DATAMOTIVE
Eliminating Cloud Boundaries

Product Deployment Guide

Hybrid, Multi-Cloud Workload Portability Platform

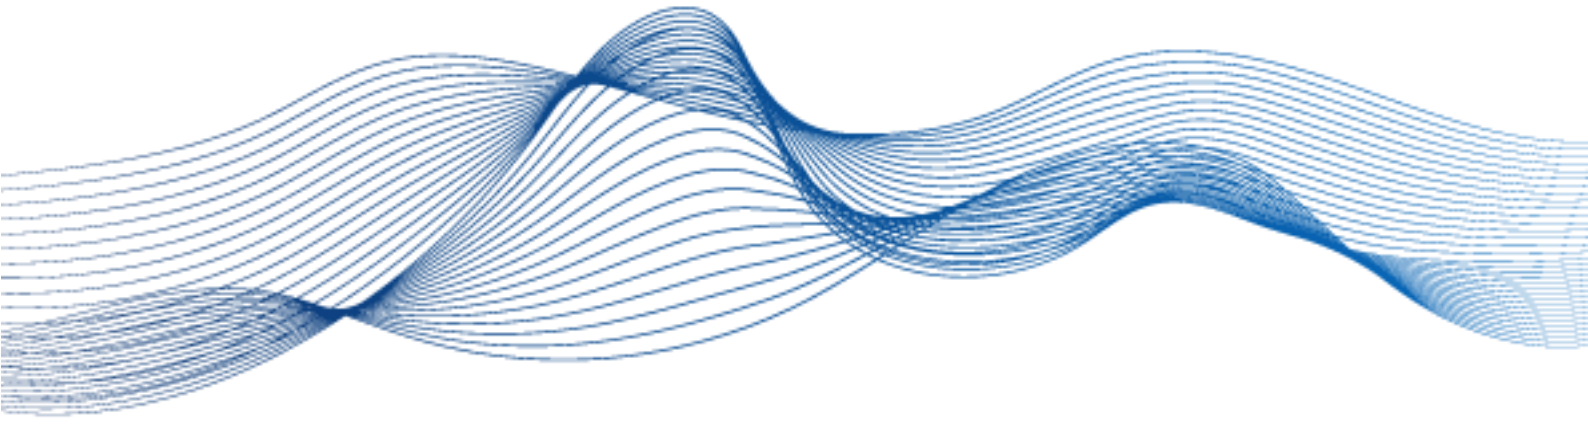

Datamotive Product Deployment Guide

The Datamotive Product Deployment Guide provides:

- a general overview of Datamotive solution and its different components
- information about how to install and configure Datamotive solution

Intended Audience

This guide is intended for anyone who wants to install or configure Datamotive solution. The guide is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacentre operations.

Datamotive Solution Versions

- 1.1.1
- 1.1.2

Contents

1. Introduction

Datamotive is industry's first product enabling organizations to move workloads across public and private clouds with a guaranteed 10-minute recovery SLA irrespective of source cloud, platform, and size. With its technology, Datamotive aims to redefine the hybrid cloud operating model by eliminating cloud boundaries that exist due to hypervisor dependencies.

Datamotive can be used to protect virtual machines/instances on primary site by replicating them periodically to the recovery site. The protected virtual machines/instances can then be recovered as needed in the recovery site as native instances (e.g., protected virtual machine from other public or private cloud is recovered as a native AWS EC2 instance).

2. Datamotive Solution Deployment Process

This section will discuss the deployment process of Datamotive solution.

2.1. Overview

Datamotive solution supports various cloud and virtualization platforms such as protected and recovery sites. Deployment of Datamotive solution is similar on all supported platforms. There are pre-requisites for deploying and configuring the Datamotive solution. Configuration of these pre-requisites vary depending on the platforms.

2.2. 3-Step Deployment Process

The deployment of Datamotive solution is a 3-step process as shown in **Figure 1**.

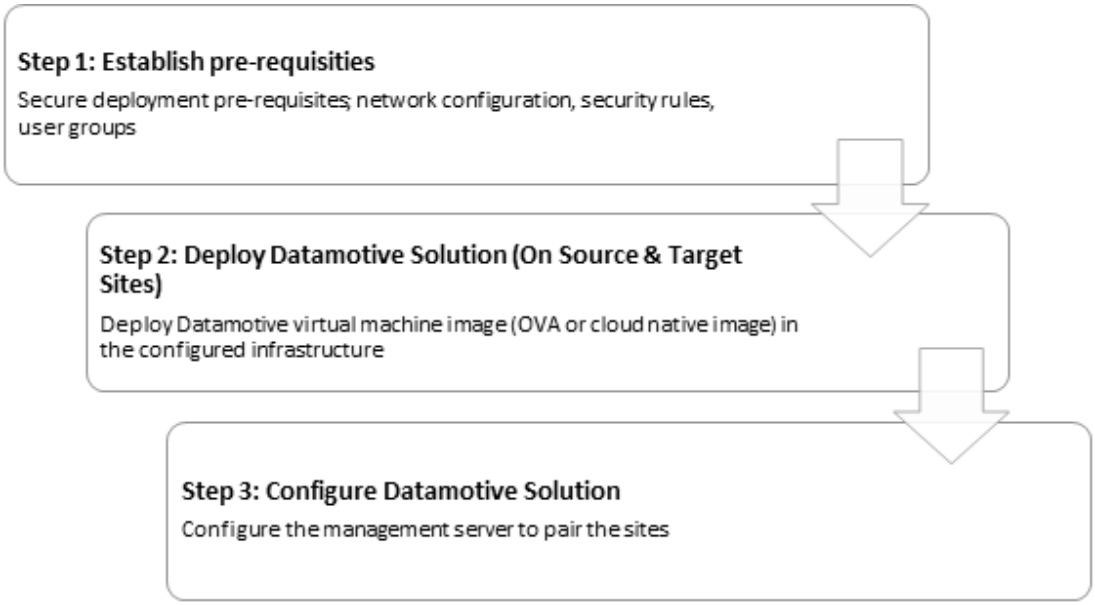

Figure 1: 3-step deployment process

3. Datamotive Solution Deployment Architecture

Datamotive is a software only solution with no dependency on additional hardware or software components. The architecture is a built adopting best practice around network security, data protection and scale.

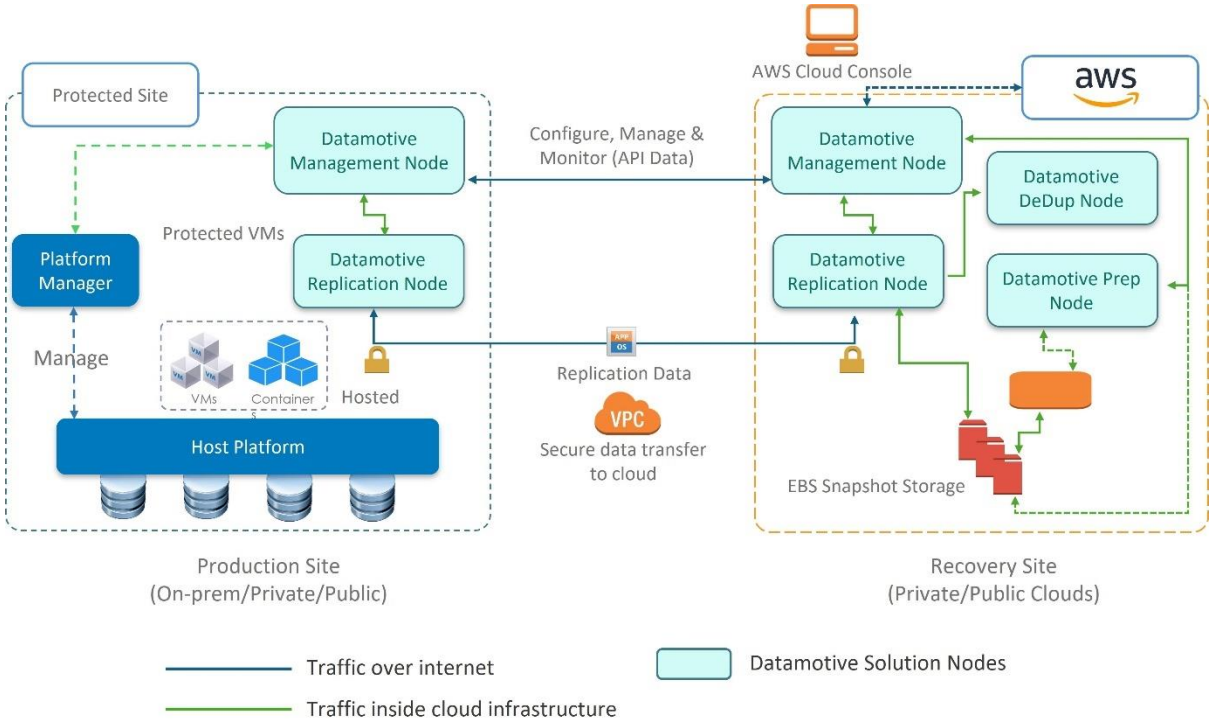

Figure 2: Datamotive Solution Deployment Architecture

3.1. Deployment Architecture: Key Terms

Below is the list of deployment architecture's key terms with their definitions:

1. **Protected site**: This is the source cloud environment hosting workloads and associated instances
2. **Protected virtual machines**: These are virtual instances within the source cloud environment marked for migration and / or disaster recovery
3. **Recovery site**: This is the target site where the workloads need to be migrated and / or recovered
4. **Recovered instances**: These are the recovered instances in the target site as part of the disaster recovery process.

3.2. Datamotive Solution Components

The Datamotive solution components include:

1. **Datamotive Management Node** is a virtual appliance deployed in a pair configuration—one on the protected site and other on the recovery site. The Datamotive management server hosts the Datamotive intuitive user interface (UI), CLIs and REST APIs for the IT administrators to perform Day0–DayN activities. For small scale deployments (up to 40 virtual machines disks), the management server also acts as a replication node. It is shipped as an OVA for VMware environment and cloud native machine image for AWS, GCP, and Azure environments.
2. **Datamotive Replication Node** is a virtual appliance for large scale deployment and is deployed in protected and recovery sites to perform data replication operations. This node executes the data replication jobs and can be leveraged to increase the overall replication capacity of the solution based on the number of protected virtual machines/instances. The maximum number of parallel replication jobs (one replication job per protected virtual machine disk/volume) supported by each node is defined by the limit provided by cloud platforms. For optimum performance, Datamotive solution limits the capacity of each replication node to 40 parallel disks. It is shipped as an OVA for VMware environment and cloud native machine image for AWS, GCP or Azure environment. Datamotive solution scales horizontally using the replication nodes.
3. **Datamotive Prep Node** is a Windows virtual appliance deployed in the recovery site infrastructure. This appliance is required only during recovery operation of Windows-based virtual machines. Once deployed and registered with Datamotive management node, it is powered-off so that it does not incur any additional compute cost while running in public clouds. It is automatically powered-on and used only when Windows virtual machines are recovered or migrated. Recoveries of Linux-based virtual machines are handled by the Datamotive management and

Datamotive replication nodes. It is shipped as an OVA for VMware environment and cloud native machine image for AWS, GCP and Azure environments.

4. **Datamotive DeDupe Node** is a virtual appliance deployed in the recovery site (public cloud) infrastructure where the protected virtual machines will be recovered in the event of disaster or when virtual machines are migrated. It is shipped as a native cloud image for AWS, GCP and Azure environments.

4. Deployment Pre-requisites

A set of pre-requisites to ensure a smooth deployment and configuration of Datamotive solution are followed for deployment in terms of:

- Deployment Location - Datamotive node deployment location
- Network - Network configuration and firewall rules
- Access Control - User roles and privileges

4.1. Deployment Location - Datamotive Node Deployment Location

The pre-requisites in terms of deployment location are as below:

1. Datamotive management server & replication nodes support replication source or target within a Cloud region or vCenter Server in case of VMware.
2. When the source infrastructure is public cloud, Datamotive nodes must be deployed in the same region as that of instances to be protected.
3. When the target infrastructure is cloud, Datamotive nodes must be deployed in the region where the protected instances are expected to be recovered. If recoveries are to be made across multiple availability zones within the target region, then Datamotive Replication/Prep Nodes must be deployed in each availability zone.

4.2. Networking - Network Configuration and Firewall Rules

The pre-requisites in terms of compute are as below:

1. **Bandwidth requirement**: Data replication between sites is dependent on the amount of data to be replicated between the two sites. Datamotive nodes are designed to support various network bandwidths. Datamotive provides a bandwidth throttling feature to match dynamic bandwidth usage needs of organizations. Datamotive recommends having a minimum dedicated bandwidth of 50 MBPS between the replication nodes to support replications for maximum supported disks per Datamotive node. An improper allocation of bandwidth will impact the time taken to replicate the data which in-turn may affect the configured RPO for the protection plan. The changed data, however, will eventually sync for all the protected applications.
2. **Datamotive inter-node connectivity**: Connectivity between Datamotive nodes within and across sites is used for transferring metadata and protected virtual

machine data. All these data transfers are secured. This connectivity can happen within private subnets or secured VPN tunnels depending on the infrastructure's network & security configuration. **Table 1** describes the ports required for communication across Datamotive nodes.

3. **Outbound connectivity**: For orchestrating the DR/Migration workflows, Datamotive nodes need to communicate with respective platform managers. (e.g. vCenter Server for VMware, AWS console APIs for AWS etc.). The operations are performed over secured APIs provided by the platform managers like vCenter APIs, AWS APIs etc. and include actions like Create virtual machine, Fetch Security Groups, etc. Datamotive nodes require outbound connectivity to reach to these platform managers. In case of public cloud platforms, the platform managers are accessed over internet and no organization data is transmitted over this connection. In such cases, Datamotive nodes need outbound internet connectivity.
4. **Connectivity with recovered entities**: Datamotive allows admin to provide custom scripts that can be executed in recovery workflows. Typically, these are pre or post recovery scripts trying to reconfigure recovered instances. In such cases, Datamotive will need network connectivity with the recovered instances so that the scripts can connect to them and perform required actions.

4.2.1 Network - Firewall Rules

Datamotive nodes operate on specific ports for their communication needs. The table below provides details of the ports. It is recommended to create security groups based on these definitions.

Table 1: Ports required for communication across Datamotive nodes

Node	Port	Description	Access
Security Group 1: Management Server	3308	Database connection	IN
	5000	Datamotive Management Server controller (GUI/REST) port	IN/OUT
	5001	Datamotive Replication Server/Node data transfer port	IN/OUT
	5002	Datamotive Replication Server encrypted data transfer port	IN/OUT
	5003	Datamotive Replication Node controller (GUI/REST) port	IN/OUT
	5005	Datamotive DeDupe server connection	OUT
	443	AWS API port	OUT
	-	vCenter Server API Port (In case protection site is VMware)	OUT
	5985	HTTP Connection to Windows Prep Node	OUT
	5986	HTTPS Connection to Windows Prep Node	OUT
	22	Used for ssh connection during upgrade process	IN/OUT
	902	VMware VDDK connection on the EXSI Host	IN/OUT

Security Group 2: Replication Node	5000	Datamotive Replication Server/Node Data transfer port	IN/OUT
	5001	Datamotive Replication Server/Node Data transfer port	IN/OUT
	5002	Datamotive Replication Server encrypted data transfer port	IN/OUT
	5003	Datamotive Replication Node controller (GUI/REST) port	IN/OUT
	5005	Datamotive DeDupe server connection	OUT
	443	AWS API port	OUT
	3308	Database connection	OUT
	-	vCenter Server API Port (In case protection site is VMware)	OUT
	5985	Connection to Windows Prep Node	OUT
	5986	HTTPS Connection to Windows Prep Node	OUT
	22	Used for ssh connection during upgrade process	IN
	902	VMware VDDK connection on the EXSI Host	IN/OUT
Security Group 3: DeDupe Node	5005	Datamotive Dedupe Node controller (REST)	IN
Security Group 4: Windows Prep Node	5985-5986	WinRM communication port	IN
	3389	RDP	IN

4.3. Access Control - User Roles and Privileges

Datamotive nodes integrate with platform managers (vCenter Server, AWS Console, GCP Console etc.) for orchestrating replication, DR and Migration workflows. The nodes leverage platform manager's APIs. The below table explains the roles required to be created based on the source and recovery sites.

Table 2: Roles required to be created based on the source and recovery sites

Platform	Permissions

Privileges on VMware site	Datastore
	• Allocate space
	• Low level file operations
	
	Global
	
	• Disable methods
	• Enable methods
	
	Host
	• Local operations
	  ○ Create virtual machine
	  ○ Delete virtual machine
	  ○ Reconfigure virtual machine
	• vSphere Replication
	  ○ Manage replication
	
	Network
	• Assign network
	
	Resource
	• Assign virtual machine to resource pool
	• Migrate powered off virtual machine
	
	Virtual machine
	• Change Configuration
	  ○ Add existing disk
	  ○ Add new disk
	  ○ Add or remove device
	  ○ Advanced configuration
	  ○ Change CPU count
	  ○ Change Memory
	  ○ Change Settings
	  ○ Change Swapfile placement
	  ○ Configure Host USB device
	  ○ Configure Raw device
	  ○ Extend virtual disk
	  ○ Modify device settings
	  ○ Remove disk

	○ Rename ○ Toggle disk change tracking • Edit Inventory ○ Create from existing ○ Create new ○ Move ○ Register ○ Remove • Interaction ○ Connect devices ○ Create screenshot ○ Power off ○ Power on ○ Reset • Provisioning ○ Allow disk access ○ Allow file access ○ Allow read-only disk access ○ Allow virtual machine download ○ Clone template ○ Clone virtual machine ○ Customize guest ○ Modify customization specification ○ Read customization specifications • Snapshot management ○ Create snapshot ○ Remove snapshot ○ Rename snapshot ○ Revert to snapshot • vSphere Replication ○ Configure replication ○ Manage replication vSphere Tagging (Select All)
Privileges on AWS site	• Datamotive recommends creating a new IAM Policy with permissions mentioned below. • Create a new IAM User with 'Attach policies directly' option. • Attach the new Datamotive policy to the user. • For the new user, under Security Credentials option, create access key. • Choose option 'Application running on an AWS compute service' and proceed with creating access key. • Save the user's Access key and Secret access key. These user details will be configured with Datamotive node while creating AWS Site in the Datamotive Management Console. • If there are encrypted disks, add the user to all KMS keys in primary and DR sites if disk encryption is used. The user should have below permission on the KMS Key.

	"kms:Encrypt", "kms:Decrypt", "kms:GenerateDataKey", "kms:GenerateDataKeyWithoutPlaintext", "kms:ReEncrypt*", "kms:CreateGrant", "ec2:CreateTags", "kms:DescribeKey" • List of permissions for creating IAM Policy is as below { "Version": "2012-10-17", "Statement": [{ "Sid": "AllowAllResources", "Effect": "Allow", "Action": ["ec2:DescribeInstances", "ec2:DescribeSnapshots", "ec2:DescribeInternetGateways", "ec2:DescribeVolumeStatus", "ec2:ModifySnapshotAttribute", "ec2:StartInstances", "ec2:DescribeVolumes", "ec2:DescribeKeyPairs", "ec2:DescribeRouteTables", "ec2:CreateTags", "ec2:DescribeSnapshotAttribute", "ec2:RunInstances", "ec2:CreateVolume", "ec2:DescribeSecurityGroupRules", "ec2:CreateSnapshots", "ec2:DescribeInstanceTypes", "ec2:AssociateAddress", "ec2:DescribeSubnets", "ec2:DescribeVpnGateways", "ec2:AttachVolume", "ec2:DescribeAddresses", "ec2:DescribeRegions", "ec2:DescribeInstanceTypeOfferings", "ec2:DescribeNetworkInterfaces", "ec2:DescribeAvailabilityZones", "ec2:CreateSnapshot", "ec2:DescribeInstanceStatus", "ec2:DescribeTags", "ec2:DescribeNatGateways", "ec2:DescribeLocalGateways", "ec2:DescribeSecurityGroups", "ec2:CreateLaunchTemplateVersion", "ec2:DescribeImages", "ec2:DescribeVpcs",

```
                "sns:GetTopicAttributes",
                "sns:ListTopics",
                "sns:CreateTopic",
                "cloudwatch:PutMetricAlarm",
                "cloudwatch:DescribeAlarms",
                "cloudwatch:PutMetricData",
                "cloudwatch:SetAlarmState",
                "cloudwatch:EnableAlarmActions",
                "cloudwatch:TagResource",
                "kms:ListKeys",
                "kms:ListAliases",
                "ebs:GetSnapshotBlock",
                "ebs:ListChangedBlocks",
                "ebs:ListSnapshotBlocks",
                "ebs:CompleteSnapshot",
                "ebs:PutSnapshotBlock",
                "ebs:StartSnapshot"
            ],
            "Resource": "*"
        },
        {
            "Sid": "AllowCopyFromSource",
            "Effect": "Allow",
            "Action": [
                "ec2:CreateSubnet",
                "ec2:UpdateSecurityGroupRuleDescriptionsIngress",
                "ec2:UpdateSecurityGroupRuleDescriptionsEgress",
                "ec2:RevokeSecurityGroupIngress",
                "ec2:RevokeSecurityGroupEgress",
                "ec2:AuthorizeSecurityGroupIngress",
                "ec2:AuthorizeSecurityGroupEgress",
                "ec2:CreateSecurityGroup",
                "ec2:DeleteSecurityGroup"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:ResourceTag/Protected-By-Datamotive": "Datamotive
protected resource"
                }
            }
        },
        {
            "Sid": "AllowDMResources",
            "Effect": "Allow",
            "Action": [
                "ec2:DetachVolume",
                "ec2:DeleteVolume",
                "ec2:DeleteSnapshot",
                "ec2:DeleteTags",
                "ec2:StopInstances",
```

	"ec2:RebootInstances",
	"ec2:ModifyLaunchTemplate",
	"ec2:TerminateInstances",
	"sns:DeleteTopic",
	"cloudwatch:DeleteAlarms",
	"iam:CreateServiceLinkedRole",
	"iam:AttachRolePolicy",
	"iam:PutRolePolicy",
	"cloudwatch:DisableAlarmActions"
	],
	"Resource": "*",
	"Condition": {
	"StringEquals": {
	"aws:ResourceTag/Protected-By-Datamotive": "Datamotive protected resource"
	}
	}
	}
	]
	}

Roles On Azure Site	• Create new registration in Azure AD -> App Registration, e.g. "DatamotiveApp" • Note down the Tenant ID and Client ID for the newly created app • In the newly created App, Create new Client Secret in section "Certificates & Secrets" • Note down the text in "Value" field of the secret. • The App can be assigned roles at either Subscription or Resource Group level. Select resource groups or subscription (where source and target VMs, networks are expected) and assign following roles to the newly created App. ○ Contributor ○ Storage Blob Contributor (E.g. Resource Group -> Access Control (IAM) -> Add Role Assignment)

4.4. Sandboxed Environment for Test Drills

Datamotive solution provides a mechanism to perform Test Drills. Test Drill functionality can be utilized to ensure the sanity of replicated copy. Using Test Drills, the recovering of protected workloads is similar to what is done during a disaster event. Test Drill operations are non-intrusive to production workloads and to the replication activities. As a part of Test Drill operation, Datamotive recovers the workloads with their last consistent replicated copy in the target site.

To ensure that the test-recovered virtual machines do not intrude into production or the DR infrastructure, Datamotive recommends creating a separate infrastructure including VPCs, networks, subnets which do not have outbound connectivity to either Production

or DR infrastructures. This will ensure that the recovered virtual machines can be tested in isolation and not impact any of the infrastructures in any way.

5. Datamotive Deployment Models

Datamotive solution is built in a highly scalable model. The solution components can be deployed and configured based on number of workloads to be protected. Datamotive recommends following three models of deployment based on the number of workloads to be protected. The recommendations are consistent across supported platforms. For deployment of individual nodes on specific platforms, refer to section Deploy Datamotive Nodes.

5.1 Small and Medium Businesses (SMB)

As seen in the architecture diagram in **Figure 3**, Datamotive solution consists of different types of nodes. However, for a small setup (up to 15 virtual machines or 40 virtual disks), admin will have to deploy only the Datamotive Management Server. This server consists of replication engine capable of transferring changed data. This eliminates the need to deploy additional replication nodes. If the number of virtual machines increase at a later point in time, then additional replication nodes can be added at runtime to handle the increased load. The below diagram illustrates the overall small deployment mode.

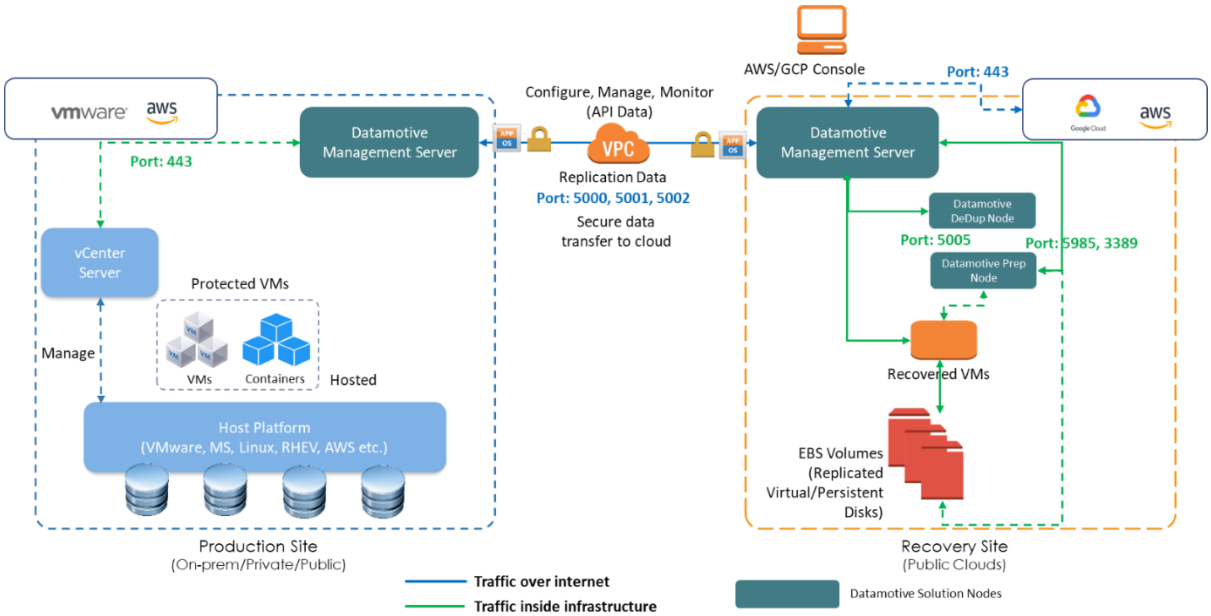

Figure 3: Single management node deployed on source and target sites

5.2. Enterprise

For an enterprise size setup, it is recommended to add additional replication nodes to the Management Server. This helps in balancing the load of data transfer from source to destination. The replication nodes need to be added in pairs on source and target sites. It is recommended to add 1 replication node per 40 virtual disks, e.g., if there are 120

virtual disks, then total replication nodes required would be 120/40=3. The management node also acts as a replication node handling 40 disks. Thus, additional replication nodes required will be 3–1=2. The below diagram illustrates overall enterprise deployment mode.

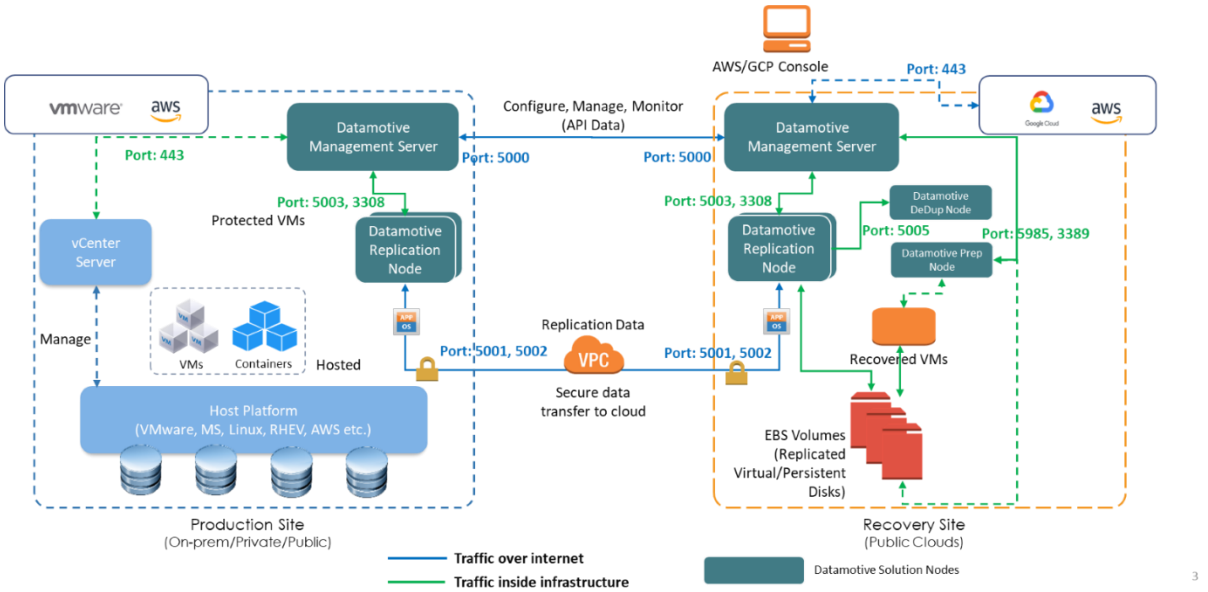

Figure 4: Multiple replication nodes deployed along with Management Server on source and target sites

6. Datamotive Node Instance Configuration

In this section, Datamotive node instance configuration for different platforms is discussed.

6.1. VMware

Table 3: Datamotive node instance configuration for VMware

Node	Instance size	Storage space	Storage type	Description
Management Server	2 vCPUs, 8GB RAM	50 GB	VMFS Datastore	Storage for running Datamotive services and an internal database
Replication Node	2 vCPUs, 8GB RAM	50 GB	VMFS Datastore	Storage for running Datamotive services
DeDupe Node	2 vCPUs, 8GB RAM	500 GB	SDD preferred	Storage size is dependent on anticipated unique replication data size
Windows Prep Node	2 vCPUs, 16GB RAM	50 GB	VMFS Datastore	Storage for running Datamotive services

6.2. AWS

Table 4: Datamotive node instance configuration for AWS

Node	Instance size	Storage space	Storage type	Description

Management Server	2 vCPUs, 8GB RAM	50 GB	GP2	Storage for running Datamotive services and the internal database
Replication Node	2 vCPUs, 8GB RAM	50 GB	GP2	Storage for running Datamotive services
DeDupe Node	2 vCPUs, 8GB RAM	500 GB	GP3 (1000 IOPS preferred)	Storage size is dependent on anticipated unique replication data size
Windows Prep Node	2 vCPUs, 16GB RAM (ex.t3.xlarge)	50 GB	GP2	Storage for running Datamotive services

6.3. Azure

Table 5: Datamotive node instance configuration for Azure

Node	Instance size	Storage space	Storage type	Description
Management Server	2 vCPUs, 8GB RAM	50 GB	Standard SSD	Storage for running Datamotive services and the internal database
Replication Node	2 vCPUs, 8GB RAM	50 GB	Standard SSD	Storage for running Datamotive services
DeDupe Node	2 vCPUs, 8GB RAM	500 GB	Standard SSD	Storage size is dependent on anticipated unique replication data size
Windows Prep Node	2 vCPUs, 16GB RAM (16 data disks which can be attached (ex.F8)	50 GB	Standard SSD	Storage for running Datamotive services

7. Deployment and Configuration of Datamotive Solution

This section discusses the deployment and configuration of Datamotive solution.

7.1 Deploy Datamotive Nodes

Once the pre-requisites are implemented and tested, the admin can deploy Datamotive nodes in it. The below table lists packaging and deployment details for each of the Datamotive nodes.

Table 6: Packaging and deployment details for Datamotive nodes

Datamotive node	Platform	Form-factor	Name	Deployment tool
Management Server	VMware	OVA	dm-***mgmt***-VMware-<version>.ova	OVFTool, vCenter Server GUI, ESXi, Custom script (Steps to Deploy OVA)
	AWS	AMI	dm-***mgmt***-AWS-<version>	AWS Console
	Azure	Machine Image	dm-***mgmt***-Azure-<version>	Azure Console
Replication Node	VMware	OVA	dm-***repl***-VMware-<version>.ova	OVFTool, vCenter Server GUI, ESXi, Custom script (Steps to Deploy OVA)
	AWS	AMI	dm-***repl***-AWS-<version>.ova	AWS Console
	Azure	Machine Image	dm-***repl***-Azure-<version>	Azure Console
Dedupe Node	VMware	OVA	dm-***dedupe***-VMware-<version>.ova	OVFTool, vCenter Server GUI, ESXi, Custom script (Steps to Deploy OVA)
	AWS	AMI	dm-***dedupe***-AWS-<version>	AWS Console
	Azure	Machine Image	dm-***dedupe***-Azure-<version>	Azure Console
Windows Prep Node	VMware	OVA	dm-***win-prep***-VMware-<version>	OVFTool, vCenter Server GUI, ESXi, Custom script (Steps to Deploy OVA)
	AWS	AMI	dm-***win-prep***-AWS-<version>	AWS Console
	Azure	Machine Image	dm-***win-prep***-Azure-<version>	Azure Console

7.2. Configure Datamotive Nodes

Once the required Datamotive nodes are deployed, the admin must follow a simple node-specific setup process as described below. For each node, configure the deployed virtual

machine's networking by using network configurations created during configuration of platform.

1. Management Node
   1.1. Deploy the virtual machine following platform specific tools
   1.2. Restart the virtual machine and open URL https://<VM IP_Address>:5000 in browser to validate the GUI. Datamotive Login screen should be rendered
   1.3. Default credentials of the system are Administrator/admin. Login using the default credentials and change the default password
   1.4. For more information, refer to Datamotive-User-Guide
2. Replication and DeDupe Node
   2.1. Deploy the virtual machine following platform specific tools
   2.2. Open GUI of Datamotive Management Server deployed earlier. Navigate to Nodes section in left navigation bar and click on Add Node. Enter the details of the node in the form. Make sure that the name of the node is same as name of the deployed virtual machine. Default credentials of the node are Administrator/admin.
   2.3. If the node gets added successfully, the replication node is deployed correctly.
   2.4. For more information, refer to Datamotive-User-Guide
3. Windows Prep Node
   3.1. Deploy the virtual machine following platform specific tools
   3.2. Default credentials of the system are Administrator/M0v3@nywh3r3.
   3.3. Open GUI of Datamotive Management Server deployed earlier. Navigate to Nodes section in left navigation bar and click on Add Node. Enter the details of the node in the form. Make sure that the name of the node is same as name of the deployed virtual machine.
   3.4. If the node gets added successfully, the windows prep node is deployed correctly.
   3.5. For more information, refer to Datamotive-User-Guide

Note:

- The deployed Datamotive nodes must have ssh service up and running
- For cloud platforms, all the nodes must be deployed in same region
- For Azure, the DM nodes should be created with No-Zone resiliency. This would enable multiple zone recovery support

7.3. SSL Certificate Management

Datamotive nodes communicate with each other using REST APIs. The APIs and GUI are served over TLS1.3 connection. By default, Datamotive nodes are shipped with self-signed certificate. Admins can update the Datamotive nodes to use CA signed or other custom certificate of their choice. Currently, the feature is provided using Datamotive CLI. To replace the default self-signed SSL certificate, follow the below steps:

Step 1: Copy the certificate file and private key file to a temporary folder in deployed Datamotive management node, e.g., /tmp. The supported certificate file formats are .crt and .pem. In case the CA has provided the certificate in a different format, convert it to .pem format. Check Cryptography material conversion and verification commands GitHub for details on how to convert SSL certificate formats.

Step 2: ssh into Datamotive management node and navigate to /opt/dmservice/bin folder.

```
ssh user@<Datamotive Management Node IP>

sudo su

cd /opt/dmservice/bin
```

Step 3: Execute Datamotive CLI to update the certificates

```
./dmcli update_certificate --cert-file <Path to certificate file> --cert-private-key <Path to certificate private key file> --nodeType Management --username administrator
e.g../dmcli update_certificate --cert-file bbcert.crt --cert-private-key bbcert.key --nodeType Management --username administrator
```

This operation restarts the API server on all Datamotive nodes post certificate replacement. Thus, it checks if there are any activities being performed in the server. The server and the nodes must not perform any data movement activities while executing this command.

8. Verification

To verify if installation was successful, open browser and type URL: https://<IP-address of virtual machine for Datamotive Replication Server node>:5000. The login screen is displayed as shown in

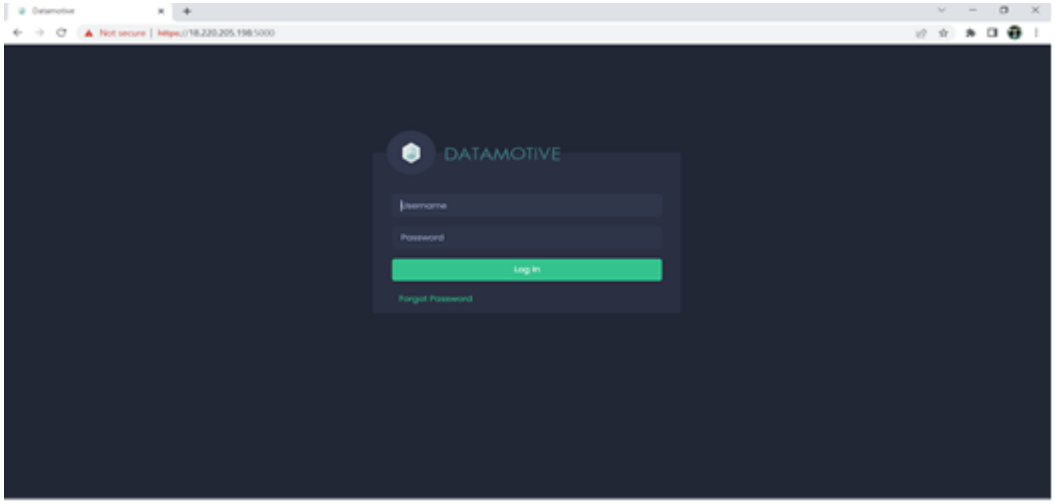

Figure 5.

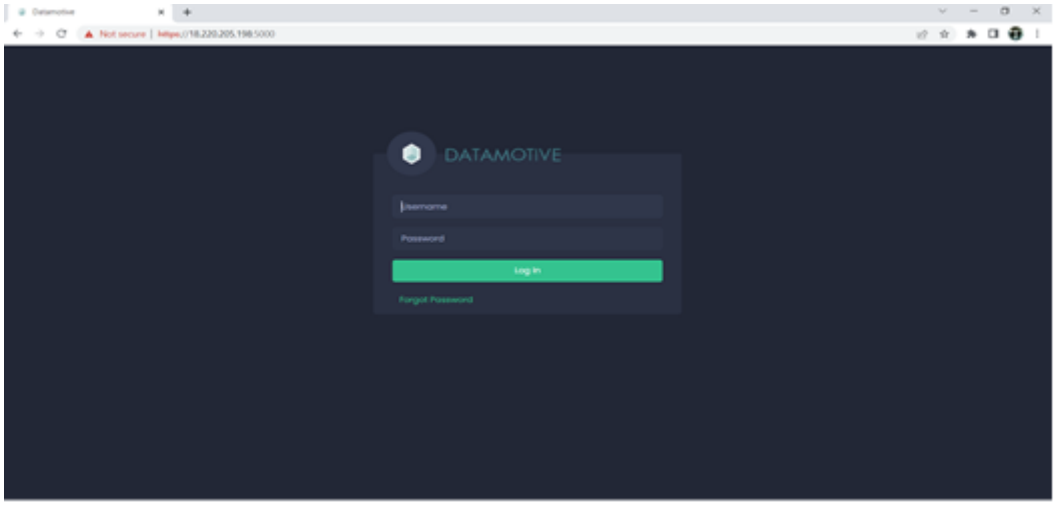

Figure 5: Login screen

Note:
Datamotive ships with self-signed certificate so first time, the browser will show untrusted certificate warning.

8.1. Datamotive Services

Once the Datamotive UI is displayed on screen, login using credentials provided by Datamotive team. If the URL is unreachable, or the UI is not displayed on screen, ssh into the Datamotive management node and check service status using following commands:

For Datamotive Replication Management Server Node:

```
systemctl status mysqld
systemctl status dm-mgmt
systemctl status dm-repl-server
systemctl status dm-repl-client
systemctl status dm-mon
```

All the above services must be up and running. In case the services are not running, restart them in following order:

```
systemctl restart mysqld
systemctl restart dm-mgmt
systemctl restart dm-repl-server
systemctl restart dm-repl-client
systemctl restart dm-mon
```

Contact Datamotive support support@datamotive.io if there are issues starting any of these services.

8.2. Networking

Once the Datamotive services are accessible, validate if the networking is configured appropriately with below checks:

1. **Management Node:**
   1.1. ssh into management node with given credentials
   1.2. Validate management node to replication node communication
   `telnet <replication node IP address>:5003`
   `telnet <replication node IP address>:5002`
   1.3. Validate management node to remote management node communication
   `telnet <Remote management node IP address>:5000`
   1.4. Validate management node to prep node communication
   `nc -z -w1 <Prep node IP or host name> 5985; echo $?`
   1.5. Validate management node to DeDupe node communication
   `telnet <DeDupe node IP address>:5005`
   1.6. In case of cloud deployment, validate external access
   - for example, ping www.google.com
2. **Replication Node:**
   2.1. ssh into replication node with given credentials
   2.2. Validate replication node to management node communication
   `telnet <management node IP address>:5000`
   `telnet <management node IP address>:3085`
   2.3. Validate replication node to remote replication node communication
   `telnet <remote replication node IP address>:5001`
   `telnet <remote replication node IP address>:5002`
   2.4. In case of cloud deployment, validate external access
   - for example, ping www.google.com

If the problem persists, contact support team of Datamotive: support@datamotive.io

9. Upgrade

The Upgrade feature is used to upgrade the Datamotive services, features and the Datamotive UI files. Datamotive provides upgrade packages as tar bundles. The upgrade feature is provided through easy-to-use CLI available as part of Datamotive Management node. Upgrades need to be done individually on management nodes of all sites. To upgrade Datamotive solution, follow below mentioned steps on the Management server only for server and configured replication nodes.

Note:

Datamotive management nodes with different versions on different sites may result in inconsistent behaviour and may impact replication and recovery operations

9.1. All Nodes

Steps to upgrade the Datamotive Server and configured replication nodes:

- Copy the "UPGRADE_PKG <version>.tar.gz" package to the Management server via Win-scp tool/scp command or any other tool at "/home/dmadmin"
- Once the package is uploaded to the server, ssh into Datamotive management node with user as 'dmadmin' in VMware, Azure and GCP and "ubuntu" in AWS.
- Get into sudo mode, sudo su.
- Navigate to directory /home/dmadmin and untar the bundle
- Use the dm-upgrade (executable to upgrade DM nodes) to upgrade the Datamotive Management and associated replication nodes.
- Use below command to upgrade.
  ./dm-upgrade <upgrade_bundle_name> (extracted upgrade bundle name) Management
- When asked for password, enter the password for given username. The username is Datamotive application username.
- Upgrade command details
  The extracted tar bundle name is DM_UPGRADE_<version>.tar.gz on current management server.
  For example, if the Datamotive Upgrade file DM_UPGRADE_<version>.tar.gz is observed after extracting the UPGRADE_PKG_<version>.tar.gz then the command would be

  ./dm-upgrade DM_UPGRADE_<version>.tar.gz Management

After the successful upgrade of the server and the node, the success message is displayed as shown in

```
root@dm-mgmt-rk:/home/dmadmin# ./dm-upgrade DM_UPGRADE-1.0.0-1436.tar.gz Management
Creating DM Upgrade directory
Copying upgrade package
Extracting upgrade package
Datamotive@123
Invoking dmcli to upgrade services
2023/07/04 10:12:57 Upgrading Datamotive nodes ...
password: **************
2023/07/04 10:13:03 Login successful
2023/07/04 10:13:03 Node: 20.204.1.64 current installed version: 1.0.0-1427 can be upgraded to: 1.0.0-1436
2023/07/04 10:13:03 If you continue, the replication jobs will be paused and resume after upgrade. Or else please stop the plan and wait
 for the replication jobs to finish before trying upgrade again.
? Do you want to continue?2023/07/04 10:13:21 Upgrading node:  20.204.1.64
2023/07/04 10:13:22 Creating backup for current service version( 1.0.0-1427 ) ...
2023/07/04 10:13:22 Performing action: backup on DM Node DB ...
2023/07/04 10:13:22 Operation: backup on database completed successfully
2023/07/04 10:14:45 Service backup created successfully
2023/07/04 10:14:45 Starting services update ...
2023/07/04 10:14:45 Performing action: upgrade on DM Node DB ...
2023/07/04 10:14:45 Operation: upgrade on database completed successfully
2023/07/04 10:16:18 Updated system with latest Datamotive binaries
2023/07/04 10:16:19 Datamotive services updated successfully
2023/07/04 10:16:19 Upgrade successful for node: 20.204.1.64 in 2m58.688853914s
2023/07/04 10:16:22 Upgrade process completed in: 3m24.903380832s
2023/07/04 10:16:22 Datamotive nodes  [20.204.1.64]  upgraded successfully
root@dm-mgmt-rk:/home/dmadmin#
```

Figure 6.

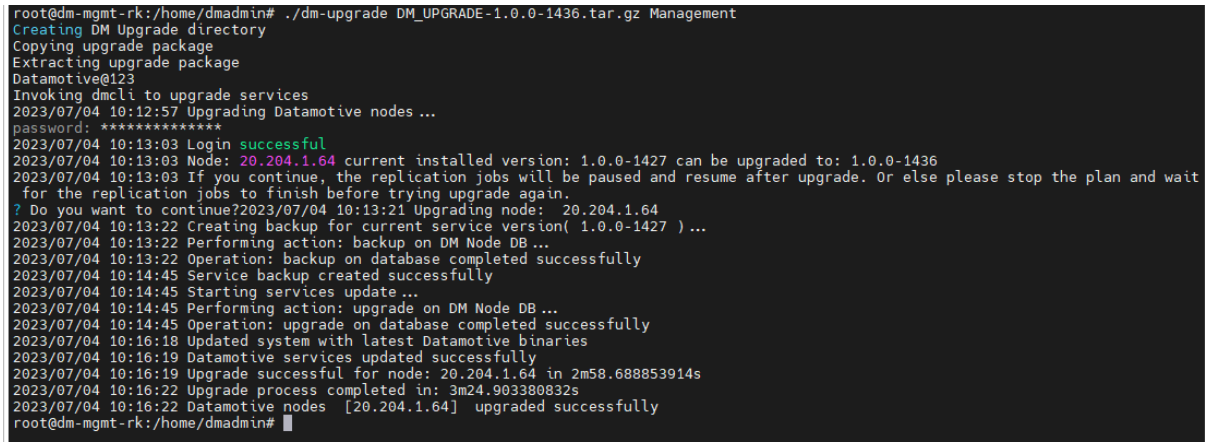

Figure 6:Success message

9.2. Replication Nodes Only

There could be cases where the Datamotive management nodes are upgraded along with configured replication nodes. When adding a new replication node to scale up the environment, after registering the replication node with the management node, make sure to upgrade the replication node first before using it in protection plans. To upgrade specific replication node, follow below steps:

1. Copy the latest UPGRADE_PKG_ <version>.tar.gz bundle/package to the replication node which needs to be upgraded via Win-scp tool /scp command or any other tool
2. ssh into the replication node. If on VMware, the default username/password would be dmadmin/M0v3@nywh3r3. If on cloud, use the key-based ssh authentication provided by the cloud natively
3. Once the package is uploaded to the server, go to datamotive 'dmadmin' user directory and untar the bundle
4. Use the dm-upgrade (executable to upgrade DM nodes) to upgrade the Datamotive replication and associated replication nodes.
5. Enter following command to upgrade:
   ./dm-upgrade <upgrade_bundle_name> (extracted upgrade bundle name) Replication

6. Once the node is successful, a success message is displayed on command line.
7. If there are errors, refer to /opt/dmservice/logs/ directory for the issues.

10. Support

In case there are any issues or queries while using the solution, contact support@datamotive.io.