchestration, and resilience for cross-cloud DR.

Key Outcomes

The Datamotive integration supports diverse workloads, including Windows, Linux, IIS, MS SQL Server, and Oracle, ensuring broad application compatibility across both cloud and on-premises environments. Disaster recovery drills can be conducted without impacting source workloads or interrupting active replication, enabled by isolated network-based recovery testing. Incremental bidirectional replication significantly reduces data egress costs by up to 70%, while eliminating the need for a dedicated recovery site with the same platform as the source, resulting in an overall total cost of operations (TCO) reduction of approximately 65%. Together, these capabilities deliver cost-effective, non-disruptive, resilient disaster recovery operations across heterogeneous infrastructure environments. [9][10][5]

How the Datamotive DR Platform Leads the DR Landscape

The Datamotive DR platform will transform modern disaster recovery to unprecedented agility, speed, and workload portability unmatched in traditional and conventional cloud-based services. Built for dynamic, multi-cloud, and hybrid environments, it removes legacy limitations and delivers faster recovery, simplified operations, and consistent results. The strengths discussed below demonstrate how Datamotive is a leader in the disaster recovery environment:

01 **AnytoAny Recovery**
Seamlessly moves workloads across diverse environments without conversions or rearchitecture.

02 **Fast, Predictable RTO**
Consistently achieves near instant recovery through a high performance, uniform replication engine.

03 **Conversion Free Recovery**
Maintains workloads in a neutral, cloud agnostic format for instant boot ready failover.

04 **Agentless Architecture**
Eliminates agents and kernel modules, reducing overhead and simplifying operations.

05 **No LockIn**
Enables recovery to any environment, enhancing flexibility and resilience.

06 **Low TCO**
Uses snapshot based, incremental replication with no active DR compute, minimizing DR cost.

07 **DeltaBased Failover/Failback**
Transfers only incremental changes for faster sync, lower bandwidth use, and smoother returns.

08 **MultiCloud Ready**
Supports hybrid and distributed workloads with a unified orchestration approach.

09 **Strong Security and Compliance**
Keeps data within the customer environment while preserving enterprise grade governance and encryption.

Risks and Mitigation

The Datamotive DR solution is subject to risks like temporary non-production downtime during migration, lack of an existing disaster recovery framework, as well as issues of proper workload sizing, which could influence planning and optimization. The risks are efficiently mitigated through Datamotive's incremental replication, which reduces downtime to a minimum and maintains data consistency. Isolated recovery testing allows one to verify failover and failback operations without disrupting production loads, thereby providing operational readiness. Moreover, the flexible and scalable architecture of Datamotive is designed to meet the needs of various clients and allows them to scale up or down accurately, optimize workload sizing, and support disaster recovery planning.

Use Cases

Datamotive DR solution "EasyHybrid DR" is designed to accommodate a variety of disaster recovery and migration scenarios for hybrid and multi-cloud environments. Its secure, highly scalable, agentless architecture allows organizations to safeguard their vital workloads, verify their recovery preparedness, and optimize cost savings and business continuity. The table below uses the following use cases to demonstrate how Datamotive can deliver concrete, quantifiable value across a wide spectrum of operational requirements.

User Case 1

A global leader in advanced mobility solutions, with eight VMwarebased manufacturing sites, needed to establish a Disaster Recovery solution on AWS public cloud.

Challenges

- Multiple manufacturing locations across the US and Canada.
- Looking for cross-cloud DR from VMWare to AWS with an RTO of 30-minutes and one hour across 200 servers for both failover and failback.
- Existing solutions like Zerto did not meet the RTO requirements. Solutions like Zerto lite proved to be costly from AWS cost perspective.
- Had plans to move from VMWare to Nutanix.

Solution

Details of the solution provided to the customer, the justification for the recommendation, and the implementation approach.

- Datamotive DR solution "EasyHybrid DR" was the solution proposed.
- Selection was based on the following criteria:
    1. Agentless, always incremental replication.
    2. Meeting and improving the RTO requirements of 30-minutes and one hour with a 10-minute RTO irrespective of size.
    3. Commitment to provide Nutanix platform support with the same user experience and the same RTO.
    4. Ability to conduct failover and fallback incrementally and with the same 10-minute RTO in both directions.

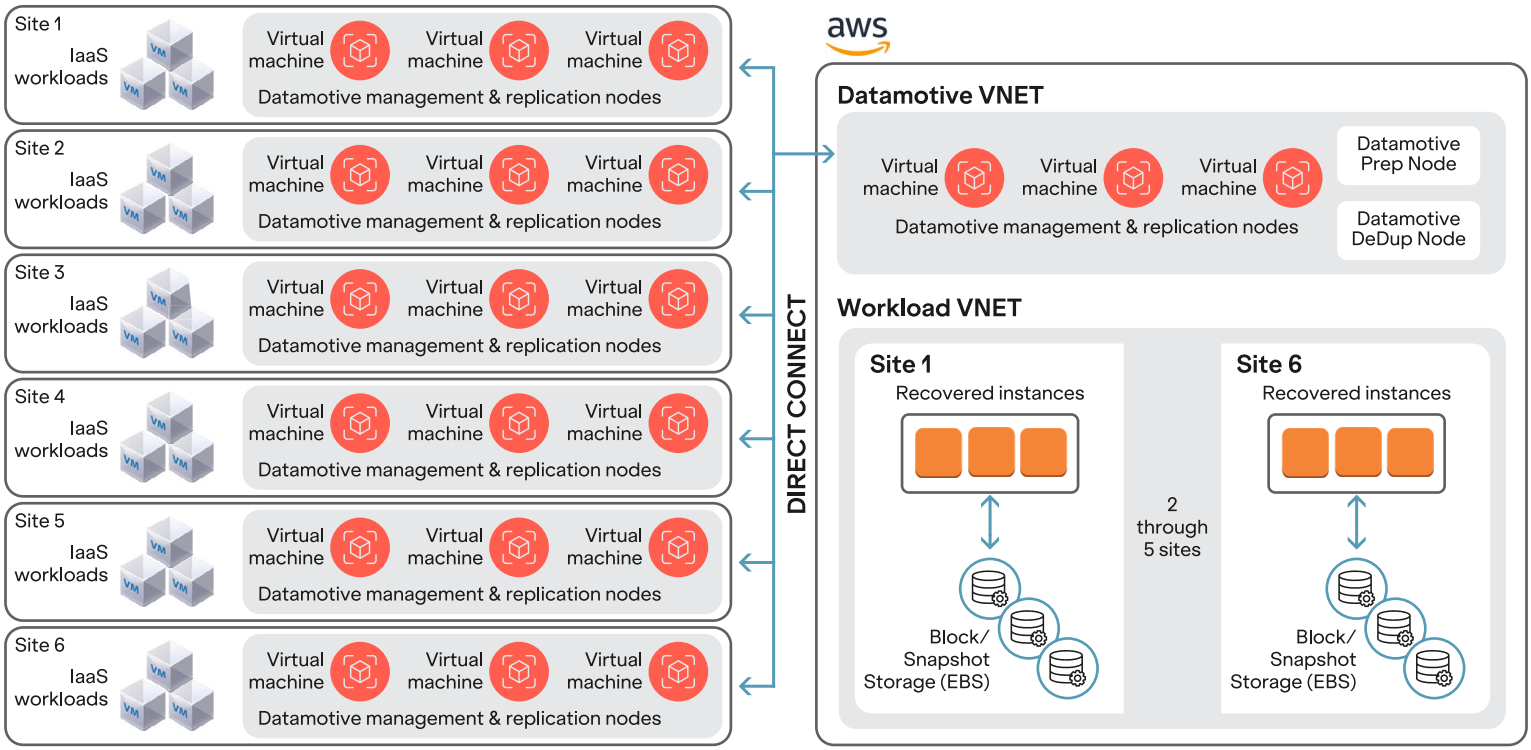

Figure 5: DR Solution for 6 VMware-based Manufacturing Sites

Benefits

- 6 manufacturing locations replicating data into AWS, leveraging the same user experience (minimal learning curve)
- With an incremental fallback, savings of up to 70% for fallback operations
- A single Day-2 portal to manage and operate all 8 distinct locations.
- Reduced the overall DR TCO by 60% without compromising the RTO requirements.

User Case 2

A U.S. East Coast–based international cruise liner leveraging VMware on-prem and AWS, Azure, and Google Cloud needed a unified Disaster Recovery solution.

Challenges

- A multi-cloud setup that included VMWare, AWS, and Azure.
- Currently leveraging different clouds, different SOPs, and a different set of tools based on the cloud.
- Very high cost of ownership, inconsistent RTOs, and a large amount of scripting.

Solution

Overview of the solution offered to the customer, why it was recommended, and the implementation approach followed:

- Datamotive DR solution “EasyHybrid DR” was the solution proposed.
- Selection was based on the following criteria:
    1. Reduced operational costs (people and time) by 45%.
    2. Freedom to choose the cloud based on workload criticality providing more negotiating leverage.
    3. Hosted solution securing data is within their network, meeting compliance and security requirements.

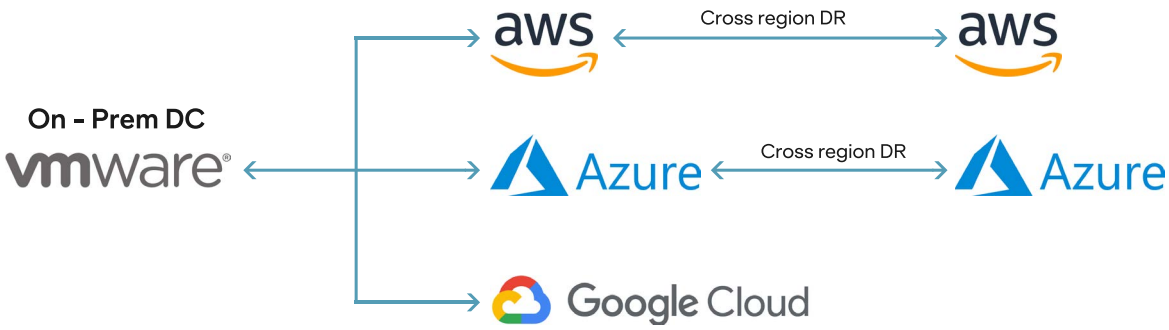

Figure 6: DR Solution for on-prem and multiple cloud sites

Benefits

- Single tool across all clouds leading to a single SOP, and consistent RTOs.
- A full capacity DR at 45% lower cost with workload compute reserved but not running.
- Reduced SRE team effort by 60% with a single pane of operations across all clouds.

Industry Aligned Methodology and Strategy for Datamotive

Datamotive and LTM takes a strategic, incremental approach, focusing on modernization, cloud use (Snowflake, Azure, AWS), and AI/ML automation. The use of proprietary service options, e.g., Mosaic, Design Thinking, and a clean core strategy, provides a safe and smooth transition of the old systems to the new, agile, and scalable cloud infrastructure. This approach reduces technical debt, improves governance, and opens up to improved analytics and business value.

Table 3: LTM Datamotive Methodology

Phase	Description
CloudFirst - Transformation	Migrate from on-prem systems to scalable, cost -efficient cloud data platforms using modern cloud services and accelerators. E nsure readiness with high performance analytics and future transformation initiatives.
Design Thinking and Discovery	Conduct collaborative discovery workshops to learn about business problems, pain points, and objectives. Verify and identify the critical data domains, their attributes, use cases, and personas with Design Thinking approaches.
Phased Migration Roadmap	Introduce a systematic migration strategy comprising of Data Discovery, Data Build and Data Validation. Adhere to established procedures to ensure accuracy, minimal disruption, and a smooth transition.
Automation and AI/ ML Enablement	Leverage automation, AI capabilities, and intelligent tooling to accelerate data access, improve decision insights, and support the creation of advanced analytics models.
Clean Core and Modernization	Maintain a clean, modular core system by reducing customizations. Modernize through cloudbased extensions and micro-services to minimize technical debt and enhance agility.
Data Governance and Security	Establish strong governance frameworks and enforce advanced security controls. Ensure data quality, compliance, lineage tracking, and protection throughout the migration lifecycle.
Proprietary Platforms and Accelerators	Utilize mature internal accelerators and IPs to fasttrack cloud adoption, analytics transformation, and modernization. Enable predictable, repeatable, and agile delivery outcomes.

LTM Datamotive Strategy

The LTM Datamotive DR solution follows a structured, step-by-step approach to cloud adoption, modernization, and resilience by focusing on gradual change and operational predictability. The strategy will allow enterprises to implement secure and consistent continuity in both hybrid and multi clouds by integrating the cross-cloud disaster recovery offerings of Datamotive, design-led methodology, and automation-driven frameworks of LTM. The focus of this approach is lower technical debt, better governance and increasingly quantifiable business value, and agnostic to cloud disaster recovery operation, which is scalable, efficient and aligned to long term transformation goals.

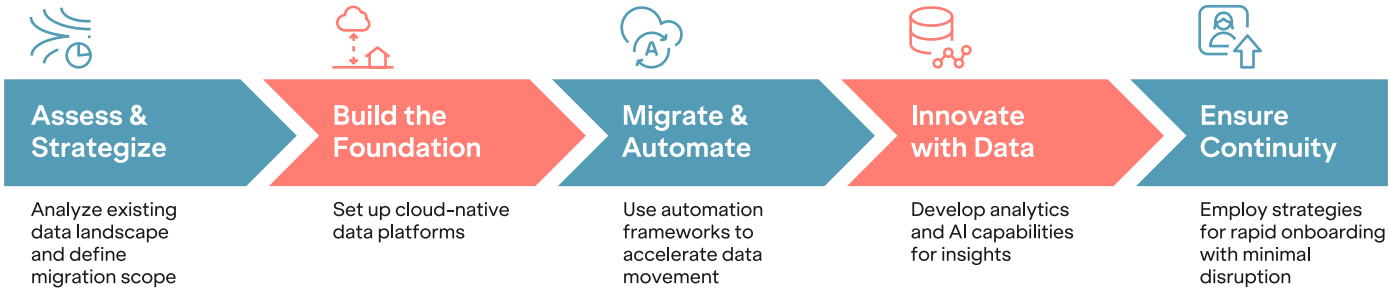

Figure 8: LTM Datamotive Strategy

Conclusion

The concept of ensuring business continuity amid exceptionally sensitive and disruptive incidents will go beyond the conventional disaster recovery frameworks, especially in current hybrid and multi-cloud structures. As infrastructure landscapes become more distributed, operational complexity, rising costs, and declining efficiency become harder to manage, and resilience and agility harder to achieve. Datamotive's DR solution overcomes these issues by providing an efficient and secure as well as a predictable cross-cloud disaster recovery solution at a cost-effective price. Its agentless design, autopilot orchestration, and a 10-minute recovery time per workload give companies the understanding and confidence in the event of an outage. The solution minimizes the overhead of recovery, platform-independent failback effort, and disaster recovery costs with small cost and high security and compliance to environments by using native cloud services and incremental replication. Datamotive's proven DR solution has been demonstrated with actual deployments; it empowers organizations to transform and modernize their legacy physical DR systems, unify recovery activities, and realize significant cost savings. More to the point, it enhances business resilience and scalability, enabling organizations to adopt cloud-native disaster recovery plans that best fit the evolving digital environment.

References

[1] Datamotive, Cross Cloud Disaster Recovery Overview, Datamotive, 2024, https://www.datamotive.io

[2] Datamotive, Agentless Replication and Cross Cloud DR Architecture, Datamotive Resources, 2024, https://www.datamotive.io/resources

[3] National Institute of Standards and Technology (NIST), Contingency Planning Guide for Federal Information Systems (Special Publication 800 34 Rev. 1), NIST, 2010, https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final

[4] ISO/IEC, ISO/IEC 22301:2019 – Security and Resilience — Business Continuity Management Systems, International Organization for Standardization, 2019, https://www.iso.org/standard/75106.html

[5] Gartner, Best Practices for Disaster Recovery in Hybrid and Multi-Cloud Environments, Gartner Research, 2023

[6] Amazon Web Services, AWS WellArchitected Framework – Reliability Pillar, Amazon Web Services, 2024, https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/

[7] Microsoft, Azure WellArchitected Framework – Reliability and Business Continuity, Microsoft Learn, 2024, https://learn.microsoft.com/azure/architecture/framework/resiliency/

[8] Google Cloud, Google Cloud Architecture Framework – Reliability, Google, 2024, https://cloud.google.com/architecture/framework/reliability

[9] VMware, Disaster Recovery Best Practices for Virtualized and Hybrid Environments, VMware Technical Whitepapers, 2023, https://www.vmware.com/resources.html

[10] Nutanix, Business Continuity and Disaster Recovery Reference Architecture, Nutanix Documentation, 2023, https://www.nutanix.com/products/business-continuity

Authors

Nitin Kumar

Senior Principal – Architecture, Cloud & Infrastructure, LTM

Nitin is a senior cloud and infrastructure architect with deep expertise in hybrid and multicloud architectures, disaster recovery, and largescale infrastructure modernization. He has led multiple enterprise grade DR and cloud transformation initiatives, helping global organizations achieve resilience, cost optimization, and operational excellence.

Tushar Tarkas

Director of Technology, Datamotive

Tushar is a seasoned technology leader with over two decades of experience in building and scaling enterprise-grade products across cloud, datacenter, and infrastructure domains. With deep expertise in product engineering, virtualization, and hybrid cloud environments, he has led large-scale transformation initiatives focused on resilience, performance, and scalability. Tushar has a proven track record of driving end-to-end product development, aligning engineering with business strategy, and enabling organizations to deliver high-impact, market-ready solutions.