

Nutanix Database Service Administration Guide

Nutanix Database Service (formerly Era) 2.5
March 15, 2024

Contents

NDB Overview.....	6
NDB Workflow.....	6
NDB Limitations and Notes.....	7
NDB Services.....	10
One-Click Provisioning.....	10
Copy Data Management (CDM).....	10
Database Protection.....	11
One-Click Patching.....	11
NDB Terminology Reference.....	12
NDB Control Plane Configuration and Scalability.....	13
Modifying VM Resource Configurations.....	17
Modifying the Tomcat Configuration.....	17
Modifying the Repository Database Configuration.....	18
Modifying Log Copy and Log Catchup Configurations.....	19
Modifying NDB Service Configuration.....	19
Modifying HAProxy Settings.....	19
Modifying File Descriptor Limits.....	21
Modifying the NDB UI's Polling Interval for Operation Summary API Call.....	22
NDB with Nutanix Objects.....	22
Registering Nutanix Objects for Storing NDB Log Backups.....	22
Updating a Storage Resource Registered with NDB.....	23
Removing a Storage Resource Registered with NDB.....	24
Provisioning Databases that Use Nutanix Objects.....	25
Registering Databases that Use Nutanix Objects.....	26
Configuring a Data Access Management Policy for a Time Machine.....	27
 NDB Installation.....	 29
NDB Network Requirements.....	29
NDB Network Requirement Diagrams.....	31
Installing NDB on AHV.....	34
Installing NDB on ESXi.....	37
 NDB Initial Configuration.....	 38
Assigning a Static IP Address to the NDB VM by Using the Console.....	38
Configuring an SSL Certificate.....	38
Logging on to NDB.....	39
Checking NDB Version.....	39
Changing the Language Settings.....	40
 Get Started with NDB.....	 42
Welcome to NDB Wizard.....	42
NDB Profiles.....	48
Getting Started with Provisioning Databases.....	49
 NDB Dashboard.....	 55

NDB Administration.....	57
NDB Service Management.....	57
NDB High Availability.....	65
NDB Service Resiliency.....	71
Enabling NDB Pulse Telemetry.....	73
Role-based Access Control.....	74
Privileges and Permission.....	75
Adding a User.....	76
Updating an Existing User.....	78
Removing a User.....	78
Configuring User Profile Settings.....	78
Adding, Updating, and Removing a Group.....	80
Refreshing a Group.....	80
Viewing Users' Roles.....	81
Creating a Role.....	81
Updating a Role.....	84
Removing a Role.....	84
Configuring Active Directory Access.....	84
NDB Entity Sharing.....	85
Entity Sharing Policies.....	86
Creating an Entity Sharing Policy.....	88
Updating an Entity Sharing Policy.....	90
Nutanix Cluster Management.....	91
Enabling NDB Multi-Cluster.....	91
Registering a Nutanix Cluster with NDB.....	93
Updating a Nutanix Cluster.....	97
Setting VM Host Affinity.....	99
NDB Network Management.....	101
Adding a VLAN to NDB.....	102
Updating a Static VLAN.....	103
Adding IP Addresses to a Static VLAN.....	105
Displaying the Static IP Address Assignment.....	105
Adding a Stretched VLAN to NDB.....	106
Updating a Stretched VLAN.....	107
Adding an AHV-managed VLAN as NDB-managed.....	110
Configuring NTP Servers.....	111
Configuring DNS Servers.....	111
Configuring the Time Zone.....	112
Configuring the NDB Proxy.....	113
NDB Tags Management.....	113
Creating a Tag.....	114
Updating a Tag.....	115
Deleting a Tag.....	115
NDB Upgrade Management.....	116
Upgrading NDB (One-Click Upgrade).....	116
Upgrading NDB (Offline Upgrade).....	117
Downloading the Diagnostics Bundle.....	118
Downloading the Diagnostics Bundle using the NDB CLI.....	119
Operation Logs.....	120
Viewing Operation Logs.....	120
Show Logs Page.....	122
Maintenance Window.....	123
Creating a Maintenance Window.....	124
Associating a Maintenance Window.....	125

Operating System Patching.....	125
Patching an Operating System.....	126
Database Management.....	128
Updating a Database.....	128
Removing a Database.....	128
Scaling a Database.....	129
Updating a Database Server VM or Database Server Cluster.....	129
Updating the IP Address of a Registered Database Server VM.....	130
Removing a Database Server VM or Database Server Cluster.....	131
Vertical Scaling of Database Server VMs.....	132
Clone Management.....	134
Updating Database Clone.....	134
Refreshing Database Clones (Manual).....	134
Creating a Clone Refresh Schedule.....	135
Deleting a Clone Refresh Schedule.....	136
Removing a Database Clone (Manual).....	136
Creating a Clone Removal Schedule.....	137
Deleting a Clone Removal Schedule.....	138
NDB Time Machine Management.....	139
Time Machine Behaviour and Functionality.....	139
Updating a Database Time Machine.....	140
Data Access Management.....	141
Adding Time Machine Data Access to a Nutanix Cluster.....	142
Updating Time Machine Data Access to a Nutanix Cluster.....	143
Snapshot Management.....	144
Creating Snapshots (Manual).....	145
Replicating Snapshots (Manual).....	146
Sanitized Snapshots.....	148
Database Log Management.....	151
Database Log Backup Management in NDB.....	151
Performing a Log Catch-up Operation (Manual).....	153
SLA Management.....	153
Built-in SLAs.....	154
Creating Custom SLAs.....	155
Comparing SLAs.....	156
Authorizing Database Server VMs.....	158
Deauthorizing Database Server VMs.....	158
Time Machine Heal Snapshot.....	158
Alert Management.....	160
Alerts Dashboard.....	160
Alert Policies Management.....	160
Disabling or Enabling Alert Policies.....	161
Setting Alert Retention.....	161
Resolving Alerts.....	162
Alert Notifications.....	163
Sending Alert Notifications to an Email Recipient.....	163

NDB Security.....	166
NDB Architecture.....	166
NDB Security Implementation.....	168
Account Management.....	168
Authentication.....	175
Network Placement.....	176
Role-based Access Control.....	176
Patch Management.....	176
SSL Server Certificate Authentication.....	176
Entity Sharing.....	178
NDB Utilities.....	179
Database Server VM Registration Pre-requirement Checks.....	179
Downloading Pre-requirement Script (CLI).....	179
Downloading Pre-requirement Script (GUI).....	180
Running the Pre-requirement Script (Linux).....	180
Downloading the Uninstall Script (GUI).....	181
NDB Uninstallation.....	182
Uninstalling NDB.....	182
NDB Shutdown.....	183
Shutting Down NDB.....	183
Copyright.....	184

NDB OVERVIEW

Nutanix Database Service (NDB) automates and simplifies database administration, bringing one-click simplicity and invisible operations to database provisioning and life-cycle management.

NDB enables database administrators to perform operations such as database registration, provisioning, cloning, patching, and restore. It allows administrators to define provisioning standards with end-state driven functionality that includes network segmentation, high availability (HA) database deployments, and much more.

With NDB multi-cluster you can easily manage databases across multiple locations, both on-prem and in the cloud, with Nutanix clusters.

For information about the supported operating systems, database server VMs, and Nutanix software versions, see *NDB Software Compatibility and Feature Support in Nutanix Database Service Release Notes*.

For information on database operations, see the following database management guides:

- [Oracle Database Management Guide](#)
- [SQL Server Database Management Guide](#)
- [PostgreSQL Database Management Guide](#)
- [MongoDB Database Management Guide](#)
- [MySQL Database Management Guide](#)
- [MariaDB Database Management Guide](#)

NDB Workflow

The following diagram explains how you use NDB as a database service.




Figure 1: NDB Workflow Diagram

The following is the database-as-a-service (DBaaS) workflow in NDB.

1. Install NDB.

Install the NDB VM on a Nutanix cluster.

2. Select a profile.

Either use the sample profiles available in NDB, update the sample profiles to suit your requirements, or create a profile.

3. Provision or register a database VM.

Either provision a new source database VM on your Nutanix cluster by using NDB or register an existing source database VM with NDB.

Note: The existing source database VM must be running on the same Nutanix cluster as NDB.

As part of the database provisioning or registration process:

1. NDB registers the VM used by the database as a database server.
2. NDB creates a time machine for that source database. The time machine takes snapshots and copies transaction logs as defined in the schedule when you provision or register the source database.
4. Once you have a source database VM, you can easily clone the source database to create more database instances.
5. Once you have the database instance set up, you can proceed to use other NDB features such as copy data management.

NDB Limitations and Notes

Limitations

- NDB currently supports English, Japanese, and German languages for database server VM operating system and database software.
- NDB supports the database name on the database server VM, VM name, graphical user interface (GUI), and command-line interface (CLI) only in English and Japanese.
- NDB does not support database server VMs protected by NearSync, metro availability, Leap, or Xi Leap.
- NDB does not support forward or reverse proxy between:
 - Database server VMs and NDB.
 - A client accessing NDB API or NDB user interface through a browser.
- NDB currently supports SCSI disks and does not support SATA disks.
- NDB does not support new registrations of database server VMs if a mix of SCSI and non-SCSI disks (having CD-ROMs does not affect the registration) are present. NDB does not support the registration even if the storage is of non-OS, non-database, or non-software disks. This restriction does not affect existing database server VMs with non-SCSI disks that are already registered with NDB.
- NDB does not support the manual addition of disks to the staging drive of database server VMs.
- NDB supports a maximum of 15 disks per controller on an ESXi VM.
- NDB does not support the following migration use cases:
 - Migrating NDB server and NDB agent to a different storage container within the same Nutanix cluster.
 - Migrating NDB server and NDB agent to a different Nutanix cluster.
 - Migrating NDB managed UVMs to a different Nutanix cluster.
- NDB does not support the addition of a database server VM that is provisioned with NDB to a protection domain (PD). If the database server VM is already a part of a protection domain, you can register the VM with NDB. However, this does not provide disaster recovery for the database server VM.
- NDB control plane does not have a disaster recovery solution.

- NDB does not support the use case where database snapshots and ESX VM snapshots co-exist.

Notes

- Perform the following steps to change the underlying hardware for any Nutanix cluster where NDB control plane VMs and database server VMs managed by NDB run:

- Add the new Nutanix nodes to the existing Nutanix cluster.
- Remove the old Nutanix nodes from the existing Nutanix cluster.

After the change, the Nutanix cluster stays the same but runs with new hardware.

- NDB supports a combination of compute-only and storage-only nodes in a Nutanix cluster.
- NDB supports Nutanix Cloud Clusters on AWS (NC2) for SQL Server, PostgreSQL, and MongoDB database engines. For more information on NC2, see [Nutanix Cloud Clusters on AWS Deployment and User Guide](#).
- NDB user interface can display content in English, Spanish, Japanese, and Chinese locale. For any other locale, NDB user interface displays content in English.

To change the language settings of the NDB user interface see [Changing Language Settings](#) topic in this guide.

- Deploy and configure the NDB VM on one of the Nutanix clusters before it can start managing the database server VMs on that cluster.
- NDB supports dark site deployments or deployments which are disconnected from the internet.
- NDB does not support databases running on platforms other than Nutanix.

Note: To clone databases running on other platforms, you must first replicate the source database to a database server VM. You can either replicate the database server VM and the database to Nutanix or you can replicate a database to a database server VM already running on the Nutanix platform. For example, you can use Oracle Data Guard to replicate the Oracle databases.

- Correct directory and file permissions are critical for NDB to function properly. Do not change directory or file permissions for the following directories, sub-directories, and files included in the directories:
 - /opt/era_base
 - /home/<ndb-user-name>
 - Binaries for Patroni, etcd, Keepalived, HAProxy
 - /etc/etcd
 - /etc/patroni
 - /etc/haproxy
 - /etc/keepalived
 - /etc/systemd/system/patroni.service
 - /usr/lib/systemd/system/etcd.service
 - /usr/lib/systemd/system/patroni-watchdog.service
 - /var/lib/etcd/data
 - /var/log/etcd
 - /var/log/patroni
 - /home/postgres/era_custom_pg_params.conf
 - /tmp
 - /tmp/perfrom_work.lock
 - /var/tmp
 - /etc/mongod.conf
 - Database engine and NDB software related directories, for example: database data, database logs, archive logs, and software_home
 - POSTGRES_SOFTWARE_HOME path (e.g. /usr/pgsql-10.4)
 - ORACLE_HOME path (e.g. /u02/app/oracle/product/19.0.0/dbhome_1)
 - GRID_HOME path (e.g. /u01/app/19.0.0/grid)
 - NDB log drive
 - C:\NTNX\
 - SQL Server database engine related directories:
 - Default software path : C:\Program Files\Microsoft SQL Server\MSSQL13.CDMINSTANCE\MSSQL\Binn\sqlservr.exe
 - Default database data\log path : C:\Program Files\Microsoft SQL Server\MSSQL13.CDMINSTANCE\MSSQL\DATA
 - NDB Provisioned default database data and log path : C:\NTNX\ERA_DATABASES\<database_name>

Note:

- You must not set the noexec option for /tmp and /var/tmp directories. You must give read, write, and execute permissions to /tmp and /var/tmp directories for 'other' users.
 - You must not remove any files from the /opt/era_base directory and sub directories.
- Nutanix recommends the source code (unzip) method for installing PostgreSQL and MongoDB. Installation using Linux package managers such as YUM or DNF is not recommended. Patching through NDB is not supported if you use a Linux package manager to install the database engine. This includes both patching (minor upgrade) via Linux package manager and a new software profile, and patching from outside of NDB.
 - You must assign unique IP addresses for database server VMs managed by the same NDB instance.
 - You must install ifup/ifdown packages in RHEL 8.x software profiles.
 - Nutanix recommends enabling fstrim operations for volume groups. For more information, see [Nutanix Volumes](#) documentation.

NDB Services

NDB offers the following key services for several database engines:

- One-Click Provisioning.
- Copy Data Management.
- Database Protection.
- One-Click Patching.

One-Click Provisioning

NDB enables you to easily provision database environments (either production or otherwise) on your Nutanix clusters.

NDB provisioning service includes the following components:

- **Database engines:** Custom software images that can be tailored to fit your needs. (See [Oracle Database Management Guide](#), [SQL Server Database Management Guide](#), [PostgreSQL Database Management Guide](#), [MySQL Database Management Guide](#), [MariaDB Database Management Guide](#), and [MongoDB Database Management Guide](#) for more information)
- **Database profiles:** Customizable database profiles for software, compute, networking, and database parameters. (See [NDB Profiles](#) on page 48 for more information)
- **Database recovery SLAs:** Customizable recovery SLAs for continuous, daily, weekly, monthly, and quarterly Recovery Point Objectives (RPOs). (See [SLA Management](#) on page 153 for more information)

Copy Data Management (CDM)

NDB enables you to clone your databases and refresh the database clones by using snapshots or transaction logs.

NDB uses storage-efficient Nutanix snapshots, thus lowering the cost of storing multiple copies of your databases. NDB allows you to clone and refresh your clones to a point in time, making the clone and refresh operations highly granular.

The NDB copy data management service includes the following components:

- **Time machine:** Collects logs and snapshots from the database. See [NDB Time Machine Management](#) on page 139 for more information.

- **One-click cloning and refresh:** Zero-byte database clones to any point in time. See [Clone Management](#) on page 134 for more information.

With NDB, you can do the following:

- Create copies of your production databases for development, testing, reporting, or auditing purposes.
- Create copies (clones) of your databases on newly created database server VMs.

To use CDM, do the following:

1. Register a source database VM running on a Nutanix cluster with NDB.
2. Either register a target database server VM with NDB or allow NDB to create a database server VM during the cloning of the source database.

NDB creates a time machine for the database after you have registered a database VM. The time machine takes periodic database snapshots and log catch-ups (see [Time Machine Behaviour and Functionality](#) on page 139). The time machine performs the snapshot and log catch-up operations as defined in the SLA that you have selected or created during the registration of the database (see [SLA Management](#) on page 153).

Database Protection

NDB protects your database with full database consistent backups within a matter of minutes.

With NDB, you can do the following:

- One-click snapshots for simplifying database rollback.
- Explore your backup data through a robust time machine.
- Back up your databases instantaneously with no overhead.

Database Restore

Restoring allows you to restore a source database registered with NDB to a snapshot or point in time supported by the source database time machine. With NDB you can restore a database by using a snapshot ID, the point-in-time recovery (PITR) timestamp, any recent snapshot, or the maximum recovery point possible.

A database restore operation replaces the current database with data as of the specified snapshot or point in time. The time machine for the source database pauses before a restore operation is initiated. After a successful restore, the time machine automatically resumes, triggering a new snapshot and log catch-up operation for the restored database. If the restore operation fails, the database remains in the state that it was in before the restore was initiated.

See [NDB Time Machine Management](#) on page 139 for more information.

One-Click Patching

NDB one-click out-of-band patching allows you to efficiently apply critical OS and database software updates. There are two distinct modes of database patching in NDB: OS and database software.

- **OS patching** - This mode allows you to patch the OS and other software. This does not include database software or the components associated with implementing the database, for example, Oracle ASM or database related HA drivers are not patched during OS patching. NDB supports OS patching for Oracle, PostgreSQL, and MongoDB database server VMs. For more information, see [Operating System Patching](#) on page 125.
- **Database software patching** - This mode allows you to patch or upgrade database software and associated components (for example, HA software) using an NDB software profile version. Creating a software profile version takes different forms. For SQL server, you can upgrade the database software by providing an ISO. For other databases, you update the software profile.

NDB enables you to patch database server VMs and database software with ease. You can apply the updates to the database server VM on-demand or in a scheduled manner. For more information, see the topics on patching for Oracle, Microsoft SQL Server, PostgreSQL, and MongoDB databases in the respective database management guides.

NDB Terminology Reference

This section defines the terms and technologies referenced in this guide. The terms and technologies are described within an NDB-specific context.

Databases

Databases contain your application data. Source databases are registered or provisioned in NDB, and cloned databases contain the data cloned from your source databases.

Database Server VM

A database server VM is a VM that contains an instance of the database engine (Oracle, SQL Server, PostgreSQL, MariaDB, MySQL, or MongoDB) software.

Greenfield Database

A database that is provisioned in NDB.

Brownfield Database

An existing database that is registered with NDB.

Greenfield Database Server VM

A database server VM provisioned in NDB.

Brownfield Database Server VM

An existing database server VM that is registered with NDB.

Clone or Cloned Database

A copy of data from your source database captured at a point in time by using a time machine. You can modify the cloned database and refresh clones as needed from the latest set of data available in a time machine.

Clustered Database

A clustered database consists of two or more database server VMs. In a clustered database, multiple instances of the same database run on different database server VMs, which can be set up across different Nutanix clusters. For example, RAC for Oracle, AG for SQL Server, and PostgreSQL HA instance for PostgreSQL.

NDB Database Agent

An NDB service that runs on a database server VM. It is responsible for scheduling and monitoring the NDB-related operations on the database server VMs.

NDB Drive

An NDB drive is attached to the source database server VM. It temporarily stores the transaction logs of the source database before NDB copies the logs.

Log Catch-up Operation

A log catch-up operation copies transaction logs from the source database based on a specified schedule. The schedule can be provided during database registration or provisioning or can be modified later. For more information, see [Database Log Management](#) on page 151.

Nutanix Cluster

Hosts all the components, such as source databases and database server VMs, target databases and database servers, and NDB).

Profiles

Templates of database software, networking, compute, database parameters, and Windows domains that NDB uses to provision databases or database server VMs.

Source Cluster

A Nutanix cluster associated with a time machine hosting one or more source databases for the same time machine.

SLAs

SLAs are data retention policies that specify for how long NDB retains the log backups and snapshots taken by the time machine (daily, monthly, weekly, or quarterly snapshots). For more information, see [SLA Management](#) on page 153.

SQL Server Database Group

A database group is a logical entity containing one or more SQL Server databases (provisioned or registered) which share a single time machine. For more information, see [SQL Server Database Management Guide](#).

Time Machine

A time machine maintains database snapshots and transaction logs. For every source database you provision or register with NDB, NDB creates a time machine. You can create clones and refresh clones from the snapshots or transaction logs captured by the time machine. For more information, see [Time Machine Behaviour and Functionality](#) on page 139.

Windows Server Failover Cluster (WSFC)

A group of independent servers that work together to increase the availability of applications and services. Microsoft SQL Server uses WSFC services to support Always On Availability Groups (AG) and SQL Server failover cluster instances.

NDB Control Plane Configuration and Scalability

Control Plane Configuration

The following are the NDB service topology options:

- Single cluster, multi cluster not enabled, NDB HA not enabled: single Nutanix cluster managed without HA.
- Multi cluster enabled, NDB HA not enabled: multiple Nutanix clusters managed without HA.
- Multi cluster enabled, NDB HA enabled on a single cluster.
- Multi cluster enabled, NDB HA enabled across clusters: multiple Nutanix clusters managed, NDB control plane HA enabled across three Nutanix clusters.

For more information, see [NDB High Availability](#) on page 65.

The number of time machines supported by NDB on a single Nutanix cluster are as follows:

- NDB control plane non-HA: up to 100 time machines per Nutanix cluster without using Nutanix Objects for log backups.
- NDB control plane non-HA: up to 400 time machines per Nutanix cluster using Nutanix Objects for log backups.
- NDB control plane HA: up to 100 time machines per Nutanix cluster without using Nutanix Objects for log backups.

Note:

- Nutanix does not support a two cluster topology for NDB HA.

- You need a subscription to Nutanix Objects to store log backups in the Nutanix Objects store. For more information, see [KB 13612](#).
- A database can only be managed by one NDB instance.

For more information about NDB single cluster and multi cluster architecture, see [NDB Architecture](#) on page 166.

The following table shows the IP address requirements and VM distribution for different NDB topologies.

IP Address Requirements and VM Distribution for Different NDB Topologies

Topology Description	IP Address per VM		VM distribution	
	NDB service VMs	IP addresses	Total VM Count	VM cluster placement
Single cluster (Multi cluster not enabled, NDB HA not enabled)	NDB API server	1	1	Primary cluster
Multi cluster enabled (NDB HA not enabled)	NDB API server	1	1	Primary cluster
	NDB agent	1	1**	1 per cluster
Multi cluster enabled, NDB HA enabled on a single cluster	NDB API server	1	3	Primary cluster
	NDB agent	1	1**	1 per cluster
	HAProxy	1*	2	Primary cluster
	Repository DB	1	3	Primary cluster
Multi cluster enabled, NDB HA enabled across clusters	NDB API server	1	3	1 per cluster
	NDB agent	1	1**	1 per cluster
	HAProxy	1*	2	1 per cluster
	Repository DB	1	3	1 per cluster

* - HAProxy VMs cumulatively require one additional IP address for NDB HAProxy virtual IP.

** - One NDB agent VM is deployed on each cluster registered with the NDB server.

Note:

- In multi-cluster topologies:
 - **NDB server** - In a non-HA configuration, the NDB server hosts all components of the NDB service control plane. When you enable HA, the original NDB server becomes the first NDB API server as part of the HA scale-out, which is why it does not exist in HA topology options.
 - **NDB agents** - The table assumes a three-cluster topology. One additional NDB agent VM and IP address are required per additional cluster added.

Scalability Values for up to 200 and up to 800 Time Machines

The configuration tables in this section depict setups of up to 200 and up to 800 time machines (TMs) with the following features:

- 200 TMs spread across three clusters.
 - Multi cluster enabled, with or without enabling HA.
 - HA (if enabled) across three clusters.
 - A maximum of 100 TMs in one cluster.
 - 30 minutes log catchup frequency.
- 800 TMs spread across three clusters.
 - Multi cluster enabled (NDB HA not enabled).
 - 30 minutes log catchup frequency.
 - Log backup using Nutanix Objects store.

Note: You must apply the recommended configurations for NDB to function at scale.

Table 1: Recommended VM Resource Configurations for up to 200 TMs

Service	vCPU	Cores	Memory
NDB Server	8	2	16 GB
Repository DB*	4	2	16 GB
HAProxy VMs*	4	2	16 GB
NDB Agent	8	2	16 GB

* - Applies only to an HA configuration.

For the configuration steps, see [Modifying VM Resource Configurations](#) on page 17.

Table 2: Recommended VM Resource Configurations for up to 800 TMs

Service	vCPU	Cores	Memory
NDB Server	16	2	32 GB
NDB Agent	16	2	32 GB

Table 3: Recommended File Descriptor Limits for up to 800 TMs

File	Parameter	Value
/etc/security/limits.conf	soft nofile	16384
	hard nofile	65536
/etc/systemd/system/era_monitor.service	LimitNOFILE	16384

For information on how to configure file descriptor limits, see [Modifying File Descriptor Limits](#) on page 21.

Table 4: NDB UI's Polling Interval for Operation Summary API Call for up to 800 TMs

Property	Default Value	Value
ui_operation_summary_api_call_interval_millisec	20000	60000

For information on how to modify this property, see [Modifying the NDB UI's Polling Interval for Operation Summary API Call](#) on page 22.

Table 5: Other Recommended Configurations for up to 200 and up to 800 TMs

Configuration	Parameter	Default Value	200 TMs with HA	800 TMs (no HA)
Tomcat Configuration	maxIdle (in context.xml)	85	135	450
	maxTotal/ maxActive (in context.xml)	95	150	500
	maxThreads (in server.xml)	165	NA	500
Repository Database Configuration	max_connections	100 for 200 TMs with HA and 250 for 800 TMs without HA	350	525
NDB Service Configuration	timeout_perform_curation	70	250	250
	priority_perform_curation	30	10	10
	staging_snapshot_curation_interval_min	180	60	60
	era_resiliency_interval_minutes	60	360	360
	era_resiliency_backup_time_machine	true	NA	false
	cleanup_preserve_system_operations_days	30	NA	3
	preserve_system_operations_days	30	NA	7
	operation_cleanup_interval_minutes	1440	NA	360
	skip_and_display_commands_for_log_catchup	false	true	true
Log Catchup Configurations	skip_refresh_metadata :: log_catchup	false	true	true
HAProxy Settings	timeout connect	3s	15s	NA
	timeout check	3s	15s	NA
	“Inter” under “listen ServiceHAHttps” for GET /era/v0.9/ha/leader	2s	15s	NA
	“Inter” under “listen ServiceHAHttp” for GET /era/v0.9/ha/leader	2s	15s	NA

Modifying VM Resource Configurations

This topic describes the steps to modify the VM resource configuration.

About this task

Make the changes (update vCPU, cores per vCPU, and memory) from the Prism UI. Apply the changes for one VM at a time. Updating values from Prism requires the VM to be powered off.

Procedure

1. From the drop-down list of the main menu, select **Administration** and go to **NDB Service**. Three VMs of type 'NDB Repository' appear under **NDB Service VMs**.
2. From among the NDB repository VMs, use the IP address of the VM marked as the leader to establish an SSH connection.
3. On the command prompt of the NDB repository, ensure that no PostgreSQL repository node shows a lag in the Lag column by executing the following command. You can execute this command from any of the repository VMs.

```
patronictl -c /etc/patroni/patroni.yml list
```

4. Change the VM specification of the ASYNC replica (node not having State as Leader or Sync standby) and restart the VM. When the VM is back online, check the patronictl list output to ensure the cluster is back to a healthy state.
5. Change the VM specification of the SYNC replica and restart the VM. Once the VM is back online, check the `patronictl list` output to ensure the cluster is back to a healthy state. At this point, one SYNC replica must be present in the cluster.
6. Perform switchover to ensure minimum downtime:

```
patronictl -c /etc/patroni/patroni.yml switchover
```
7. Once the switchover is complete, ensure the cluster is back to a healthy state. All nodes in the output should have State as running and there should be no lag in the Lag column. There should be one node with Role as Leader and one node with Role as Sync standby.
8. Change the VM specification of the primary VM and restart the VM. Once the VM is back online, ensure that the cluster is back to a healthy state.

Modifying the Tomcat Configuration

This topic describes the steps to modify Tomcat configurations.

About this task

You need to increase the Tomcat configurations to meet the increased demand during NDB upgrade.

Procedure

1. SSH to the NDB server or each of the NDB API servers if NDB HA is enabled.
2. Navigate to the directory: `/home/era/era_base/deploy/apache-tomcat-<version>/conf`.
3. Open the file `context.xml` or `server.xml` in a text editor.
4. Change the values for the properties mentioned in the `<Connector>` or `<Resource>` element.

5. Restart Tomcat service:

- a. Find the process ID of the Tomcat process using the command:

```
ps -aef | grep tomcat
```

- b. kill the process using the process ID fetched from the above step:

```
kill -9 <process_id>
```

- c. A new Tomcat process automatically starts within 30 to 60 seconds. Validate the process using the command:

```
ps -aef | grep tomcat
```

Modifying the Repository Database Configuration

This topic describes the steps to modify the PostgreSQL configuration.

Procedure

1. If the setup has NDB HA enabled:

- a. From the drop-down list of the main menu, select **Administration** and go to **NDB Service**.
- b. From among the **NDB Service VMs**, use the IP address of one of the three VMs with type 'NDB Repository' to establish an SSH connection.
- c. On the command prompt of the NDB repository, use the following commands to edit the configuration:

```
export EDITOR=vi  
patronictl -c /etc/patroni/patroni.yml edit-config
```

- d. Update the `max_connections` value. Save and exit from the vi editor.
- e. Validate and confirm the displayed changes.
- f. Execute the list command:

```
patronictl -c /etc/patroni/patroni.yml list
```

- g. Copy the name of the cluster from the "Cluster" column and replace `<cluster_name>` in the following command and execute:

```
patronictl -c /etc/patroni/patroni.yml restart <cluster_name>
```

- h. Run the list command to confirm that the cluster is healthy. All nodes in the output should have State as running and there should be no lag in the Lag column. There should be one node with Role as Leader and one node with Role as Sync Standby.

```
patronictl -c /etc/patroni/patroni.yml list
```

2. If the setup does not have NDB HA enabled:

- a. SSH to the NDB server.
- b. Open the file `/postgres10/postgresql.conf`.
- c. Update the `max_connections` value. Save the file and exit.

```
patronictl -c /etc/patroni/patroni.yml list
```

- d. Stop the Postgres process using the command:

```
sudo -H -u postgres bash -c /usr/pgsql-10/bin/pg_ctl stop -D /postgres10 -m fast
```

- e. A new process automatically starts within 30 to 60 seconds. Validate this process using the command:

```
ps -aef | grep postgres
```

Modifying Log Copy and Log Catchup Configurations

This topic describes how to modify the log copy and log catchup configurations.

Procedure

1. Run the following command on the leader API server node if NDB HA is enabled or on the NDB server VM if NDB HA is not enabled:

```
era-server -c "config driver_config list name=common_layer_config  
output_file=common_driver.json"
```

2. Open the `common_driver.json` file in a text editor and update the log copy and log catchup properties.
3. Run the following command to update the configuration:

```
era-server -c "config driver_config update name=common_layer_config  
input_file=<file_path>/common_driver.json"
```

4. Reload the configuration by using the command:

```
era-server -c "config reload"
```

Modifying NDB Service Configuration

This task describes the steps to modify the NDB service configuration.

Procedure

1. For each property that needs to be modified, execute the following command on the leader API server node if NDB HA is enabled or on the NDB server VM if NDB HA is not enabled.

```
era-server -c "config era_config update name=<property_name> value=<value>"
```

2. Run the following command to reload the configuration:

```
era-server -c "config reload"
```

Modifying HAProxy Settings

This topic describes the steps to modify HAProxy settings.

Procedure

1. From the drop-down list of the main menu, select **Administration** and go to **NDB Service**.
2. From among the **NDB Service VMs**, use the IP address of the VMs with type 'NDB Repository HA Proxy' to establish an SSH connection as "era" user.
3. On the command prompt of the HAProxy node, execute the following command to open the HAProxy configuration file:

```
sudo vi /etc/haproxy/haproxy.cfg
```

4. Under ServiceHAHttps for GET/era/v0.9/ha/leader, do the following:

- a. Add the following values:

```
timeout connect 15s
timeout check 15s
```

Replace `default-server inter 3s fall 3 rise 2 on-marked-down shutdown-sessions` with `default-server inter 15s fastinter 3s downinter 3s fall 3 rise 2 on-marked-down shutdown-sessions`.

Note: The resultant configuration is:

```
listen ServiceHAHttps
  timeout connect 15s
  timeout check 15s
  bind <ip>:443
  bind <ip>:8443
  option httpchk GET /era/v0.9/ha/leader
  http-check expect status 200
  default-server inter 15s fastinter 3s downinter 3s fall 3 rise 2 on-marked-down shutdown-sessions
    server server_<ip>_443 <ip>:443 maxconn 1000 check check-ssl verify none port 443
    server server_<ip>_443 <ip>:443 maxconn 1000 check check-ssl verify none port 443
    server server_<ip>_443 <ip>:443 maxconn 1000 check check-ssl verify none port 443
```

5. Under ServiceHAHttp for GET/era/v0.9/ha/leader, do the following:

- a. Add the following values:

```
timeout connect 15s
timeout check 15s
```

Replace `default-server inter 3s fall 3 rise 2 on-marked-down shutdown-sessions` with `default-server inter 15s fastinter 3s downinter 3s fall 3 rise 2 on-marked-down shutdown-sessions`.

Note: The resultant configuration is:

```
listen ServiceHAHttp
  timeout connect 15s
  timeout check 15s
  bind <ip>:80
  option httpchk GET /era/v0.9/ha/leader
  http-check expect status 200
  default-server inter 15s fastinter 3s downinter 3s fall 3 rise 2 on-marked-down shutdown-sessions
```

```
server server_1<ip>_443 <ip>:80 maxconn 1000 check check-ssl verify none
port 443
server server_<ip>_80 <ip>:80 maxconn 1000 check check-ssl verify none
port 443
server server_<ip>_80 <ip>:80 maxconn 1000 check check-ssl verify none
port 443
```

6. Repeat steps 2 through 5 on the second HAProxy node.
7. Restart the HAProxy service with the following command on both the HAProxy nodes:

```
sudo systemctl restart haproxy
```

Modifying File Descriptor Limits

This topic describes the steps to modify the file descriptor limits.

Procedure

1. SSH to the NDB server or each of the NDB API servers if NDB HA is enabled.
2. Navigate to the `/etc/security/limits.conf` directory.
3. Open the `limits.conf` file in a text editor.
4. Add the following lines to the file:

```
* soft nofile 16384
* hard nofile 65536
```

5. Navigate to the `/etc/systemd/system/era_monitor.service` directory.
6. Open the `era_monitor.service` file in a text editor.
7. Add the following line to the file:

```
LimitNOFILE=16384
```

8. If you only modified the `era-monitor.service` file, stop and restart the `era-monitor` service. Skip this step if you also modified the `limits.conf` file.

```
sudo systemctl daemon-reload
sudo service era_monitor stop
sudo service era_monitor start
```

9. Identify the process ID of the Tomcat process:

```
ps -ef |grep tomcat
```

10. Kill the Tomcat process on the NDB server VM:

```
kill -9 process_id
```

Replace *process_id* with the process ID for Tomcat. The `era-monitor` service restarts Tomcat after 15-20 seconds.

11. Validate the limit values for the Tomcat process:

```
cat /proc/<process_id>/limits
```

Note: Steps 8 through 9 are valid if you modified only the `era_monitor.service` file. If you modified the `limits.conf` file also, restart the VM for the changes to take effect globally.

Modifying the NDB UI's Polling Interval for Operation Summary API Call

This topic describes the steps to modify the NDB UI's polling interval for operation summary API call.

Procedure

1. Access the NDB configuration page at the following link:

```
https://<ndb_ip_address>/#/eraconfig
```

Replace `<ndb_ip_address>` with the IP address of the NDB server.

2. Modify the value of the `ui_operation_summary_api_call_interval_milliseconds` property.
3. Click **Update**.

NDB with Nutanix Objects

NDB supports Nutanix Objects for storing log backups.

NDB supports restoring and cloning of databases to a given point in time. For these functionalities to work, NDB needs to store the log backups of the databases as defined by the SLAs. By default, NDB attaches one volume group (VG) per database to the NDB agent (or NDB server in the case of single Nutanix cluster) and stores the log backups in the VGs. These VGs act as log drives for the database.

However, this method of storing log backups can face multiple bottlenecks while scaling beyond 200 time machines on a single Nutanix cluster. To enable seamless scaling of databases, you can integrate NDB with Nutanix Objects for storing log backups. NDB supports Objects for all database engines including Oracle, SQL Server, PostgreSQL, MongoDB, MySQL, and MariaDB.

Note:

- NDB does not support the migration of databases and time machines with VG-based log drives already registered with NDB to Nutanix Objects-based log drives. To convert a VG-based database to an Objects-based database, deregister the database, register an Object Store with NDB, and register the database again.
- NDB curates the logs from the Object Store depending on the SLA of the associated time machine. If you set the retention period to 30 days, NDB automatically curates the logs after 30 days. You need not configure a WORM policy on Objects for the curation of database logs.

Registering Nutanix Objects for Storing NDB Log Backups

This task describes the steps to enable Nutanix Objects for storing log backups from NDB.

About this task

To use Nutanix Objects with NDB, implement an Object Store in a Nutanix cluster where NDB hosts its databases or in a separate Nutanix cluster. NDB connects to this Object Store and consumes it.

Before you begin

Create an Object Store.

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.

2. Generate the input file for registration of Object Store.

```
era > storage_resource add generate_input_file output_file=storage_resource.json
```

Replace *storage_resource.json* with the name of the output file.

The contents of the output JSON file are:

```
{
  "name": "",
  "description": "",
  "type": "NUTANIX_OBJECT_STORE",
  "info": {},
  "fqdn": "",
  "accessKey": "",
  "secretKey": ""
}
```

3. Exit the *era* CLI and modify the content of the JSON file with the details of the Object Store that you created. The *name*, *fqdn*, *secretKey*, and *accessKey* parameters are necessary for Object Store registration with NDB.

```
era > exit
[era@localhost~]$ vi storage_resources.json
```

4. Enter the *era* CLI again and add the storage resources information.

```
era > storage_resource add input_file=storage_resource.json
```

Replace *storage_resource.json* with the name of the JSON file.

5. List the storage resources registered with NDB.

```
era > storage_resource list
```

The following sample information displays:

Storage Resources Information :

```
+-----+-----+-----+-----+
+-----+-----+-----+-----+
|          ID          |          Name          | Description |
| Type      | Info                   |                  |
|-----|-----|-----|
| 90103ce0-fbc8-479d-b93f-94734f05c35a | EraCluster |
| NUTANIX_CL | {}                  |
| USTER      |                      |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 7ba151c7-a9f0-4526-b957-2b89b6e45761 | Storage Resource 1 | SR2        |
| NUTANIX_OB | {}                  |
| JECT_STORE |                      |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
Total records: 2
```

Updating a Storage Resource Registered with NDB

This task describes the steps to update a storage resource registered with NDB.

About this task

You must update the storage resource if the Nutanix Objects credentials used while on-boarding the storage resource with NDB expires or becomes invalid. This is a common requirement as access keys expire frequently.

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.
2. Generate the input file for modifying the storage resource.

```
era > storage_resource update generate_input_file output_file=storage_resource.json
```

Replace `storage_resource.json` with the name of the output JSON file.

The sample contents of the output JSON file are:

```
{
  "id": "69521aab-560b-4378-a7b9-5a51da16d93e",
  "name": "Storage Resource 1",
  "description": "SR2",
  "type": "NUTANIX_OBJECT_STORE",
  "ownerId": "eac70dbf-22fb-462b-9498-949796ca1f73",
  "info": {}
}
```

3. Exit the `era` CLI and update the content of the JSON file with the new details of the Object Store. You can update the `name`, `fqdn`, `secretKey`, and `accessKey` parameters, if needed.

```
era > exit
[era@localhost~]$ vi storage_resources.json
```

4. Enter the `era` CLI again and update the storage resources information.

```
era > storage_resource update input_file=storage_resource.json
```

Replace `storage_resource.json` with the name of the JSON file.

5. List the updated storage resources registered with NDB.

```
era > storage_resource list
```

The following sample information appears:

```

Storage Resources Information :
+-----+-----+-----+
+-----+-----+-----+
|                                     |                                     |                                     |
|      ID                               Name                               Description                               |
| Type |                               Info |                               |                                     |
|                                     |                                     |                                     |
+-----+-----+-----+
+-----+-----+-----+
| 69521aab-560b-4378-a7b9-5a51da16d93e | Storage Resource 2 | SR2 |
| NUTANIX_OB | {} | |
| JECT_STORE | | |
+-----+-----+-----+
+-----+-----+-----+
Total records: 1

```

Removing a Storage Resource Registered with NDB

This task describes the steps to remove a storage resource registered with NDB.

Before you begin

A Nutanix Objects storage resource cannot be removed from NDB if there is a time machine using the storage resource. For removing the storage resource from NDB, remove all the time machines using the storage resource from NDB.

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.
2. Remove the storage resource.

```
era > storage_resource remove id=resource_id
```

Replace *resource_id* with the ID of the storage resource.

The following response appears indicating that you removed the storage resource successfully.

```
{'status': 'success', 'message': 'success', 'info': None, 'errorCode': 0}
```

Provisioning Databases that Use Nutanix Objects

This task describes the CLI steps to provision a database that uses Nutanix Objects for storing log backups.

About this task

Perform the following steps to provision a database that uses Nutanix Objects for storing log backups. For the API steps, see [KB 13612](#).

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.
2. Create the input file for database provisioning:

```
era > database provision engine=database_engine provision_instance  
generate_input_file single_instance create_dbserver=create_dbserver  
software=software_profile compute=compute_profile network=network_profile  
nx_cluster_id=nutanix_cluster_id db_parameter=db_parameter_profile  
output_file=file_name.json
```

Replace *database_engine* with the database engine name.

Replace *create_dbserver* with `true` to create a new database server. Replace with `false` to use a pre-provisioned database server.

Replace *software_profile* with the name software profile.

Replace *compute_profile* with the name compute profile.

Replace *network_profile* with the name network profile.

Replace *nutanix_cluster_id* with the Nutanix cluster ID.

Replace *db_parameter_profile* with the name of the database parameter profile.

Replace *file_name.json* with the name of the input file.

3. Provision a database with a given SLA and specify the `storage_resource_id` of the Nutanix Object Store you want to use.

```
era > database provision engine=database_engine provision_instance name=database_name  
input_file=input_file_name sla=sla_name storage_resource_id=resource_id  
force_vg_based_log_drive=vg_based_log_drive
```

Replace *database_engine* with the database engine name.

Replace *database_name* with the name of the database.

Replace *input_file_name* with the name of the input JSON file from the previous step.

Replace *sla_name* with the name of the SLA for the time machine.

Replace *resource_id* with the ID of the storage resource.

Replace *vg_based_log_drive* with `true` or `false`.

This CLI command will submit a provisioning operation to NDB. You can view the status of the operation from the **Operations** menu in the UI.

Note:

- The parameters `storage_resource_id` and `force_vg_based_log_drive` are mutually exclusive, and you can only provide one in the CLI input. If you provide both the parameters, the NDB server throws an exception.
- If you register an Object Store with NDB server and do not specify a `storage_resource_id`, NDB uses the first reachable Object Store for storing log backups of the newly provisioned database.
To overrule this behaviour, specify the `force_vg_based_log_drive` parameter as `true` and do not provide a `storage_resource_id` in the CLI input. This ensures that a VG-based log drive is created for the newly provisioned database.
- If you did not register an Object Store with NDB, a VG-based log drive is created for the newly provisioned database.
- If you submit the provisioning operation from the UI and there are Object Store Services registered with the NDB Server, then NDB chooses the first Object Store that passes the connectivity checks and creates a log drive on it.

Registering Databases that Use Nutanix Objects

This task describes the CLI steps to register a database that uses Nutanix Objects for storing log backups.

About this task

Perform the following steps to register a database that uses Nutanix Objects for storing log backups. For the API steps, see [KB 13612](#).

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.
2. Create the input file for database registration.

```
era > database generate_input_file type=database_engine output_file=file_name.json
```

Replace *database_engine* with the database engine name.

Replace *file_name.json* with the name of the input JSON file.

3. Register a database with a given SLA and specify the `storage_resource_id` of the Nutanix Object Store you want to use.

```
era > database add clustered=clustered_database name=database_name  
nx_cluster_id=nutanix_cluster_id input_file=input_file_name  
sla_name=sla_name schedule_name=schedule_name storage_resource_id=resource_id  
force_vg_based_log_drive=force_vg_based_log_drive
```

Replace `clustered_database` with `true` or `false` to specify if you are registering a clustered database.

Replace `name` with the name of the database.

Replace `nutanix_cluster_id` with the Nutanix cluster ID.

Replace `input_file_name` with the name of the input JSON file from the previous step.

Replace `sla_name` with the name of the SLA for the time machine.

Replace `schedule_name` with the name of the schedule.

Replace `resource_id` with the ID of the storage resource.

Replace `vg_based_log_drive` with `true` or `false`.

This CLI command will submit a registration operation to NDB. You can view the status of the operation from the **Operations** menu in the UI.

Note:

- The parameters `storage_resource_id` and `force_vg_based_log_drive` are mutually exclusive, and you can only provide one in the CLI input. If you provide both the parameters, the NDB server throws an exception.
- If you register an Object Store with NDB server and do not specify a `storage_resource_id`, then NDB uses the first reachable Object Store for storing log backups of the newly registered database.
To overrule this behaviour, specify the `force_vg_based_log_drive` parameter as `true` and do not provide a `storage_resource_id` in the CLI input. This ensures that a VG-based log drive is created for the newly registered database.
- If you did not register an Object Store with NDB, a VG-based log drive is created for the newly registered database.
- If you submit the registration operation from the UI and there are Object Store Services registered with the NDB server, then NDB chooses the first Object Store that passes the connectivity checks and creates a log drive on it.

Configuring a Data Access Management Policy for a Time Machine

This task describes the CLI steps to configure a data access management (DAM) policy for a time machine.

About this task

Perform the following steps to configure a DAM policy for a time machine.

Procedure

1. SSH to the NDB server and enter the `era` command to enter the CLI.

2. Create a DAM Policy for a time machine with a given SLA and specify the `storage_resource_id` of the Nutanix Objects store.

```
era > time_machine nx_cluster_association add id=time_machine_id  
      nx_cluster_id=nutanix_cluster_id sla_name=sla_name type=database_engine  
      storage_resource_id=resource_id force_vg_based_log_drive=force_vg_based_log_drive
```

Replace *time_machine_id* with the time machine ID.

Replace *nutanix_cluster_id* with the Nutanix cluster ID.

Replace *sla_name* with the name of the SLA for the time machine.

Replace *database_engine* with the name of the database engine.

Replace *resource_id* with the ID of the storage resource.

Replace *force_vg_based_log_drive* with `true` or `false`.

Note:

- The parameters `storage_resource_id` and `force_vg_based_log_drive` are mutually exclusive, and you can only provide one in the CLI input. If you provide both the parameters, the NDB server throws an exception.
- If you register an Object Store with NDB server and do not specify a `storage_resource_id`, NDB uses the first reachable Object Store for the DAM policy.

To overrule this behaviour, specify the `force_vg_based_log_drive` parameter as `true` and do not provide a `storage_resource_id` in the CLI input. This ensures that a VG-based log drive is created for the DAM policy.

- If you did not register an Object Store with NDB, a VG-based log drive is created for the DAM policy.
- If you submit the DAM policy creation from the UI and there are Object Store Services registered with the NDB server, NDB chooses the first Object Store that passes the connectivity checks to create the log drive.

NDB INSTALLATION

NDB is a virtual appliance that is installed on either AHV or ESXi. To install NDB on an AHV cluster, use the Prism Element web console. To install NDB on an ESXi cluster, use vCenter Server.

For information on the supported AOS, AHV, and vSphere software versions for NDB and NDB multi cluster, see *NDB Software Compatibility with Nutanix and VMware Products* in *Nutanix Database Service Release Notes*.

NDB Network Requirements

This section provides information on the network requirements for NDB server, database engine, database server VM, and NDB multi-cluster.

Ensure that you meet the following network requirements on your cluster.

Table 6: NDB Server Port Requirements

Component	Description
REST APIs	The NDB database agent, running in each DB Server VM, communicates with the NDB server using port 443 and protocol HTTPS and communicates with the Nutanix cluster on port 9440 and protocol HTTPS for REST API calls.
Telemetry data	The NDB server sends periodic telemetry data to the central Nutanix repository. The NDB server posts data to the Pulse pipeline that is reachable at insights.nutanix.com on port 443.
Access to Nutanix portal	For checking and downloading latest NDB upgrades, the NDB server must have connectivity to the Nutanix portal.
Access to NTP time servers	The NDB server periodically synchronizes with time servers to update the server clock. The NDB server must be able to connect to an NTP server (for example, *.centos.pool.ntp.org) on port 123. For an NDB server that has no internet connectivity, you must manually perform time synchronization using the initial setup guide or at a later time through era-server CLI and ensure that the server time is correct.
iSCSI connections	Ports 3205 and 3260 must be open on the NDB server and Controller VMs (CVMs) to facilitate iSCSI connections.

Note: You must configure DSIP on the Prism Element to create iSCSI connections.

Component	Description
Nutanix cluster	- The NDB server, NDB agent, and NDB database agent communicate with the registered Nutanix clusters over port 9440 (Prism REST API calls). If any of the disks are part of the volume group that is attached through iSCSI, the NDB server uses port 3260 and 3205. Ports 3205 and 3260 must be open to both Data Services IP (DSIP) and Nutanix cluster virtual IP, and on all the Controller VMs (CVMs) to facilitate iSCSI connections. - Prism Element with cluster administrator access (AHV and ESXi) - ESXi only - ensure that you have registered vCenter with Prism Element. Note: Data replication between the Nutanix clusters needs IP ICMP, TCP ports 2009 (Stargate), and 2020 (Cerebro) to be open for every CVM IP and Cluster VIP that is between the source and the destination cluster.

For a list of required NDB ports, see the respective port reference guide. The port reference guides provide detailed port information for Nutanix products and services, including port sources and destinations, service descriptions, directionality, and protocol requirements.

This table provides links to the port reference guides.

Table 7: NDB Port Requirements

Database server VM	NDB (Database Server VM)
Multi-cluster network	NDB (Multi-cluster Network)
Oracle	NDB (Oracle)
SQL Server	NDB (SQL Server)
PostgreSQL	NDB (PostgreSQL)
MySQL	NDB (MySQL)
MariaDB	NDB (MariaDB)
NDB control plane (HA)	NDB control plane (HA)

Note: NDB on Nutanix Cloud Clusters on AWS (NC2) does not have any specific port requirements.

NDB Network Requirement Diagrams

NDB Network Requirement Diagrams for Windows




Figure 2: Windows Database Server VM Network Requirements for AHV




Figure 3: Windows Database Server VM Network Requirements for ESXi




Figure 4: Multi-Cluster Windows Database Server VM Network Requirements for AHV




Figure 5: Multi-Cluster Windows Database Server VM Network Requirements for ESXi

NDB Network Requirement Diagrams for Linux




Figure 6: Linux Database Server VM Network Requirements for AHV




Figure 7: Linux Database Server VM Network Requirements for ESXi




Figure 8: Multi-Cluster Linux Database Server VM Network Requirements for AHV




Figure 9: Multi-Cluster Linux Database Server VM Network Requirements for ESXi

Installing NDB on AHV

About this task

You can install NDB on an AHV cluster by using the Prism Element web console.

Note: NDB requires a Prism Element account with cluster administrator access.

Perform the following procedure to install NDB on an AHV cluster.

Procedure

1. Download the QCOW2 image file of NDB from the Nutanix Support portal.
2. Log on to the Prism Element web console.
3. Configure an image. For the detailed steps, see [Configuring Images](#).
4. Create a VM for NDB. Perform Steps 1 through 8 as stated in [Creating a VM \(AHV\)](#).

Note: For information on the quantity of vCPUs, cores, and memory to allocate to the VM, see [NDB Control Plane Configuration and Scalability](#) on page 13.

5. To configure the network by assigning a static IP address to the NDB VM, do the following in Step 9 of [Creating a VM \(AHV\)](#):
 - a. Select the **Custom Script** check box.
 - b. In the **Type or paste** script text box, paste the following script.

```
#cloud-config  
runcmd:
```

```
- configure_static_ip ip=ip_address gateway=gateway_address netmask=ip_netmask  
nameserver=ns1
```

All parameters except the nameserver parameter are mandatory.

Replace the variables with the following:

- Replace `ip_address` with the static IP address you want to assign to the NDB VM.
- Replace `gateway_address` with the gateway IP address for the NDB VM.
- Replace `ip_netmask` with the netmask of the network.
- Replace `ns1` with the URL or IP address of the name server.

For example:

☒ Custom Script

Provide a Cloudinit or Sysprep script to customize this VM.

☐ ADSF Path

Start typing for suggestions

☐ Upload A File

Choose File

No file chosen

☒ Type Or Paste Script

```
#cloud-config  
runcmd:  
- configure_static_ip ip=ip_address gateway=gateway_address  
netmask=ip_netmask nameserver=ns1
```

Figure 10: Assign a Static IP Address to the NDB VM

Alternatively, you can assign a static IP address to the NDB VM after you have installed the VM. For instructions, see [Assigning a Static IP Address to the NDB VM by Using the Console](#).

Note: Nutanix recommends assigning a static IP address to the NDB VM as a best practice and does not recommend IP address management services.

6. Perform Steps 10 through 12 as stated in [Creating a VM \(AHV\)](#) to complete the creation of the VM.
7. In the **Table** view of the VM dashboard, select the NDB VM and click **Power On** to start the VM.
8. Determine the IP address assigned to the NDB VM from the **IP Addresses** field of the **Table** view in the Prism Element web console.

If you assign a static IP address to the NDB VM on a VLAN that has a DHCP server, Prism Element first assigns an IP address to the NDB VM by using DHCP. Wait for a few minutes and refresh the Prism Element page to verify that the static IP address you specified has been assigned to the VM.

Installing NDB on ESXi

About this task

You can install NDB on an ESXi cluster by using the vCenter Server.

Perform the following procedure to install NDB on an ESXi cluster by using vCenter Server.

Procedure

1. Download the NDB OVA file from the Nutanix Support portal.
2. Log on to vCenter Server with your administrator credentials.
3. To deploy an OVF template in vCenter Server, follow the [VMware documentation](#). For more information on NDB scalability requirements, see [NDB Control Plane Configuration and Scalability](#) on page 13.
4. Turn on the NDB VM.
5. To determine the IP address that is assigned to the NDB VM, do one of the following.
 - » In vCenter Server, check the **Summary** tab of the VM.
 - » In the Prism Element web console, check the **IP Addresses** field of the **Table** view.
6. To assign a static IP address to the NDB VM, see [Assigning a Static IP Address to the NDB VM by Using the Console](#) on page 38.

If you assign a static IP address to the NDB VM, Prism Element uses DHCP to complete the assignment. Wait for 1 or 2 minutes and refresh the Prism Element page to verify that NDB has assigned the static IP address to the VM.

Note: Nutanix recommends assigning a static IP address as a best practice.

NDB INITIAL CONFIGURATION

The NDB VM must be deployed and configured in a Nutanix cluster before it can start managing database VMs.

Assigning a Static IP Address to the NDB VM by Using the Console

About this task

For AHV, either perform the following procedure or assign a static IP to the VM during installation. For instructions about assigning a static IP address to the NDB VM during installation, see [Installing NDB on AHV](#) on page 34.

For ESXi, perform the following procedure if you want to assign a static IP address to the NDB VM.

Procedure

1. Log on to the Prism Element web console.
2. Go to the VM dashboard, select the NDB VM, and click **Launch Console**.
3. Use the following credentials to log on to NDB.

- Username: era
- Password: Nutanix.1

4. Launch the NDB server prompt.

```
$ era-server
```

5. Assign a static IP address to NDB.

```
era-server > era_server set ip=ip_address gateway=gateway_address netmask=ip_netmask  
nameserver=ns1 ntp_server=ntp1
```

All parameters except `nameserver` and `ntp_server` are mandatory.

Replace the variables with the following:

- Replace *ip_address* with the static IP address you want to assign to the NDB VM.
- Replace *gateway_address* with the gateway IP address for the NDB VM.
- Replace *ip_netmask* with the netmask of the network.
- Replace *ns1* with the URL or IP address of the name server.
- Replace *ntp1* with the URL or IP address of the NTP server.

Alternatively, instead of typing the command, you can press the tab key to prompt you to complete the command.

For more information on changing the IP addresses of the NDB server VM and the NDB agent VM, see [KB 12164](#).

Configuring an SSL Certificate

NDB supports custom SSL certificate-based authentication. You can configure a custom SSL certificate using CLI. Configuring a custom SSL certificate also helps resolve any invalid certificate issues while accessing the NDB user interface. Nutanix recommends replacing the default self-signed certificate with a Certificate Authority (CA) signed certificate.

Before you begin

If NDB high availability is enabled, the custom certificate should contain VIP, VIP FQDN, and IP addresses and FQDN of all three API servers in the SAN extension section.

About this task

Perform the following procedure to configure a custom SSL certificate using the CLI.

Procedure

1. Launch the NDB server CLI.

```
$ era-server
```

2. Configure the certificate by using the command:

```
era-server > security ssl add_custom certificate_file=file_path private_key=file_path  
ca_certificate=file_path
```

Replace *file_path* with the security certificate file location.

Note: If NDB high availability is enabled, update the SSL certificate on all the API server VMs by executing the above CLI commands on each API server VM individually.

Logging on to NDB

About this task

Perform the following procedure to log on to NDB.

Note: Root access is disabled on the NDB VM.

Procedure

1. Open a web browser and type the following in the address bar:

https://ndb_vm_ip_address

Replace *ndb_vm_ip_address* with the IP address assigned to the NDB VM.

2. If you are signing in for the first time, do the following:
 - a. Read the Nutanix End User License Agreement (EULA) agreement, click the **I have read and agree to terms and conditions** option, and click **Continue**.
 - b. In the **Nutanix Customer Experience Program** screen, click **OK**.
By participating in the Nutanix customer experience program, Nutanix collects non-identifying information for product improvement.
 - c. At the login screen, set a password for the administrator user in the **Enter new password** and **Re-enter new password** fields and click **Set Password**.
3. In the login screen, type the password of the "admin" user and press **Enter**.
If you are signing in for the first time, you are directed to the **Welcome to NDB** wizard. Otherwise you are directed to the NDB dashboard.

Checking NDB Version

The user menu on the right of the main menu provides the option to know the current NDB version you are running.

About this task

Perform the following procedure to check the current NDB version.

Procedure

1. Click the user menu drop-down at the top-right corner of the screen.
2. Select **About NDB**.




Figure 11: About NDB Page

The **About NDB** page appears displaying the NDB version number and build information.

Changing the Language Settings

You can change the NDB user interface language settings from English (default) to other languages. This option allows you to view the user interface and provide user inputs for the available fields in the selected language.

About this task

Perform the following procedure to change the language settings of the NDB user interface.

Procedure

1. Select **Profile** from the user menu at the top-right corner of the screen.
The **Profile** page appears.
2. Select the required language from the **Language** drop-down list.

3. Click **Save** to complete.

You must logout and login again for the changes to take effect. The NDB user interface language changes to selected language.

Note: The language settings are local to a web browser. When you sign in to NDB from another browser, the language settings default to English.

GET STARTED WITH NDB

After the NDB setup is complete, the **Welcome to NDB** wizard guides you to complete the initial tasks such as registering a Nutanix cluster with NDB, configuring network segmentation, and creating a sample network profile. As part of network segmentation you can add one or more VLANs to segment the NDB server network traffic for each of the following communications:

- Prism Element
- vCenter Server (for ESXi)
- Prism iSCSI data service
- Database server VMs

Note: Network segmentation is only supported for the Oracle database engine.

This wizard then directs you to the **Getting Started** page, which guides you in provisioning a database or registering an existing database to create a time machine.

For information related to NDB service management, role-based access control (RBAC), multi-cluster management, network management, and so on, see [NDB Administration](#) on page 57.

Welcome to NDB Wizard

About this task

Perform the following procedure to complete the wizard.

Procedure

1. You complete the wizard in the following steps.

- a. **NDB's Cluster**
- b. (Optional) **Network Segmentation**
- c. **Services**
- d. **Storage Container**
- e. **Network Profile**
- f. **Setup NDB**

2. In the **NDB's Cluster** step, do the following in the indicated fields:
 - a. **Name.** Type a name for the Nutanix cluster as you want the name to appear in NDB.
 - b. **Description.** Type a description for the Nutanix cluster.
 - c. **Address.** Type the IP address of the Prism Element web console of the Nutanix cluster.
 - d. **Prism Element Administrator.** Type the username of the Prism Element user account with which you want NDB to access the Nutanix cluster.

Note:

- You must configure DSIP on the Prism Element to create iSCSI connections.
- Nutanix does not recommend using the default administrative account for NDB operations. Instead, Nutanix recommends using a separate Prism Element user account with Nutanix cluster administrative privileges as the NDB service account.

- e. **Password.** Type the password of the Prism Element user account.
- f. Click **Next**.

NDB

Welcome to NDB

1 NDB's Cluster 2 Network Segmentation 3 Services 4 Storage Container 5 Network Profile 6 Setup NDB

Register the Cluster with NDB

Name

Description

Address

Cluster Virtual IP Address

Prism Element Administrator

admin

Password

Next

Figure 12: NDB Wizard: NDB's Cluster Step

3. Do the following in the **Network Segmentation** step if you want to segment the network traffic from the NDB server.

Caution: You cannot enable high availability for NDB if you configure network segmentation for NDB.

Note:

- Network segmentation is only supported for static VLANs.
- Network segmentation is an advanced optional configuration; you can skip it if cluster segmentation is not configured in Prism Element.
- A network interface is attached to the NDB server VM for every unique VLAN that you configure.

- a. Select the VLAN access type for which you want to configure network segmentation.

- **Prism Element.** Displays the existing VLAN that is attached to the new NDB server. The IP address of this VLAN is used for all NDB-Prism communications.

Note:

- By default, all the network traffic is directed to the Prism Element network.
- The network traffic of the VLAN access types that do not have a network segmentation configured is directed to Prism Element.
- This network acts as a default gateway for the NDB server and cannot be modified.

This NIC cannot be removed once the NDB server setup is completed.

- **Prism iSCSI Data Service.** Select this VLAN access type to configure a VLAN that the NDB server can use to make connection requests to the iSCSI data service IP addresses.

Note: NDB uses one of the segmented iSCSI data service IP addresses if you have configured multiple segmented data service IP addresses in Prism Element.

- **DBServer Access from NDB server.** Select this VLAN access type to configure a VLAN that is used for communications between the NDB server VM and the database server VM on the NDB server cluster.
- (Only for an ESXi cluster) **vCenter.** Select this VLAN access type to configure a VLAN that the NDB server can use for vCenter API connections.

- b. Click **Next**.

Note: VLAN configurations made in the **Network Segmentation** step cannot be reverted.

NDB
admin

Welcome to NDB

1 NDB's Cluster
2 Network Segmentation
3 Services
4 Storage Container
5 Network Profile
6 Setup NDB

Select a vLAN that NDB should use for communication for each of the following services.
This is an optional setting for Advanced Network Segmentation. A network interface will be attached to NDBServer VM for each of the Distinct vLAN Names provided in the below section.

vLAN Access Type	vLAN	Configuration		
		Gateway	Subnet Mask	IP Address
☒ Prism Element				
☐ Prism iSCSI Data Service				
☐ DBServer Access from NDB Server				
☐ vCenter				

Skip
Back
Next



Figure 13: NDB Wizard: Network Segmentation Step

4. In the **Services** step, do the following in the indicated fields:

- Edit the **DNS Servers** and **NTP Servers** IP addresses according to your requirements.

The list contains the DNS servers and NTP servers configured in the NDB server and the cluster (Prism Element) registered in the NDB server. If no DNS servers are configured in the NDB server, then the list

only shows the DNS servers pulled from the cluster. You can add or remove IP addresses according to your requirements, and the changes only apply to the NDB server.

b. Configure the **SMTP Server** to allow NDB to send alert notifications through email.

- Type the **IP address** and **Port** number for the SMTP server. For more information about configuring the SMTP server, see [Sending Alert Notifications to an Email Recipient](#) on page 163.

Note: SMTP server configuration details are pulled from the cluster and auto-populated in the **Services** screen. You can edit the details and save your changes.

- Select the cryptographic protocol type to use for authentication and data encryption when sending alert notifications through email.
- Type the **Username**, **Password**, and **Sender's email** address when configuring the SMTP server.

c. Only for ESXi - Type your vCenter username and password.

The fields to enter the vCenter credentials appear if the cluster runs on ESXi hypervisor.

d. **NDB server's OS Time Zone.** Select a time zone for the NDB server.

e. Only for ESXi - Enter the **vCenter Username** and **vCenter Password**.

f. Click **Next**.

The screenshot displays the 'Services' step of the NDB Wizard. The interface is divided into three main sections for configuration:

- DNS Servers:** Includes a text area for 'Addresses' (comma-separated values), a 'Hypervisor' section with 'Type: ESX', and fields for 'vCenter IP Address', 'vCenter Username', and 'vCenter Password'.
- NTP Servers:** Includes a text area for 'Addresses', a 'Time Zones' section with 'From: Asia/Calcutta', and a dropdown for 'NDB Server's OS Time Zone' (currently set to 'US/Pacific').
- SMTP Server:** Includes a text field for 'SMTP Server Address', a 'Port' field, a 'Type' section with radio buttons for 'None', 'SSL' (selected), and 'TLS', and fields for 'Username', 'Password', and 'Sender's Email'.

At the bottom right, there are 'Back' and 'Next' buttons. The top of the wizard shows a progress bar with steps: 1. NDB's Cluster, 2. Network Segmentation, 3. Services (current), 4. Storage Container, 5. Network Profile, and 6. Setup NDB. The top left shows 'NDB' and the top right shows 'admin'.

Figure 14: NDB Wizard: Services Step

5. In the **Storage Container** step, select the storage container for NDB to provision a new database VM. Click **Next**.




Figure 15: NDB Wizard: Storage Container Step

6. In the **Network Profile** step, do the following in the indicated fields if you want to create a sample network profile. Select the **Create Profile Later** option to skip this step.

a. Select a **VLAN** from the drop-down list of available VLANs.

This VLAN becomes the default network profile that NDB uses when provisioning new database server VMs.

Note: A VLAN must be created on your Nutanix cluster for your database environment.

b. **Manage IP Address Pool in NDB.** Select this option to allow NDB to assign static IP addresses to your database server VMs.

If you select this option, NDB assigns a static IP address from this IP address pool to your newly provisioned database server VMs that use the network profile.

c. Type the **Gateway, Subnet Mask, Primary DNS, Secondary DNS, DNS Domain**, and the **Initial Range of IP Addresses for New Database Servers**.

d. Click **Next**.




Figure 16: NDB Wizard: Network Profile Step

7. In the **Setup** step, click **Get Started**.

The **Getting Started** page describes how to register and provision databases in NDB. You can also open the main menu and start using the product.

NDB Profiles

NDB profiles simplify the database provisioning process. Profiles enable you to create templates that define database software, networking, compute, and database parameters that you can use to provision databases or database server VMs on NDB.

NDB provides the following profiles to get you started with provisioning your database or database server VMs.

Software Profile

Contains an image of the database and operating system generated from the registered database server VM. NDB uses images to provision the database server VM. For PostgreSQL, MariaDB, and MySQL, a sample software profile is available that contains the PostgreSQL, MariaDB, and MySQL

software. For more information about supported built-in profiles, see *Nutanix Database Service Release Notes*.

For Oracle and SQL Server, you must create a software profile. The software profiles for these engines are created using the previously registered database server VM that is running the type of database needed with NDB.

Note: Before creating a software profile, you must register a database server VM for use in the software profile.

Compute Profile

Specifies the number of vCPUs, cores per vCPU, and the amount of memory for the database server VM.

A sample compute profile **DEFAULT_OOB_COMPUTE** is available for all types of database server VMs.

Network Profile

Specifies the VLAN that you want the new database server VM to use.

Note: Network Segmentation feature for MongoDB is disabled by default. To enable this feature reach out to Nutanix Support.

Database Parameter Profile

Lets you customize database parameters for the new database instance.

Sample profiles for Oracle (DEFAULT_ORACLE_PARAMS), PostgreSQL (DEFAULT_POSTGRES_PARAMS), SQL Server (DEFAULT_SQLSERVER_PARAMS), MariaDB (DEFAULT_MARIADB_PARAMS), and MySQL (DEFAULT_MYSQL_PARAMS) are available for use.

Note: If you do not want to use the sample profiles available, you can either modify the sample profiles or create your own profiles in NDB.

Windows Domain Profile

Contains Windows domain information such as domain name (FQDN), domain user, and password. This profile also includes information related to the SQL service startup account, NDB worker service account, and Active Directory such as the Organizational Unit (OU) paths for clusters and database server VMs.

Getting Started with Provisioning Databases

After the NDB setup is complete, you can provision a database or register a database to create a time machine. The **Getting Started** page guides you through the provisioning of your first database. You can use NDB to provision database server VMs and databases on the registered Nutanix cluster, or you can register an existing source database running on the Nutanix cluster.

About this task

You can select **Getting Started** from the drop-down list of the main menu to open the **Getting Started** page.

In this example, you can provision a PostgreSQL instance. The instance contains multiple databases.

Procedure

1. From the drop-down list of the main menu, select **Getting Started** and then select the **PostgreSQL** engine.




Figure 17: Getting Started Menu

2. In the **Provision a PostgreSQL Database** page, click **Provision a Database**.




Figure 18: Provision a PostgreSQL Instance

The **Provision a PostgreSQL Database** window appears.

3. You provision an instance in the following steps:
 - a. **Database Server VM**
 - b. **Instance**
 - c. **Time Machine**
4. In the **Database Server VM** step, select **Create New Server**.
5. Select the following built-in sample profiles available in NDB.
 - a. **Software Profile**. Select the **POSTGRES_10.4_OOB** profile from the drop-down list.
 - b. **Compute Profile**. Select **DEFAULT_OOB_COMPUTE** from the drop-down list.

By using the sample profile, the PostgreSQL database server VM is created with the following compute values:

 - vCPUs: 2
 - Cores per vCPU: 8
 - Memory: 32 GB
 - c. **Network Profile**. Select **DEFAULT_OOB_POSTGRESQL_NETWORK** from the drop-down list.

The **DEFAULT_OOB_POSTGRESQL_NETWORK** profile includes the VLAN you selected when you created this profile in the **Welcome to NDB** wizard.

Provision a PostgreSQL Instance

1 Database Server VM 2 Instance 3 Time Machine

Database Server VM

☒ Create New Server
☐ Use Registered Server

New Database Server VM

Database Server VM Name *i*

Description

POSTGRES_10.4_OOB

Software Profile

POSTGRES_10.4_OOB *i*

Version: POSTGRE... [Update Version](#)

Compute Profile

DEFAULT_OOB_COMPUTE *i*

Network Profile

DEFAULT_OOB_POSTGRESQL_NETWORK *i*

Database Server VM's Time Zone

Asia/Calcutta

SSH Public Key for Node Access *i*

☒ File
[Upload File](#) Select your public key.
☐ Text

i About profiles in Era 1.3

Cancel Next

Figure 19: Select Built-in Profiles

6. Type the details for the instance including name, description, password, and name of the initial database. Select the **DEFAULT_POSTGRES_PARAMS** profile from the **Database Parameter Profile** drop-down list.

Provision a PostgreSQL Instance

?

×

1 Database Server VM

2 Instance

3 Time Machine

PostgreSQL Instance Name

postgres

Description

Listener Port

−

5432

+

Size (GiB)

−

200

+

Name of Initial Database

test_postgre

Database Parameter Profile

DEFAULT_POSTGRES_PARAM

POSTGRES password

.....

Pre-Post Commands ^

Pre-Create Command

Full OS command that needs to be executed before creating the database. For example: /usr/bin/python /var/scripts/provision.py

Post-Create Command

Full OS command that needs to be executed after creating the database. For example: /usr/bin/python /var/scripts/provision.py

About profiles in Era 1.3

Back

Next

Figure 20: Database Details

7. Type a name for the time machine, select an SLA (duration for which NDB retains the snapshots), and specify the schedule for taking snapshots and creating copy transaction logs.

NDB has the following built-in SLAs.

Name ▾	Databases ▾	Continuous ▾	Daily ▾	Weekly ▾	Monthly ▾	Quarterly ▾
☐ DEFAULT_OOB_GOLD_SLA	0	30 days	90 days	16 weeks	12 Months	35 Quarters
☐ DEFAULT_OOB_SILVER_SLA	0	14 days	60 days	12 weeks	12 Months	0 Quarters
☐ DEFAULT_OOB_BRONZE_SLA	5	7 days	30 days	8 weeks	6 Months	0 Quarters
☐ DEFAULT_OOB_BRASS_SLA	1	0 days	7 days	0 weeks	0 Months	0 Quarters
☐ NONE	0	0 days	0 days	0 weeks	0 Months	0 Quarters

Figure 21: SLAs

If you do not want to use the built-in SLAs, you can create your own SLAs in NDB. For more information, see [SLA Management](#) on page 153.

Provision a PostgreSQL Instance

1 Database Server VM

2 Instance

3 Time Machine

Name

postgres_TM

Description

SLA

DEFAULT_OOB_GOLD_SLA ▾

i

Schedule

Daily Snapshot at

01:00:00 ⌚

Snapshots Per Day

– 1 +

Log Catch Up Every

30 minutes ▾

Asia/Calcutta

Weekly Snapshot on

Friday ▾

Monthly Snapshot on the

15th ▾

Quarterly Snapshot in

Jan, Apr, Jul, Oct ▾

API Equivalent

Back

Provision

Figure 22: Time Machine

8. Click **Provision** to start provisioning your first PostgreSQL instance.

After the operation is complete, you have provisioned your first PostgreSQL instance in NDB. The instance is displayed in the list of databases. You can click the name of the instance to review the instance homepage. The instance homepage displays separate tiles for instance summary, database server VM, profiles, and time machine details.

What to do next

You can now proceed to use other NDB features such as copy data management.

Note: You can similarly provision Oracle, SQL Server, MariaDB, and MySQL instances. To create a software profile for provisioning an Oracle database, you must have previously registered an Oracle database with NDB.

For more information about database provisioning, see [One-Click Provisioning](#) on page 10.

NDB DASHBOARD

NDB provides a dashboard view that allows you to manage and monitor activities across your databases and database server VMs on NDB. To view the dashboard at any time, select **Dashboard** from the drop-down list of the main menu.




Figure 23: NDB Dashboard

Main Menu

The main menu at the top of every screen provides access to all the NDB features. The following options are displayed in the main menu:

- **View Menu.** The view menu on the left of the main menu provides a drop-down list of all the NDB entities.
- **User Menu.** The user menu on the right of the main menu provides options to update the user profile, access the REST API explorer, download the NDB CLI, change session settings, know the current NDB version, and log off from NDB.
- **Alerts.** The alerts icon on the main menu indicates the number of alert messages that have been generated and have not been marked as resolved. You can click **Alerts** to view all the alert details.
- **Operations.** The operations icon on the main menu indicates the number of operations running on NDB. You can click **Operations** to view all the operations and their status.

Dashboard

The dashboard is the opening screen that appears after logging into NDB. It provides a dynamic summary view that allows you to manage and monitor activities across your databases and database server VMs on NDB.

The dashboard displays the following widgets:

Table 8: Dashboard Widgets

Name	Description
Databases	Displays the total count of database groups, source databases, and database clones in NDB. Note: Database groups are currently supported only for SQL Server.
Clone Data Age (Days)	Displays the age range of cloned databases maintained by NDB.
Data Usage (GiB)	Displays information on the amount of data used and saved by using NDB CDM. A comparative analysis of the disk space usage when cloned traditionally (full clones) and when cloned using NDB is also displayed.
Operations (last 5 days)	Displays the summary of failed and succeeded operations from the last 5 days.
Alerts (last 5 days)	Displays the summary of critical and warning alerts from the last 5 days.
Summary Table	Displays a summary of all the source databases in NDB. The associated time machine, clones created, and data usage details are also displayed.
About	Displays the cluster name and version of NDB.

NDB ADMINISTRATION

NDB Service Management

You can view information about the NDB Service VMs, SMTP server and global notification settings, network segmentation configurations, Pulse, and the NDB service name (IP address of the NDB instance that is in use). You can also update the existing configurations, and enable high availability.

For more information, see [Enabling High Availability for NDB](#) on page 66 for NDB. See [Nutanix Cluster Management](#) on page 91) for information on NDB multi-cluster.

The **NDB Service** page view changes according to features (NDB High Availability, NDB Multi-cluster, network segmentation) enabled. The following tables describe the views available.




Figure 24: NDB Service Page: Default View (when NDB Multi-cluster, NDB High Availability, and Network Segmentation are disabled)

Table 9: NDB Service Page Default View Details

Widget	Description
NDB Service VMs	Displays the NDB server VM name, IP address, configured cluster, and the VM type information.
SMTP Server	Displays details of the configured SMTP server. For example, address of the server, username, assigned port number, and so on. Click Update to change the existing SMTP configuration or click Delete to remove the existing SMTP configuration.

Widget	Description
Global Notifications	Displays details of the configured NDB global notifications. You can configure global notifications to send alert notifications to one or more email recipients for the following operations. • Alerts • Failed Operations • Successful Operations Note: • Email notifications will only be sent to the email IDs saved under Email Recipients • By default, a Super Admin only receives notifications for their entities or operations.
Pulse	Pulse shares data about your application usage with Nutanix. All data collected is confidential. The information collected is used to help the product team understand how you are using the application and make better decisions about the direction of the product. Click Update to enable or disable Pulse. Pulse Status: • Enabled: Yes or No.. Indicates whether Pulse data collection is enabled or not. • Connected: Yes or No. Indicates if the last data upload to the Nutanix Telemetry service was successful or not. Note: Pulse is enabled by default.
NDB Service	Displays the current name of the NDB server. Click Update Service Name to update the NDB server name.

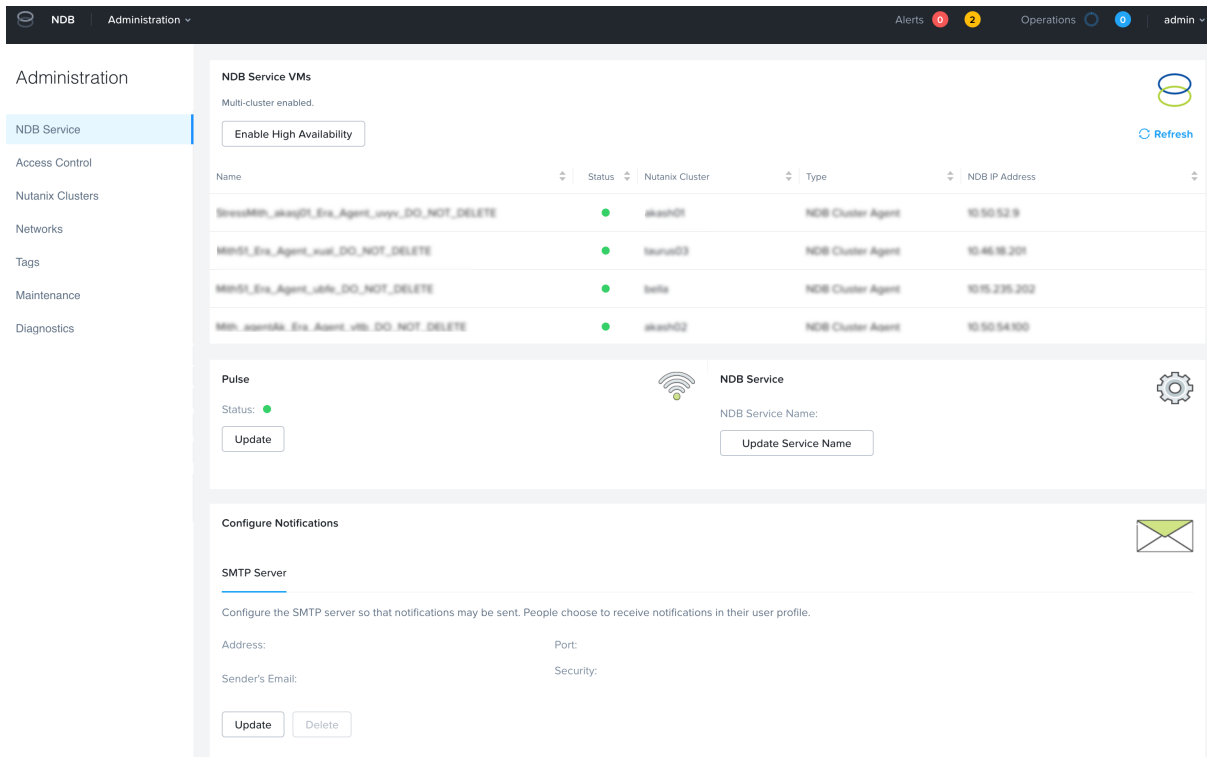


Figure 25: NDB Service Page: Multi-cluster View (when NDB Multi-cluster is enabled)

Table 10: NDB Service Page Multi-cluster View Details

Widget	Description
NDB Service VMs	Displays the list of NDB Service VM names, their IP addresses, configured clusters, and VM type information.
SMTP Server	Displays details of the configured SMTP server. For example, address of the server, username, assigned port number, and so on. Click Update to change the existing SMTP configuration or click Delete to remove the existing SMTP configuration.

Widget	Description
Global Notifications	Displays details of the configured NDB global notifications. You can configure global notifications to send alert notifications to one or more email recipients for the following operations. • Alerts • Failed Operations • Successful Operations Note: • Email notifications will only be sent to the email IDs saved under Email Recipients • By default, a Super Admin only receives notifications for their entities or operations.
Pulse	Pulse shares data about your application usage with Nutanix. All data collected is confidential. The information collected is used to help the product team understand how you are using the application and make better decisions about the direction of the product. Click Update to enable or disable the Pulse. Note: Pulse is enabled by default.
NDB Service	Displays the current name of the NDB server. Click Update Service Name to update the NDB server name.




Figure 26: NDB Service Page: Network Segmentation View (when NDB Multi-cluster is enabled and Network Segmentation is configured).

Table 11: NDB Service Page Network Segmentation View Details

Widget	Description
NDB Service VMs	Displays the list of NDB Service VM names, their IP addresses, configured clusters, network segmentation, and VM type information. Click Network Segmentation to view the current VLAN configurations of the respective cluster.
SMTP Server	Displays details of the configured SMTP server. For example, address of the server, username, assigned port number, and so on. Click Update to change the existing SMTP configuration or click Delete to remove the existing SMTP configuration.

Widget	Description
Global Notifications	Displays details of the configured NDB global notifications. You can configure global notifications to send alert notifications to one or more email recipients for the following operations. • Alerts • Failed Operations • Successful Operations Note: • Email notifications will only be sent to the email IDs saved under Email Recipients • By default, a Super Admin only receives notifications for their entities or operations.
Pulse	Pulse shares data about your application usage with Nutanix. All data collected is confidential. The information collected is used to help the product team understand how you are using the application and make better decisions about the direction of the product. Click Update to enable or disable Pulse. Note: Pulse is enabled by default.
NDB Service	Displays the current name of the NDB server. Click Update Service Name to update the NDB server name.

NDB Service Page: High Availability View (when NDB Multi-cluster and NDB High Availability is enabled).




Figure 28: NDB Service Page: High Availability View - Topology Format

Table 12: NDB Service Page High Availability View Details

Widget	Description
NDB Service VMs	Click the Table tab to display the list of NDB Service VM names, their IP addresses, configured clusters, network segmentation, and VM type information. Click the Topology tab to display the topological view of the NDB high availability structure consisting of the Nutanix clusters, API Server VMs, Repository VMs, and the HA Proxy VMs.
SMTP Server	Displays details of the configured SMTP server. For example, address of the server, username, assigned port number, and so on. Click Update to change the existing SMTP configuration or click Delete to remove the existing SMTP configuration.

Widget	Description
Global Notifications	Displays details of the configured NDB global notifications. You can configure global notifications to send alert notifications to one or more email recipients for the following operations. • Alerts • Failed Operations • Successful Operations Note: • Email notifications will only be sent to the email IDs saved under Email Recipients • By default, a Super Admin only receives notifications for their entities or operations.
Pulse	Pulse shares data about your application usage with Nutanix. All data collected is confidential. The information collected is used to help the product team understand how you are using the application and make better decisions about the direction of the product. Click Update to enable or disable Pulse. Note: Pulse is enabled by default.
NDB Service	Displays the current name of the NDB server and the assigned NDB service virtual IP address. Note: The assigned virtual IP address is set as the default IP address of the NDB server. Click Update Service Name to update the NDB server name.
NDB Virtual IP Address	Displays the assigned NDB service virtual IP address when NDB HA is enabled. Note: The assigned virtual IP address is set as the default IP address of the NDB server.

NDB High Availability

You can enable high availability (HA) for NDB to ensure that there is no loss on service or data when the NDB VM fails.

As part of enabling HA, NDB creates a cluster of API server VMs, Repository VMs, and HA proxy VMs. These cluster servers help in securing the system against unexpected failures, planned maintenance, and upgrades. When the primary API server VM or Repository VM goes down, NDB automatically fails over to a standby VM.

See [NDB Service Resiliency](#) on page 71 for information on NDB HA resiliency.

For information on HA resource requirements, see [NDB Control Plane Configuration and Scalability](#) on page 13.

Note: NDB HA is not a disaster recovery solution.

Caution:

- You cannot cancel or revert the operation to enable NDB HA once started. Similarly NDB HA cannot be disabled once enabled.
- You cannot enable NDB HA if you have configured network segmentation for NDB. Similarly you cannot configure network segmentation for NDB if you enable NDB HA.




Figure 29: Example: Three Cluster NDB HA Topology

Enabling High Availability for NDB

Before you begin

- To ensure high availability with auto-failover for the NDB control plane, Nutanix requires three Nutanix clusters to distribute the NDB control plane VMs across failure domains. To protect against site failure, these should be at separate DCs or sites with a round-trip time (RTT) of less than 5ms. A two-cluster topology is not supported as the Nutanix cluster running two API servers and two repository servers is considered a single point of failure.
- Ensure that network profiles have at least eight free IP addresses.
- Two IP addresses for new API servers and three for PostgreSQL repository servers

Note: The existing NDB server is converted into an API server and continues using its current IP address.

- Two IP addresses for HA proxy VMs and one for HA proxy VIP
- All the Nutanix clusters on which you want the HA nodes to be available must be registered with NDB.
- All the Nutanix clusters on which you want the HA VMs to be available must have the `POSTGRES_10.4_HA_ENABLED_OOB` profile.
- You can enable NDB HA across three clusters where two clusters run AHV and the third runs ESXi, or two clusters run ESXi and the third runs AHV. Install Nutanix Guest Tools (NGT) on the NDB server for cross-hypervisor support (AHV and ESXi).

- NDB multi-cluster must be enabled. See [Enabling NDB Multi-Cluster](#) on page 91

About this task

Perform the following procedure to enable High Availability for NDB.

Caution:

- You cannot cancel or revert the operation to enable NDB HA once started. Similarly, NDB HA cannot be disabled once enabled.
- You cannot enable NDB HA if you have configured network segmentation for NDB. Similarly, you cannot configure network segmentation for NDB if you enable NDB HA.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **NDB Service** and click **Enable High Availability**. The window **Enable Service HA** appears. You enable HA in the following steps:
 - a. **Nutanix Cluster**
 - b. **NDB VMs**

3. In the **Nutanix Cluster** step, do the following in the indicated fields:
- a. Select the Nutanix clusters to host the NDB VMs during a failover.

- Note:**
- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
 - Nutanix recommends selecting either one or three Nutanix clusters to host the NDB VMs. Selecting two clusters for the same is not supported.

- b. Click **Next**.

Enable Service HA? ×

1 Nutanix Cluster

2 NDB VMs

To increase availability across clusters, it is recommended to select at least 2 clusters.

Select the Nutanix Clusters to host the NDB VMs

☐	Nutanix Cluster	Description
☐	b	
☒	E	d
☒	ta	C

Cancel

Next

Figure 30: Enabling High Availability : Nutanix Cluster Step

4. In the **NDB VMs** step, do the following in the indicated fields:

a. **Name Prefix.** Type a prefix for the NDB VM name.

This field is optional. If you do not type a prefix, NDB automatically selects NDB as a default prefix name.

b. **Compute Profile.** Select a compute profile from the drop-down list of the compute profiles available.

Note: Nutanix recommends a compute profile with memory greater than 16GB.

c. **Network Profile for API Server & Repository.** Select a PostgreSQL HA network profile that contains VLANs for all the selected Nutanix clusters from the drop-down list available. This network profile will be used to host the API Server VMs and the Repository VMs.

Select a profile with NDB-managed VLANs. If you do not have a static VLAN, you can create one (For more information, see [NDB Network Management](#) on page 101).

Note:

- A HA instance network profile appears in the list only if you have created the profile on the selected Nutanix clusters. For information on creating a clustered network profile, see [Creating a HA Instance Network Profile](#).
- Enabling NDB High Availability does not require stretched VLANs.

d. **Network Profile for HA Proxy.** Select a PostgreSQL HA network profile that will be used to host the HA Proxy VMs.

Note: To host HA proxy VMs on different Nutanix clusters you must select a PostgreSQL HA network profile that contains a stretched VLAN. See [Creating a HA Instance Network Profile](#).

e. All the VMs that are created as part of the operation are displayed under **Attributes of Individual NDB VMs**. The following types of VMs are created.

- **Repository.** Runs the data persistence layer of NDB control plane.
- **API Server.** Stateless REST API gateway for NDB control plane.
- **HA Proxy.** Load balances access to API servers and repository servers.

Select the Nutanix clusters for the respective VM types.

f. **Topology.** Displays a topological view of NDB High Availability structure consisting of the Nutanix clusters, API Server VMs, Repository VMs, and the HA Proxy VMs.

g. Click **Enable High Availability**.

Enable Service HA

?

1 Nutanix Cluster

2 NDB VMs

Attributes of All NDB VMs

Name Prefix

Era

Compute Profile

DEFAULT_OOB_COMPUTE

Network Profile for API Server & Repository

st

Network Profile for HA Proxy

st

Attributes of Individual NDB VMs

Era VM Name	Nutanix Cluster	Type
E	E	Repository
E	g	API Server
E	E	API Server
srly	Era	API Server
E	E	HA Proxy

Topology

HA Proxy

E

E

API Server

E

Repository

E

E

g

API Equivalent

Back

Enable High Availability

Figure 31: Enabling High Availability : NDB VMs Step

A message indicating that an operation to enable NDB High Availability has started is displayed. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

Note: When the operation is in progress, NDB service goes down for a few minutes and all the user sessions are logged out.

NDB Service Resiliency

NDB provides service resiliency that allows you to recover from a failure to the latest stable state. NDB achieves resiliency by monitoring critical NDB services and periodically backing up the necessary datasets.

Table 13: NDB Failure Scenarios

Failure Scenario	Recoverable	Impact	Fix
NDB server process is down.	Yes, auto-recoverable.	NDB service is down for few seconds.	-
NDB database is down.	Yes, manual interaction required.	NDB service is down.	Restart the database.
NDB worker service is down.	Yes, auto or manually recoverable.	The database server VM with the NDB worker service will be impacted. NDB worker to NDB management plane communication is not possible. Log catchup operations do not work.	Start or restart the NDB worker service running in the DB Server VM. The NDB worker service is a process for Linux based database server VMs and is a service for Windows based database server VMs.
NDB agent VM is down.	Yes, manual interaction required.	NDB operations happening in the Nutanix clusters managed by the agent VM are impacted.	Manually bring the NDB agent VM to an online state. For further assistance, call Nutanix Support.
NDB server VM is down.	Yes, manual interaction required.	NDB service is down and inaccessible.	Start or restart the VM. For further assistance, call Nutanix Support.
NDB log drive is corrupted or deleted.	Partially recoverable, manual interaction required. Note that the NDB service internal Recovery Point Objective (RPO) takes 1 hour.	Point-in-time recovery (PITR) is not available for some time.	For assistance, call Nutanix Support.
NDB server VM is lost or deleted.	Partially recoverable, manual interaction required. Note that the NDB service internal Recovery Point Objective (RPO) takes 1 hour.	NDB service is down and inaccessible.	For assistance, call Nutanix Support.

Table 14: NDB HA Failure Scenarios

Failure Scenario	Recoverable	Impact	Fix/Verification
A single API Server VM is down.	Yes, auto-recoverable.	If the failed API server VM is receiving end user connections, then the API service will be unavailable for a few seconds. If the failed API server VM is not receiving end user connections, then the API service will not be unavailable.	Ensure that the failed API server VM is started in case it does not automatically.
More than one API Server VMs are down.	Yes, auto-recoverable.	If anyone of the failed API server VMs is receiving end user connections, then that API service will be unavailable for a few seconds. If none of the failed API server VMs receive end user connections, the API service will not be unavailable.	Ensure that the failed API server VM is started in case it does not automatically.
All the API Server VMs are down.	Yes, manual interaction required.	NDB API service will be unavailable.	Ensure that at least one API server VM is online to restart the API service.
Note: NDB HA can only tolerate up to two service VM failures.			
A single HA proxy VM is down.	Yes, auto-recoverable.	If the failed HA proxy VM was managing a virtual IP address (VIP), then the HA proxy functionality will be unavailable for a few seconds. If the failed HA proxy VM was not managing a VIP, then there will be no HA proxy service interruption.	Ensure that the failed HA proxy VM is started in case it does not automatically.

Failure Scenario	Recoverable	Impact	Fix/Verification
More than one HA proxy VMs are down.	Yes, manual interaction required. Note: NDB HA can only tolerate up to a single HA proxy VM failure.	NDB service will be unavailable since the end user will be disconnected from the API service. The API service therefore cannot connect with the PostgreSQL repository.	Ensure that at least one HA proxy VM is online to restart the HA proxy functionality.
A single repository VM is down.	Yes, auto-recoverable.	If the failed repository VM is the primary or secondary (synchronised copy) node, then the NDB repository service will be unavailable for a few seconds. If the failed repository VM is not the primary or secondary (synchronised copy) node, then there will be no service interruption.	Ensure that the failed repository VM is started in case it does not automatically.
More than one repository VM are down.	Yes, manual interaction required. Note: NDB HA can only tolerate up to a single repository VM failure.	NDB service is down.	Call Nutanix Support to bring the NDB repository service online.

Enabling NDB Pulse Telemetry

This topic describes the steps to enable or disable NDB Pulse telemetry.

About this task

Pulse shares data about your application usage with Nutanix. The information collected helps the product team understand how you are using the application and make better decisions about the direction of the product.

Note: Pulse is enabled by default and shares the following data with Nutanix.

Operation:

```
["id", "type", "status", "cluster_uuid", "sequence_number",
"system_triggered", "feature_type", "failed_step", "error_message"]
```

Database:

```
["id", "name", "type", "status", "cluster_uuid", "sequence_number",
"is_clone", "timemachine_id", "sizeTB"]
```

Cluster:

```
["id", "name", "type", "status", "cluster_uuid", "sequence_number",
"nx_cluster_uuid", "version", "hypervisor_type", "hypervisor_version",
"num_nodes", "max_cores"]
```

Time machine:

```
["id", "name", "type", "status", "cluster_uuid", "sequence_number", "size"]
```

Database server VMs:

```
["id", "name", "type", "status", "cluster_uuid", "sequence_number",  
"era_version", "is_internal"]
```

Pulse sends the telemetry data to Nutanix at <https://insights.nutanix.com> on port 443. Perform the following steps to enable or disable Pulse.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **NDB Service** and click **Update** under **Pulse**. The **Update Pulse** window appears.
3. To enable (disable) Pulse, check (uncheck) the **Enable Pulse** check-box and click **Update**.
4. To verify the configuration, hover the cursor over the green status light.
Enabled: Yes indicates that Pulse is turned on. **Connected: Yes** indicates that NDB is able to send telemetry data.

Role-based Access Control

NDB employs role-based access control (RBAC) to govern access to an NDB system. RBAC allows system administrators to restrict user access and limit the operations they can perform. For example, a database operator can only be given access that is related to database-related operations.

With RBAC, you can add users either in NDB or Active Directory and then later you can add those users to a group. The group allows you to assign the same roles to multiple users at the same time. After adding a user, you can assign one or more roles to the user based on the operations they can perform.

Terminology

Users

A user represents a person who interacts with NDB. Users can view only the operations that they have initiated. Only a super administrator user is allowed to view all the operations that any user has initiated.

Roles

A role describes the level of access that is associated with the account of user. A role grants privileges to perform the specified actions on the resource. Each privilege is specified explicitly in the role.

Table 15: Type of Roles

Role	Description
Super administrator	Provides the ability to have full control over NDB from setting up NDB, creating users, upgrading and patching NDB, setting disaster recovery and High Availability for NDB, alert policy management, telemetry, and managing all entities and metadata in NDB.

Role	Description
Infrastructure administrator	Provides the ability to perform infrastructure-related tasks such as adding a Nutanix cluster to NDB in a multi-cluster environment, adding VLAN and IP pool to NDB, and creating compute and network profile.
Database infrastructure administrator	Provides the ability to perform DBA-related tasks such as creating software and database parameter profile.
Database administrator (DBA)	Provides the ability to perform database, database server, clone, and time machine operation such as provisioning, registering, cloning a database, creating and managing tags.

Privileges

A privilege consists of a specified entity and the actions permitted on the entity. A user assigned a role receives all the privileges of that role. A user can have one or more roles assigned.

Privileges limit the actions that users can perform on an entity. For example, only an infrastructure administrator has the ability to add a new VLAN or IP pool.

Privileges and Permission

Each role has a list of privileges that a user has access to. Only a super administrator user has full access to all the privileges. If an entity is owned by a user, by default the user can view and modify that entity including operations and alerts. Privileges are granted based on the role a user plays in an organization. For example, the ability to add a new VLAN or IP Pool should be allowed only if the user role is Infrastructure administrator or above.

For more information on NDB entities and entity sharing policies, see [NDB Entity Sharing](#) on page 85.

Table 16: Privileges and Permission

Privilege	Super Administrator	Infrastructure Administrator	Database Infrastructure Administrator	Database Administrator (DBA)
Databases	Yes	Yes	Yes	Yes
Time Machines	Yes	Yes	Yes	Yes
Compute Profiles	Yes	Yes	Yes	No
Network Profiles	Yes	Yes	Yes	No
Software Profiles	Yes	Yes	Yes	No
Database Parameter Profiles	Yes	Yes	Yes	No
Clones	Yes	Yes	Yes	Yes
Database Servers	Yes	Yes	Yes	Yes
Access Control	Yes	No	No	No
NDB Configuration	Yes	No	No	No

Privilege	Super Administrator	Infrastructure Administrator	Database Infrastructure Administrator	Database Administrator (DBA)
SLAs	Yes	Yes	Yes	Yes
Tags	Yes	Yes	Yes	Yes
Alerts	Yes	Yes	Yes	Yes

Note: The super administrator user cannot access tags created by other users from the UI. From the API, the super administrator can view the tags created by other users only by specifying the tag name or ID.

Adding a User

About this task

Roles grant users access to the NDB entities. NDB provides four roles that administrators can use to control access to an NDB system. You must have the super administrator role assigned to create a user and assign one or more roles to a user. Only a super administrator user is allowed to view all the operations that any user has initiated. Users can view only the operations that they have initiated.

Perform the following procedure to add a user.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Users**.

2. To add a user, click **Add** and select one of the following:

- » **Local User.** A local user is authenticated by a password provided by the user in NDB. Select this option to add a local user.
- » **User from Active Directory.** The user is authenticated by Active Directory. Select this option if you want to create a user from Active Directory.

Note: The **User from Active Directory** option is available when Active Directory is configured with NDB.




Figure 32: Adding a New User

The **Add User** window appears.

3. Do the following in the indicated fields:

- a. **Username.** Type a name for the user. If you selected **User from Active Directory** in the previous step, type a name for the user and then click **Verify** to verify the user in the Active Directory list.
- b. **Password.** Type a temporary password for the username entered. Ensure that you reset your password after first login.

Note: User from Active Directory is authenticated by Active Directory. This field is not populated if you selected remote user in the previous step.

- c. **Super Admin.** Super Administrator has full control over NDB and can manage all entities, configure, and upgrade NDB. Select this check box if you want to assign the Super Administrator role to a user.
- d. **Roles.** Select the user roles from the available lists of roles. You can assign one or more roles to a user. See [Privileges and Permission](#) on page 75 for information on the role privileges and permissions.

Note: If you have assigned the **Super Admin** role to a user, the **Roles** field is automatically disabled.

- e. **Email.** Type an email address of the user.

4. Click **Add** to create a user.

This step displays a list of the users that you have created. The list includes the name of a user, the role that a user is assigned to, and email address of a user. You can use the filter option to filter the list based on the roles.

Updating an Existing User

About this task

Modifying an existing user allows you to edit the email address, password, and the role that you have previously assigned to a user. You can also grant or remove super administrator privileges to a user. Only a user with the super administrator role is allowed to update existing user details. Perform the following procedure to update an existing user.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Users**.
2. Select the user that you want to update and click **Update**.
The **Update User** window appears.
3. Update the indicated fields and click **Update** to update a user.
This step updates the user data and the users' access to the operations that a user can perform.

Removing a User

About this task

You can remove a user and transfer the access that user had to either administrator or some other user with the same role. Perform the following procedure to delete a user.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Users**.
2. Select the user that you want to remove and click **Remove**.
The **Remove User** confirmation window appears.
3. Select a user from the drop-down list to transfer all the access.
4. Click **Remove**.
The user is removed from NDB and all the access privileges are transferred to the selected user.

Configuring User Profile Settings

NDB supports Role-based Access Control (RBAC). Once logged in, you can modify your profile according to your needs.

About this task

Perform the following procedure to configure the user profile settings.

Procedure

1. Select **Profile** from the **User Menu**.

The **User Profile** page appears.

The screenshot shows the 'admin's Profile' configuration window. It includes fields for Email, Time Zone, Language, and Timeout. There are also checkboxes for 'Display in High Contrast Mode', 'Enable debug logging', and 'Change Password'. A section for 'Receive email notifications for' includes checkboxes for Alerts, Failed Operations, and Successful Operations. A 'Roles and Privileges' section shows a list with '+ Super Admin'. At the bottom right, there are 'Cancel' and 'Save' buttons.

Figure 33: User Profile Page - Configure Settings

2. In the **Email** box, enter a valid email address of a user.
3. In the **Time Zone** drop-down list, make sure to select the time zone that you want to use.
4. In the **Timeout** drop-down list, select an appropriate timeout value for the user sessions.
5. Select the **Display in High Contrast Mode** check box if you want to present the NDB GUI on large screens.
6. Select the **Enable debug logging** check box when logging into the GUI in the browser's console. You can select this option to perform GUI troubleshooting.
7. Select the **Change Password** check box if you want to change the current user password. Enter the current and new passwords in the fields that appear.
8. In the **Receive email notifications for** area, you can select any one or all the following options - **Alerts**, **Failed Operations**, **Successful Operations**. You will receive an email notification when the selected events occur.

Note: These options appear after you configure the SMTP server. Ensure that the details of the SMTP server are configured correctly to receive email notifications. For more information on configuring the SMTP server, see [Sending Alert Notifications to an Email Recipient](#) on page 163.

9. In the **Roles and Privileges** area, you can view the role and privileges assigned to you.
10. Click **Save** to apply your settings and close the **User Profile** page.

Adding, Updating, and Removing a Group

The group allows you to assign the same roles to multiple users at the same time. For this operation, you must create a group in Active Directory and add users to that group. You can then select the group and assign roles to all the users in that group in NDB.

Before you begin

Ensure that you create a group in Active Directory and add users to the group.

About this task

Perform the following procedure to add a group and then assign roles to all the users in that group.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Groups**.

2. To add an Active Directory group, click **Add**. The **Add Active Directory Group** window appears.

3. In the **Group** step, under **Active Directory Group Name**, type an Active Directory group name and click **Verify**.

After the group name is verified, NDB displays a list of users that are part of the given group in the Active Directory and the existing users if the user had been already added previously.

4. Click **Next** to proceed to the next step.

5. In the **Roles** step, select the roles that you want to assign to the group.

The roles that you select are assigned to all the users that are added to that group. If the user was already assigned some roles previously, then the respective user gets assigned with the additional roles.

6. Click **Add**.

After the group is created, the NDB displays the group details on the **Access Control > Groups** page. You can click the number displayed in the **Users** column to view the details of users that are added in that group.

If you want to update the group that you have created previously, select the group name and click **Update**. You can only update the roles that are assigned to all the users in that group.

If you want to remove a group, you can select the group and click **Remove**. The **Remove Group** dialog box appears. In **Transfer all entities to**, select the user to transfer all the access and click **Remove**. This operation removes the users who do not have any roles assigned outside of that group but retains users that are added to another group or have additional roles assigned in NDB.

Refreshing a Group

Refresh operation synchronizes the users in NDB with the users in active directory.

About this task

You can add users either in NDB or Active Directory. When you add a user to a group in the active directory, then the refresh operation adds the user in NDB and assigns the same roles to user that are granted for the entire group. When you remove a user from the group in the active directory, then the refresh operation removes the user from the group in NDB and all the roles that were assigned to the user as part of that group are revoked.

Perform the following procedure to refresh a group.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Groups**.

2. Select the group that you want to refresh and click **Refresh**.
This operation refreshes the NDB user status of all the users in that group.

Viewing Users' Roles

Only a user with the super administrator role is allowed to view other users' information and the operations initiated by other users. Users with other than the super administrator roles can view only the operations that they have started.

About this task

Each role is assigned with certain privileges that a user has access to. Only a super administrator user can view the information for all the roles and number of users that are assigned to that role. Perform the following procedure to view the list of roles and privileges that each role has.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Access Control > Roles** and then select the role for which you want to view the list of privileges.
This step displays the role details that include role name, description, number of users assigned to that role, and list of privileges that role has. You can click the number next to Users to view the list of users that are assigned that role.

Creating a Role

About this task

Each role is composed of privileges that a user has access to. You can create a role and assign privileges to it in NDB. If you create a role with create role privilege, this role is visible to all the users in NDB.

Procedure

1. From the Dashboard drop-down, select **Administration**.
2. Go to **Access Control**, select **Roles**, and click **Create**.

3. In the **Create Role** window that appears, enter the following details:
- Name.** Type a name for the role that you want to create.
 - Description.** Type a description for the role.
 - Select the privileges that you want to assign to the role. You can select from the following privileges:
 - Add Cluster
 - Add Group
 - Add IP Pool
 - Add User
 - Configure Authentication Directory
 - Configure Notification Settings
 - Create Access Policy
 - Create Clone
 - Create Maintenance Window
 - Create Maintenance-Tasks
 - Create Profile
 - Create Profile Version
 - Create Role
 - Create SLA
 - Create Tag
 - Enable Era Multi-Cluster
 - Enable Era service HA
 - Extend Database Storage
 - Fetch TM Capability
 - Generate Diagnostic Bundle
 - Patch Database Server
 - Patch Database Server OS
 - Provision Database
 - Provision Database Group
 - Provision Database Server
 - Provision Standby Database
 - Provision Standby Database Group
 - Refresh Clone
 - Refresh Clone Group
 - Register Database

- Register Database Group
- Register Database Server
- Restore Database
- Restore Database in Group
- Scale Database Server
- Upgrade Database Server
- Upgrade Era

d. Click the right arrow button to add the selected privileges to the role.

4. Click **Create**.

5. From the **Roles** page that appears, select the new role from the list to view the details.

Create Role ? | X

Name

Description

☐ 33 items **Privileges**

Search Privileges

- ☐ Add Cluster
- ☐ Add Group
- ☐ Add IP Pool
- ☐ Add User
- ☐ Configure Authentication Directory
- ☐ Configure Notification Settings
- ☐ Create Access Policy

☐ 0 item **Selected Privileges**

Search Privileges

No Privileges Selected

Figure 34: Creating a New Role

The role details include role name, description, number of users assigned to that role, and list of privileges that role has. You can click the number next to users to view the list of users that are assigned that role.

Updating a Role

About this task

Modifying an existing role allows you to update the role name, description, or the privileges that are assigned to a role. Perform the following procedure to update a role.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Access Control > Roles**. Select a role and click **Update**. The **Update Role** window appears.
3. Update the indicated fields and click **Update** to update a role.
This step updates the role name, description, or privileges that were assigned to this role previously.

Removing a Role

About this task

If a role is no longer required, you can remove the role from NDB. Removing a role automatically removes that role from all the users that were using it. Perform the following procedure to remove a role.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Access Control > Roles**. Select a role and click **Remove**. This step displays a confirmation dialog box.
3. Click **Yes** to remove a role.

Configuring Active Directory Access

About this task

NDB authenticates a remote user through active directory. NDB must have access to the active directory server to authenticate users. Perform the following procedure to configure active directory access.

Note: NDB uses port 389 or 636 to communicate with the active directory.

Procedure

1. From the drop-down list of the main menu, select **Administration > Access Control > Authentication**.
This page displays the current active directory configuration.
2. Click **Configure** to configure active directory access. This **Configure Active Directory Access** window appears.

3. Do the following in the indicated fields.
 - a. **Configuration Name.** Enter a name for the configuration.
 - b. **Domain.** Enter the domain name of the configuration.
 - c. **Primary URL.** Select the connection type as either **ldap** or **ldaps**, enter the IP address, and port number.
 - d. **Secondary URL.** Secondary URL is used for redundancy purposes. Select the connection type as either **ldap** or **ldaps**, enter the IP address, and port number.
 - e. **Service Account Username.** Enter a username for the service account.
 - f. **Service Account Password.** Enter a password.

Figure 35: Configuring Active Directory Access

4. Click **Verify** to check if you can connect to an active directory server.
5. Click **Save** to save the details.

Note: **Save** button is enabled only when the verification is successful.

This step displays the current active directory configuration settings. Click **Remove** to delete the current configuration. The configuration is not deleted until users are associated with the configuration. Remove the associated users and then delete the configuration. Repeat the previous steps to add a new configuration.

NDB Entity Sharing

Entity sharing in NDB allows you to share the following NDB entities that are owned by a single user.

- Databases
- Database clones
- Database server VMs
- Database server clusters
- Snapshots
- Time machines

An entity created by a user can be shared with other NDB users or active directory user groups based on the entity sharing policy configured.

Note:

- The entities are visible only to the owner of the entities unless shared with other users through an entity sharing policy.
- Users with a super admin role have access to all entities in NDB.
- A user without the super admin role cannot view the status of an operation even if you share the entity on which the operation is performed with the user through an entity sharing policy.
- You cannot share profiles, SLAs, and tags through entity sharing policies. To share a profile with all users, set the status of the profile to **Published**. For more information, see [Updating a Software Profile Version](#).

Entity Sharing Policies

You can share an NDB entity with another user or AD user group by configuring entity sharing policies. Different policies provide different access level privileges to the shared user.

The following access level privileges can be configured for the NDB entities.

Table 17: Entity Sharing Access Levels

Access Level	Description
View Access	Provide read-only access.
Use Access	Provides partial ownership access on the shared entity. Note: Users with use access cannot delete or share the shared entity.
Manage Access	Provides partial ownership access on the shared entity. Note: Users with manage access cannot delete or share the shared entity.
Full Access	Provides full ownership access on the shared entity. The following actions can be performed.

Table 18: Entity Sharing Policies: NDB Entities and the Access levels, Privileges

NDB Entity	Access Levels and Privileges			
	View Access	Use Access	Manage Access	Full Access
Databases	View all metadata	Not supported	Not supported	Not supported
Database Clones	View all metadata	Not supported	Not supported	Not supported

NDB Entity	Access Levels and Privileges			
	View Access	Use Access	Manage Access	Full Access
Database server VMs and database server clusters	View all metadata	- Provision a database in the database server VM or database server cluster. - Update database server VM or database server cluster details.	- Provision a database in the database server VM or database server cluster. - Update database server VM or database server cluster details. - Apply patch. - Upgrade the database server VM or the database server cluster.	- Provision a database in the database server VM or database server cluster. - Update database server VM or database server cluster details. - Delete or remove a database from the database server VM or database server cluster. - Apply patch. - Upgrade the database server VM or the database server cluster. - Share the database server VM or the database server cluster.
Snapshots	View all metadata.	- Create clones using the snapshot. - Update snapshot details.	Not supported	Not supported

NDB Entity	Access Levels and Privileges			
	View Access	Use Access	Manage Access	Full Access
Time Machines	View time machine details including snapshots.	- View time machine details including snapshots. - Create clones of the time machine.	- View time machine details including snapshots. - Create clones of the time machine. - Manually create and replicate snapshots and log catchups. - Update, pause or resume time machine. - Manage replication policy.	- View time machine details including snapshots. - Create clones of the time machine. - Manually create and replicate snapshots and log catchups. - Create/Update snapshot removal schedule. - Update, pause or resume time machine. - Delete time machine or snapshot. - Share time machine. - Manage replication policy.

Creating an Entity Sharing Policy

About this task

Perform the following procedure to create an entity sharing policy.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Access Control > Entity Sharing Policies** and click **Create**. The **Create Entity Sharing Policy** window appears.

3. Do the following in the indicated fields:

a. **Entity Type.** Select the entity type that you want to share from the drop-down list.

The following entity types can be shared:

- **Database.** Select this option to share a database and the associated database server VM.

Note: A database can only be shared with View access policy.

- **Clone.** Select this option to share a database clone.

Note: A database clone can only be shared with View access policy.

- **DB Server VM.** Select this option to share a database server VM.

- **Snapshot.** Select this option to share a snapshot and the associated time machine.

Note: A snapshot can only be shared with Use access policy.

- **DB Server Cluster.** Select this option to share a server cluster. All the database server VMs in the server cluster are also shared.

- **Time Machine.** Select this option to share a time machine.


b. **Entity.** Select the entity that you want to share from the drop-down list.

c. **Access Recipient.** Select the user or AD user group with whom you want to share the selected entity.

d. **Access Level.** Select the access level privileges that you want to assign to the entity sharing policy that you are creating.

e. **Description.** Type a description for the policy.

4. Click **Create**.



The image shows a 'Create Entity Sharing Policy' dialog box. At the top, the title 'Create Entity Sharing Policy' is centered, with a close button (X) on the right. The form contains several fields: 'Entity Type' with a dropdown menu showing 'Database'; 'Entity' with a dropdown menu showing 'Select Entity'; 'Access Recipient' with a dropdown menu showing 'Select Recipient'; 'Access Level' with a dropdown menu showing 'View'; and a 'Description' text area. At the bottom, there are three buttons: 'API Equivalent' (disabled), 'Cancel', and 'Create Policy' (highlighted in blue).

Figure 36: Create Entity Sharing Policy

A message indicating that the policy is successfully created is displayed.

The entity sharing policy appears in the **Entity Sharing Policies** list. Click the name of the database to open the summary page for the selected entity.

Updating an Entity Sharing Policy

About this task

Perform the following procedure to update the entity sharing policy details.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Access Control > Entity Sharing Policies**, select a cluster and click **Update**. The **Update Policy** window appears.

3. Do the following in the indicated fields.
 - a. **Description.** Type an updated description for the policy.
 - b. **Access Level.** Select the access level privileges that you want to assign to the entity sharing policy that you are updating.
 - c. Click **Update**.

The screenshot shows a modal window titled "Update Policy" with a close button in the top right corner. The window is divided into two columns. The left column contains the "Entity Name" field (with a partially visible value "m...20") and the "Description" field (which is empty). The right column contains the "Access Recipient" field (with a partially visible value "dl...m") and the "Access Level" dropdown menu (which is currently set to "Use"). At the bottom of the window, there are three buttons: "API Equivalent", "Cancel", and "Update".

Figure 37: Updating Entity Sharing Policy

Nutanix Cluster Management

NDB allows multi-cluster database management, with NDB multi-cluster you can easily manage databases across multiple locations both on-prem and cloud with Nutanix clusters.

As part of enabling NDB multi-cluster, the NDB server creates an NDB agent that runs on a dedicated VM. All the remote Nutanix clusters added to NDB have an instance of the NDB agent created to manage, schedule, perform, and execute database operations of that cluster.

Enabling NDB Multi-Cluster

Before you begin

Ensure that the Nutanix cluster AOS version is supported for multi cluster. See *NDB Software Compatibility and Feature Support* in *Nutanix Database Service Release Notes* for more information.

About this task

Perform the following procedure to enable NDB multi-cluster.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Nutanix Clusters**, click **Enable Multi-Cluster**. The **Enable Multi-Cluster** window appears.

Caution: You cannot disable NDB multi-cluster once enabled.

3. Do the following in the indicated fields:

- a. **Name Prefix.** Type a prefix for the NDB agent name.

This field is optional. If you do not type a prefix, NDB automatically selects NDBAgent as a default prefix name.

- b. The following VLAN options are displayed based on the existing network segmentation configurations of the NDB server.

The following VLAN option is displayed if you do not have network segmentation configured on your NDB server.

- **VLAN.** Select a VLAN to provision the NDB Agent VM from the drop-down list.

Select an NDB-managed static VLAN. If you do not have a static network, you can create one (For more information, see [NDB Network Management](#) on page 101).

Enable Multi-Cluster

?

×

Before another Nutanix Cluster can be registered in NDB, the NDB Agent must be installed on a new VM.

Warning:

- This operation will quiesce NDB Service for approximately 10 minutes
- This operation may not be undone

Provision Agent VM

Name Prefix

vLAN

Select a vLAN

⌵

vCPUs: 4

Memory: 4 GiB

Storage: 50 GiB

API Equivalent

Cancel

Enable Multi-Cluster

Figure 38: Enabling NDB Multi-Cluster: Without Network Segmentation

The following network segmentation options are displayed if you have network segmentation configured on your NDB server. By default, the **vLAN**, **Gateway**, and **Subnet Mask** values are populated based on the existing NDB server network configurations.

- Type the **IP Address** for the configured VLAN access types.

NUTANIX

Nutanix Database Service (formerly Era) | NDB Administration | 92

Enable Multi-Cluster

Provision Agent VM

Name Prefix

ErstAgent

vCPUs: 4

Memory: 4 GiB

Storage: 50 GiB

Agent Network Segmentation

This is an optional setting for Advanced Network Segmentation. A network interface will be attached to NDBCluster Agent VM for each of the Distinct vLAN Names provided in the below section.

vLAN Access Type	vLAN	Configuration		
		Gateway	Subnet Mask	IP Address
☒ Prism Element	v1000	10.10.1.1	255.255.255.0	
☐ Prism iSCSI Data Service	v1000	10.10.1.1	255.255.255.0	
☒ DBServer Access from NDB Agent	v1000k	10.10.1.1	255.255.255.0	
☒ vCenter	v1000atic	10.10.1.1	255.255.255.0	

API Equivalent

Cancel

Enable Multi-Cluster

Figure 39: Enabling NDB Multi-Cluster: With Network Segmentation

c. Click **Enable Multi-Cluster**.

A message indicating that an operation to enable NDB multi cluster has started is displayed. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

You can see the NDB agent details in the **NDB Service** page.

Registering a Nutanix Cluster with NDB

About this task

This task describes the steps to add a Nutanix cluster to NDB. With network segmentation, you can add one or more NICs to segment the network traffic on the NDB agent of the new cluster.

NDB requires Nutanix Guest Tools (NGT) on the source cluster for cross-hypervisor support (AHV to ESXi). When enabling multi-cluster, NGT is installed by default on the NDB agent VM of the AHV cluster where the NDB server exists.

Before you begin

- Ensure that you have enabled NDB multi-cluster (see [Enabling NDB Multi-Cluster](#) on page 91).

Caution:

- If the NDB server storage container and the storage container used by NDB for provisioning are different on the NDB server cluster, then you must configure mapping between the remote storage container and the storage container of the NDB server.
- If the database server VMs or the source databases contain disks from storage containers other than the NDB server storage container or the storage container used by NDB for provisioning, then you

must configure the storage mappings between the source and target clusters for such containers before performing any replication operation.

- If disaster recovery network segmentation is enabled on the cluster, then you must configure the remote sites and the corresponding storage container mappings between all the clusters.
- You cannot enable high availability for NDB if you configure network segmentation for NDB. Similarly you cannot configure network segmentation for NDB if you enable high availability for NDB.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Nutanix Clusters**, click **Register Nutanix Cluster**. The **Register Nutanix Cluster** window appears. You register a cluster in the following steps:
 - a. **Nutanix Cluster**
 - b. **Agent VM**
 - c. **Storage Container**

3. In the **Nutanix Cluster** step, do the following in the indicated fields.

- a. **Cluster Name.** Type a name for the Nutanix cluster you want to register.
- b. **Description.** Type a description for the cluster.
- c. **IP Address of Prism Element.** Type the Prism Element IP address.
- d. **Prism Element Administrator.** Type the username of the Prism Element administrator.
- e. **Password.** Type the Prism Element password.
- f. Click **Verify**.

NDB verifies the Prism Element and cluster information provided. The following fields are displayed if an ESXi hypervisor is found on the provided cluster.

- **Address.** Displays the vCenter IP address.
- **vCenter Administrator Username.** Type the username of the vCenter administrator.
- **Password.** Type the vCenter password.

g. Click **Next**.

Register Nutanix Cluster

1 Nutanix Cluster 2 Agent VM 3 Storage Container

Cluster Name
be

Description

IP Address of Prism Element
10

Prism Element Administrator
admin

Password

Verify

An ESXi hypervisor was found on this cluster. Provide vCenter credentials.

Address
10 9

vCenter Administrator Username
Ac

Password

Cancel Next

Figure 40: Registering a Nutanix Cluster

4. In the **Agent VM** step, do the following in the indicated fields.

a. **Name Prefix.** Type a prefix for the NDB agent name.

This field is optional. If you do not type a prefix, NDB automatically selects NDBAgent as a default prefix name.

b. Configure network segmentation to segment the network traffic of the agent VM.

Note:

- Network segmentation is only supported for static VLANs.
- Network segmentation is an advanced optional configuration; you can skip it if cluster segmentation is not configured in Prism Element.
- A network interface is attached to the NDB agent VM for every unique VLAN that you configure.

Select the VLAN access types for which you want to configure network segmentation:

- **Prism Element.** Select this VLAN access type to configure a VLAN that the NDB agent VM can use to communicate with Prism.

Note:

- Prism Element access VLAN is mandatory.
- By default, all the network traffic is directed to the Prism Element network.
- The network traffic of the VLAN access types that do not have a network segmentation configured is directed to Prism Element.
- This network will act as a default gateway for the NDB agent VM. This cannot be modified once the cluster registration is complete and the agent VM is created.

- **Prism iSCSI Data Service.** Select this VLAN access type to configure a VLAN that the agent VM can use to make connection requests to the iSCSI data services IP.

Note: NDB uses one of the segmented iSCSI data service IP addresses if you have configured multiple segmented data service IP addresses in Prism Element.

- **DBServer Access from NDB server.** Select this VLAN access type to configure a VLAN that is used for communications between the NDB agent VM and the database server VM on the newly registered NDB server cluster.
- **vCenter.** (Only for a ESXi cluster) Select this VLAN access type to configure a VLAN that the agent VM can use for vCenter API connections.

c. **DNS Servers.** Displays the list of DNS servers available.

d. **NTP Servers.** Displays the list of NTP servers available.

e. Click **Next**.

Register Nutanix Cluster X

1 Nutanix Cluster 2 **Agent VM** 3 Storage Container

Name Prefix
EraAgent

Agent Network Configuration

Agent Static IP Assignment Support:

- Static vLAN: Supported
- IPAM vLAN: Supported outside of managed pools
- DHCP vLAN: Not supported

The following section allows for configuring Advanced Network Segmentation. Prism Access vLAN is a mandatory input. Others can be set if the Prism Element 'EraAgent' has network segmentation enabled for the respective management services. A network interface will be attached to NDBCluster Agent VM for each of the **Distinct vLAN Names** provided in the below section.

vLAN Access Type	vLAN	Configuration		
		Gateway	Subnet Mask	IP Address
☒ Prism Element	Select a vLAN			
☐ Prism iSCSI Data Service				
☐ DBServer Access from NDB				
☐ vCenter				

DNS Servers: 10.10.10.16

NTP Servers: 0.0.0.0:rg

Back Next

Figure 41: Registering a Nutanix Cluster

- In the **Storage Container** step, select a storage container which is used for performing database operations in the cluster.
- Click **Register**.

The new cluster appears in the list of Nutanix clusters and a message appears indicating that the operation to register a cluster has started. Click the message to monitor the progress of the operation. Alternatively, select **Operations** from the drop-down list of the main menu to monitor the progress of the operation.

Select the name of the cluster to view the cluster status and database information. See [Updating a Nutanix Cluster](#) on page 97 for details on updating the cluster information. .

Updating a Nutanix Cluster

About this task

You can update the cluster information, Prism Element credentials, and the vCenter credentials of the registered cluster by doing the following.

Procedure

- Go to **Nutanix Clusters** and select the cluster which you want to update.
- From the drop-down list of the main menu, select **Administration**.

3. Click the **Cluster Actions** drop-down and select **Update**.

The **Update Cluster** window appears.

4. Do the following in the indicated fields.

- a. **Name of Cluster.** Type a new cluster name.
- b. **Description.** Type an updated description for the cluster.
- c. **Virtual IP Address.** Type an updated virtual IP address.
- d. **Storage Container.** This field cannot be updated.
- e. **Update Credentials.** Select this check box and do the following to update the Prism and vCenter credentials (for ESXi).
 - **Prism Administrator Username.** Type the updated Prism administrator username.
 - **Password.** Type the updated Prism administrator account password.

Note:

- For clusters imaged with AOS version 5.17.x or later, the default admin password expiration period (the maximum number of days between a password change) is changed from 99999 days to 60 days. For a Nutanix user, the number of days remains unchanged at 99999 days (no expiry).
- For clusters imaged with AOS versions older than 5.17.x, the password for the admin user does not expire. AOS upgrades to a version above 5.17.x do not affect the password expiration date for the admin user.
- In the event of a password expiry, the admin must change the password in Prism and configure the changed password in NDB. If NDB continues to use an expired password, Prism returns a

401 response with a password expiry message for all subsequent REST API requests. For more information on how to check the password expiry in Prism, see KB 4676.

- **vCenter IP Address.** Displays the vCenter IP address used in NDB.

Note: An alert icon is displayed for the vCenter IP address if the IP address displayed does not match the vCenter IP address in Prism Element.

- **vCenter Administrator Username.** Type the updated vCenter administrator username.
- **Password.** Type the updated vCenter administrator account password.

f. Click **Update**.

Update Cluster

Name of Cluster:

Description:

Cluster Virtual IP Address:

Storage Container:

☒ Update Credentials

Prism Administrator Username:

Password:

vCenter IP Address: 10.10.10.10

vCenter Administrator Username:

Password:

API Equivalent Cancel Update

Figure 42: Updating Nutanix Cluster

Setting VM Host Affinity

NDB sets affinity for the database server VMs that are part of a database server cluster in Prism to different nodes or hosts in order to ensure high availability in the event of any host failures.

About this task

NDB provides you the flexibility to manually set, unset, or clear the VM host affinity for AHV powered Nutanix clusters in Prism.

Perform the following procedure to set the VM host affinity of the a database server VM back to its initial pinned host.

Procedure

1. From the drop-down list of the main menu, select **Administration**.

2. Go to **Nutanix Clusters** and select the cluster for which you want to set affinity.
3. Click the **Cluster Actions** drop-down and select **Set VM-Host Affinity**.

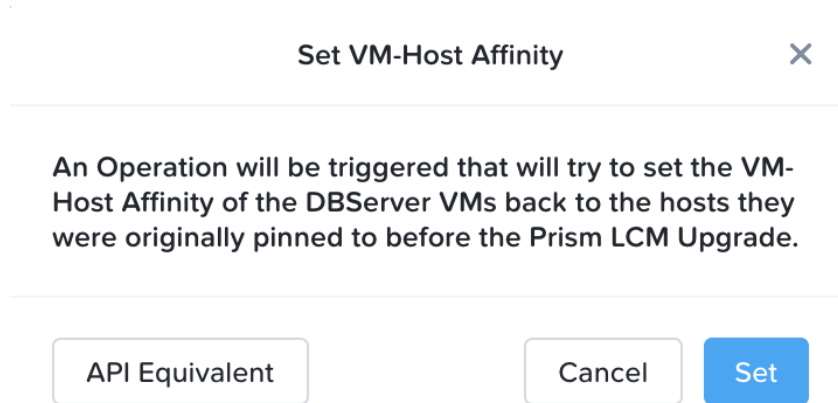


Figure 43: Set VM-Host Affinity

4. Click **Set** in the **Set VM-Host Affinity** confirmation window.
A message indicating the start of the operation is displayed. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

You can see the affinity details of the selected cluster in the **Nutanix Clusters** page.

UnSetting VM Host Affinity

About this task

Perform the following procedure to unset the VM host affinity of the a database server VM.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Nutanix Clusters** and select the cluster for which you want to set affinity.

3. Click the **Cluster Actions** drop-down and select **Unset VM-Host Affinity**.




Figure 44: Unset VM-Host Affinity

4. Click **Unset** in the **Unset VM-Host Affinity** confirmation window.
A message indicating the start of the operation is displayed. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

You can see the affinity details of the selected cluster in the **Nutanix Clusters** page.

Clearing VM Host Affinity

About this task

Perform the following procedure to clear or un-initialize the VM host affinity of the a database server VM.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Nutanix Clusters** and select the cluster for which you want to set affinity.
3. Click the **Cluster Actions** drop-down and select **Clear VM-Host Affinity**.
4. Click **Clear** in the **Unset VM-Host Affinity** confirmation window.
A message indicating the start of the operation is displayed. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

You can see the affinity details of the selected cluster in the **Nutanix Clusters** page.

NDB Network Management

The **Networks** page displays all the VLANs and IP address pools that are managed both in NDB and outside NDB.

NDB uses network profiles to associate VLANs with database server VMs. Therefore, you must first ensure that you have created VLANs on the Nutanix cluster, add the VLANs to NDB, and then create network profiles.

For Nutanix clusters, NDB also provides the flexibility to use AHV managed VLANs for static IP assignments. Using CLI, you can configure NDB to manage a set pool of IP addresses within the AHV managed IP address range.

You can add VLANs to NDB, update the VLANs, and add or remove IP address pools from the static VLANs from the **Networks** page. You can also create stretched VLANs that spans across multiple Nutanix clusters.




Figure 45: NDB Networks Page

Adding a VLAN to NDB

About this task

Perform the following procedure to add a VLAN to NDB.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Networks**, click **Add** > **Add vLAN**.
The **Add vLAN to NDB** window appears.
3. Do the following in the indicated fields.
 - a. **Nutanix Cluster**. Select the Nutanix cluster on which you want to add the VLAN.

Note: NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.

- b. **Select a VLAN**. Select a VLAN that you want to add to NDB.

Note: The VLAN appears in this list only if you have created the VLAN on the selected Nutanix cluster.

If the selected VLAN is IPAM (in the case of AHV Hypervisor), it cannot be managed by NDB using the NDB UI. Also, if you want to use DHCP to assign IP addresses to your database server VMs, then directly proceed to step 5.

4. (Static IP address assignment only) Click the **Manage IP Address Pool in NDB** option if you want to assign static IP addresses to your database server VMs and then do the following.
 - a. Type the gateway, subnet mask, primary DNS, and secondary DNS IP addresses and DNS domain for the VLAN configuration.
 - b. Type the IP address range for new database servers in the **First Address** and **Last Address** fields.

Note: NDB considers any VLAN added this way as static VLAN, irrespective of whether the VLAN is DHCP or static.

5. Click **Add** to add the VLAN to NDB.

?
✕

Add VLAN to NDB

Nutanix Cluster

mesarc0304
⌵

Select a VLAN ↻

☐
vlan_gpo

☐
default-vlan651-gpo

☐
default-vlan651-static

☒
Manage IP Address Pool

vLAN Configuration

Gateway

Verify

Subnet Mask

Primary DNS

Secondary DNS (optional)

DNS Domain (optional)

Initial Range of IP Addresses for New Database Servers

First Address

Last Address

API Equivalent

Cancel

Add

Figure 46: Adding VLAN to NDB

Updating a Static VLAN

About this task

Perform the following procedure to update a static VLAN.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
 2. Go to **Networks**, select the VLAN you want to update and click **Update**.
The **Update VLAN** window appears.
 3. Do the following in the indicated fields.
 - a. **Nutanix Cluster**. Select the Nutanix cluster on which you want to update the VLAN.
- Note:** NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
- b. **Select a VLAN**. Displays the VLAN selected for update.
 - c. Under **VLAN configuration**, update the gateway, subnet mask, primary DNS, and secondary DNS IP addresses and DNS domain for the VLAN configuration.

Note: You cannot update the gateway and subnet mask of the VLANs that belong to a stretched VLAN.

The screenshot shows the 'Update VLAN' window. The 'Nutanix Cluster' dropdown is set to 'Nutanix Cluster'. The 'Select a VLAN' list shows several VLANs, with 'v1' selected. The 'Manage IP Address Pool' checkbox is checked. The 'VLAN Configuration' section includes the following fields:

- Gateway: 1.1.1.1
- Subnet Mask: 255.255.255.0
- Primary DNS: 1.1.1.5
- Secondary DNS (optional):
- DNS Domain (optional):

The 'Verify' button is located below the Gateway field. At the bottom of the window, there are three buttons: 'API Equivalent', 'Cancel', and 'Update'.

Figure 47: Updating a VLAN

4. Click **Update** to update the static VLAN.

Adding IP Addresses to a Static VLAN

About this task

Perform the following procedure to add an IP address pool to a static VLAN.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Networks**, select the VLAN to which you want to add an IP address pool and, in the list of the IP address pools for that VLAN displayed on the right, click **Add**.




Figure 48: Add IP Address Pools

3. In the **Add IP Addresses** dialog box, type the IP address range in the **First Address** and **Last Address** fields.
4. Click **Add** to add the IP address pool to the static VLAN.

Displaying the Static IP Address Assignment

About this task

You can display the IP address to database server VM mapping of all the IP addresses belonging to an IP address pool.

Procedure

1. From the drop-down list of the main menu, select **Administration**.

- Go to **Networks**, select the VLAN whose IP address assignment you want to view and then select the IP address pool whose IP address assignment you want to view.

The IP address to database server VM mapping is displayed in a table as shown in the following screenshot.




Figure 49: IP Address Pool Assignment

Adding a Stretched VLAN to NDB

Before you begin

Add the required VLANs (which will be a part of the stretched VLAN) from each Nutanix cluster to NDB. See [Adding a VLAN to NDB](#) on page 102.

About this task

You can add a stretched VLAN to NDB by selecting the existing VLANs from each Nutanix cluster. Perform the following procedure to add a stretched VLAN to NDB.

Procedure

- From the drop-down list of the main menu, select **Administration**.
- Go to **Networks**, click **Add > Add Stretched vLAN**.
The **Add Stretched vLAN** window appears.

3. Do the following in the indicated fields.
 - a. **Stretched vLAN Name.** Type a name for the stretched VLAN.
 - b. **Description.** Type a description for the stretched VLAN.
 - c. Select the VLANs from the respective Nutanix clusters that you want to add to the stretched VLAN.

- Note:**
- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
 - The **vLAN** list only displays the VLANs added on the selected Nutanix cluster.
 - Only static VLANs that are managed in NDB can be added to a stretched VLAN.

4. Click **Add** to add the stretched VLAN to NDB.

? | ×

Add Stretched vLAN

Stretched vLAN Name

Description

Select the clusters with respective vLANs

	Name	Description	vLAN
☒	...		2.static ⌵ ⓘ
☒	E...		5.static ⌵ ⓘ
☒	S...		2.sti ⌵ ⓘ

Figure 50: Adding Stretched VLAN to NDB

Updating a Stretched VLAN

About this task

Perform the following procedure to update a stretched VLAN.

Procedure

1. From the drop-down list of the main menu, select **Administration**.

2. Go to **Networks**, select the stretched VLAN you want to update and click **Update**.
The **Update Stretched vLAN** window appears.

3. Do the following in the indicated fields.
- a. **Stretched vLAN Name.** Update the name for the stretched VLAN.
 - b. **Description.** Update the description for the stretched VLAN.
 - c. Update the stretched VLAN **Gateway** and **Subnet Mask** IP address. Click **Verify** to verify the status of the IP addresses.
- Updating the gateway and subnet mask of the stretched VLAN also updates the gateway and subnet mask of the VLANs that belong to that stretched VLAN.
- a. Add or remove any new or existing Nutanix cluster and the respective VLAN, or update the existing VLAN to the stretched VLAN.

Note:

- The **vLAN** list only displays the VLANs added on the selected Nutanix cluster.
- Only static VLANs that are managed in NDB can be added to a stretched VLAN.
- You cannot remove an existing VLAN from the stretched VLAN if that stretched VLAN is used by any profile in NDB.

Update Stretched vLAN

Stretched vLAN Name

Description

Description

Gateway

1

Subnet Mask

2

Verify

Select the clusters with respective vLANs

☐	Name	Description	vLAN
☒	b		v.static ⓘ
☒	E		c.2.st ⓘ
☐	g		d.2.st ⓘ

API Equivalent

Cancel

Update

Figure 51: Updating a Stretched VLAN

NUTANIX

Nutanix Database Service (formerly Era) | NDB Administration | 109

4. Click **Update** to update the stretched VLAN.

Adding an AHV-managed VLAN as NDB-managed

This topic describes the steps to add an AHV-managed VLAN as NDB-managed. This operation is supported only on the CLI.

About this task

This task describes how to use an AHV-managed subnet (on-premises or in NC2) for static IP address allocation where NDB performs IP address management (IPAM) instead of AHV. Perform the following steps to add an AHV-managed subnet to NDB as an NDB Managed VLAN to use the subnet as a pool of static IP addresses for database servers.

Caution: Nutanix does not recommend adding an AHV-managed VLAN as NDB-managed as this may lead to misconfigurations of IP address pools, causing provisioning operations to fail. However, this configuration is supported for use cases such as NC2 on AWS.

Procedure

1. For this procedure, assume a need for 200 IP addresses managed by NDB on the 10.10.10.0/24 subnet while AHV manages the remaining. From the Nutanix Prism Element (PE), create an AHV-managed VLAN. While creating the VLAN, stipulate the IP address range that AHV manages by specifying the start and end addresses of the IP address pool. For example, specify the range from 10.10.10.201 to 10.10.10.253 for AHV-managed IP addresses.
2. SSH into the NDB server and enter the 'era' CLI prompt.
3. Fetch the Nutanix cluster ID for the PE by using the command:

```
era> cluster list
```

4. Use the following command to add the VLAN to NDB as an NDB-managed VLAN while providing the IP address range until 10.10.10.200:

```
era> resource network add nx_cluster_id=cluster_id name=vlan_name ip_managed_in_era  
gateway=default_gateway subnet_mask=subnet_mask primary_dns=primary_dns  
start_ip=10.10.10.1 end_ip=10.10.10.200
```

- **resource.** Specify this parameter to manage resources in NDB.
- **network.** Specify this parameter to manage network resources.
- **add.** Specify this action to add a VLAN to NDB.
- **nx_cluster_id.** Specify the Nutanix cluster ID from step 3.
- **name.** Specify the VLAN name in PE.
- **ip_managed_in_era.** Specify this parameter to manage IP address pools in NDB.
- **gateway.** Specify the default gateway of the network.
- **subnet_mask.** Specify the subnet mask of the network.
- **primary_dns.** Specify the address of the primary DNS server.
- **start_ip.** Specify the first IP address in the range of IP addresses.
- **end_ip.** Specify the last IP address in the range of IP addresses.

To verify the configuration, go to **Administration** and select **Networks**. The newly added VLAN appears.

5. You can now create a network profile for your database engine of choice and use that network profile for provisioning databases using static IP addresses.
6. If the NDB instance is managing multiple Nutanix PE clusters, repeat steps 1 through 4 for every PE cluster on which you need to provision databases.

Configuring NTP Servers

This topic describes the steps to add or remove NTP servers for the NDB server and NDB agent.

Procedure

1. Launch the era-server CLI.

Configuring NTP servers for the NDB server

2. Run the following command to add a new NTP server for the NDB server:

```
era-server > era_server set ha-enabled=false ntp_server=ntp_server
```

Replace *ntp_server* with the IP address of the NTP server.

3. (Optional) To add multiple NTP servers for the NDB server, run the following command:

```
era-server > era_server set ha-enabled=false ntp_server=ntp_server1  
ntp_server=ntp_server2 ntp_server=ntp_server3
```

Set *ha-enabled* as *false* if NDB high availability is not enabled. Replace *ntp_server1*, *ntp_server2*, and *ntp_server3* with the IP addresses of the NTP servers.

4. Run the following command to remove an NTP server of the NDB server:

```
era-server > era_server delete ha-enabled=false ntp_server=ntp_server
```

Configuring NTP servers for the NDB agent

5. Get the Nutanix cluster ID by using the command:

```
era-server > cluster list
```

6. Run the following command to add a new NTP server for the NDB agent:

```
era-server > cluster agent set nx_cluster_id=nx_cluster_id ntp_server=ntp_server
```

Replace *ntp_server* with the IP address of the NTP server and *nx_cluster_id* with the Nutanix cluster ID from step 5.

7. (Optional) To add multiple NTP servers for the NDB agent, run the following command:

```
era-server > cluster agent set nx_cluster_id=nx_cluster_id ntp_server=ntp_server1  
ntp_server=ntp_server2 ntp_server=ntp_server3
```

Replace *ntp_server1*, *ntp_server2*, and *ntp_server3* with the IP addresses of the NTP servers.

8. Run the following command to remove an NTP server of the NDB agent:

```
era-server > cluster agent delete nx_cluster_id=nx_cluster_id ntp_server=ntp_server
```

Configuring DNS Servers

This topic describes the steps to add or remove DNS servers for the NDB server and NDB agent.

Procedure

1. Launch the era-server CLI.

Configuring DNS servers for the NDB server

2. Run the following command to add a new DNS server for the NDB server:

```
era-server > era_server set ha-enabled=false nameserver=nameserver
```

Replace *nameserver* with the IP address of the DNS server.

3. (Optional) To add multiple NTP servers for the NDB server, run the following command:

```
era-server > era_server set ha-enabled=false nameserver=nameserver1  
nameserver=nameserver2 nameserver=nameserver3
```

Set *ha-enabled* as *false* if NDB high availability is not enabled. Replace *nameserver1*, *nameserver2*, and *nameserver3* with the IP addresses of the DNS servers.

4. Run the following command to remove a DNS server of the NDB server:

```
era-server > era_server delete ha-enabled=false nameserver=nameserver
```

Configuring DNS servers for the NDB agent

5. Get the Nutanix cluster ID by using the command:

```
era-server > cluster list
```

6. Run the following command to add a new DNS server for the NDB agent:

```
era-server > cluster agent set nx_cluster_id=nutanix_cluster_id nameserver=nameserver
```

Replace *nameserver* with the IP address of the DNS server and *nutanix_cluster_id* with the Nutanix cluster ID from step 5.

7. (Optional) To add multiple DNS servers for the NDB agent, run the following command:

```
era-server > cluster agent set nx_cluster_id=nutanix_cluster_id  
nameserver=nameserver1 nameserver=nameserver2 nameserver=nameserver3
```

Replace *nameserver1*, *nameserver2*, and *nameserver3* with the IP addresses of the DNS servers.

8. Run the following command to remove an NTP server of the NDB agent:

```
era-server > cluster agent delete nx_cluster_id=nutanix_cluster_id  
nameserver=nameserver
```

Configuring the Time Zone

This task describes the steps to configure the time zone of the NDB server.

Procedure

1. Launch the era-server CLI.

2. Run the following command to configure the time zone of the NDB server:

```
era-server > era_server set timezone=timezone
```

Replace *timezone* with the required time zone. Press the <Tab> key to list and dynamically fill the supported time zones.

Configuring the NDB Proxy

This task describes the steps to configure the NDB proxy allowing NDB to forward any outbound HTTP/HTTPS requests through it. You can configure the NDB proxy to report the telemetry data through Pulse, if you do not allow direct outbound internet access from your NDB server.

About this task

Perform the following steps to configure a proxy for NDB.

Procedure

1. Launch the NDB CLI.
2. Issue the following command to enable proxy for NDB:

```
era > proxy update url=url port=port enabled=true
```

Replace *url* with the URL or the IP address of the proxy and *port* with the port number.

3. To verify the configuration, use the following command:

```
era > proxy show
```

```
ERA Proxy :
+-----+-----+
| Parameter | Value |
+-----+-----+
| ERA Proxy URL | 10.10.10.10 |
+-----+-----+
| ERA Proxy Port | 3128 |
+-----+-----+
| ERA Proxy Type | HTTPS |
+-----+-----+
| ERA Proxy Enabled | 1 |
+-----+-----+
```

4. To disable the NDB proxy, issue the following command:

```
era > proxy update url=url port=port enabled=false
```

Replace *url* with the URL or the IP address of the proxy and *port* with the port number.

NDB Tags Management

NDB allows you to assign metadata to entities (clones, time machines, databases, and database servers) by using tags. When you are cloning a database, you can associate tags with the database that you are creating.

Each tag is a label consisting of a user-defined name and a value that makes it easier to manage, search for, and filter entities. For example, if you search for a tag value 187a, entities that have 187a as a tag value are listed. Tags enable you to organize and categorize entities by purpose, owner, or other criteria. This organization is useful if you have many entities of the same type. You can quickly identify a specific entity based on the tags that you have assigned to it.

Tags help you browse entities based on the tags associated and to add or remove tags from multiple entities at once. You can add, update and delete tags to databases, clones, time machines, and database servers. If an entity is associated with a tag, the respective entity window displays a search bar. You can click the search bar and select the tag to view the list of entities associated with the selected tag.

Creating a Tag

About this task

Perform the following procedure to add a tag to an entity (clone, time machine, database, and database server VM).

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Tags** and click **Create**. The **Create Tag** window appears.
3. Do the following in the indicated fields.
 - a. **Name**. Type a name for the tag.
 - b. **Description**. Type a description for the tag.
 - c. **Add Tag to New and Existing**. Select an entity for the tag to be associated with. You can select one of the following.
 - » **Oracle/SQL Server Databases, PostgreSQL/MariaDB/MySQL Instances**
 - » **Time Machines**
 - » **Clones**
 - » **Database Server VMs**
 - d. **Require a Tag Value**. Select this check box if you want to provide a tag value.
Select this check box and provide a value for any new entity with this tag.

The screenshot shows the 'Create Tag' dialog box. It has a title bar with the text 'Create Tag' and icons for help (?) and close (X). The dialog is divided into two main sections. The left section has a label 'Name' above a text input field. Below this input field is a checkbox labeled 'Require a Tag Value'. The right section has a label 'Description' above a larger text input field. Below the description input field is a section titled 'Add Tag to New and Existing' which contains five radio button options: 'Oracle/SQL Server Databases', 'PostgreSQL/MariaDB/MySQL/MongoDB Instances', 'Time Machines', 'Clones', and 'Database Server VMs'. At the bottom of the dialog, there is a button labeled 'API Equivalent' on the left, and two buttons labeled 'Cancel' and 'Create' on the right.

Figure 52: Creating a Tag

4. Click **Create** to create a tag.

After the tag is created successfully, it appears in the list of created tags.

Updating a Tag

About this task

Perform the following procedure to update an existing tag.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Tags**. Select the tag and click **Update**. The **Update Tag** window appears.

Update Tag

Name Description

☒ Require a Tag Value

☒ Add/Update Tag on Oracle/SQL Server Databases
PostgreSQL/MariaDB/MySQL/MongoDB Instances

API Equivalent Cancel Update

Figure 53: Updating a Tag

3. Update the indicated fields.
 - a. **Name.** Type an updated name for the tag.
 - b. **Description.** Type an updated description for the tag.
 - c. **Add/Update Tag on '<entity name>'.** Clear this check box if you want to deprecate the tag of its associated entity.
 - d. **Require a Tag Value.** Clear this check box if you want to remove the mandatory tag value.
4. Click **Update** to update a tag.

Deleting a Tag

Before you begin

Delete the tag values for the entities (database, clone, time machine, database server VM) associated with a tag.

About this task

Perform the following procedure to delete a tag.

Procedure

1. From the drop-down list of the main menu, select **Administration**.
2. Go to **Tags**. Select the tag and click **Delete**. This step displays a confirmation dialog box.

3. Click **Yes** to confirm and delete a tag.

NDB Upgrade Management

You can upgrade to a new version of NDB to use the latest available features. As part of the upgrade process, the NDB server is upgraded first, followed by the NDB agents, database server VMs, HAProxy VMs, and repository VMs in parallel.

The **Dashboard** displays the current version and an option to upgrade the NDB server to a latest version when a new version is available for upgrade.




Figure 54: Upgrade Version Details

You can upgrade NDB in the following two ways:

- One-click Upgrade - If the NDB server is already connected to the internet, the system automatically detects the latest version and you can perform the one-click upgrade operation (see [Upgrading NDB \(One-Click Upgrade\)](#) on page 116).
- Offline Upgrade (CLI) - If the NDB server is not connected to the internet, you can upgrade the server in offline mode by performing the manual CLI upgrade operation (see [Upgrading NDB \(Offline Upgrade\)](#) on page 117).

Both methods require some application downtime to ensure that NDB loses no data during the upgrade.

Upgrading NDB (One-Click Upgrade)

If the NDB server is already connected to the internet, the system automatically detects the latest version and you can perform the one-click upgrade operation.

About this task

Perform the following procedure to upgrade NDB to a latest available version by using the one-click upgrade method.

Procedure

1. From the drop-down list of the main menu, select **Administration > Maintenance**.

Alternatively, on the **Dashboard**, click the **Upgrade Available** option.

2. Do one of the following:

- » Download the latest available version.
- » Upload an upgrade bundle.

3. (Download the latest available version). Select the check box for the version that you want to upgrade to, and click **Download**.

Wait until NDB completes the download of the new version from the Nutanix Support Portal.

NDB Version
Current Version 1.0.0 Last Modified January 1, 2019

Available Upgrades

Version	Released	Status
☒ 1.0.0.1	2018-12-06 09:49:10	Downloaded
☐ 1.0.0.2	2019-01-03 06:25:24	Available

Upload an Upgrade Bundle

Figure 55: Downloading a Latest Available Version

4. (Upload an upgrade bundle). Do the following:
 - a. Download the upgrade bundle from the Nutanix Support portal.
 - b. Under **Upload an Upgrade Bundle**, click **Upload** and select the file that you have downloaded.
5. Select the version that you want to upgrade to, and click **Upgrade Now**.

A message appears indicating that the operation to upgrade NDB has started. Click the message to monitor the progress of the operation. Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation. You cannot perform any actions in the NDB until the upgrade is complete.

After the upgrade is successful, NDB refreshes the page to reflect the latest version.

For more information, see [NDB Service Management](#) on page 57.

Upgrading NDB (Offline Upgrade)

If the NDB server is not connected to the internet, you can upgrade the server in offline mode by performing the manual CLI upgrade operation.

About this task

Perform the following procedure to upgrade NDB to a latest available version by using the offline upgrade (CLI) method.

Procedure

1. Download the upgrade bundle from the Nutanix Support portal.
2. Upload the downloaded file to the NDB server that you want to upgrade.

- ```
era > upgrade upload_bundle path=path_to_upgrade_bundle
```

- Note:** You cannot perform any actions in the NDB until the upgrade is complete.

If you are experiencing issues with NDB or with any of the operations, or with the entities registered with NDB, you can download a diagnostic bundle and then share the diagnostics bundle with Nutanix Support for troubleshooting the issues. The diagnostics bundle helps Nutanix Support to further investigate the problem.

There are two ways to download the diagnostics bundle:

1. You can generate a diagnostic bundle for any of the operations performed in NDB from the **Operations** page. Note that bundle generation from the **Operations** page provides the bundle only for the VMs impacted by a specific operation. For more information, see *Viewing Operation Logs* section in this guide.
2. You can download the diagnostic bundle from the **Administration** menu. This is a generic method that allows you to generate the bundle by selecting specific VMs when you encounter an issue. Perform the following procedure to download the diagnostic bundle from the **Administration** menu.



1. From the drop-down list of the main menu, select **Administration** > **Diagnostics**.
2. Select the NDB server, an entity, or both for which you want to download the diagnostics bundle.

3. Click **Generate Bundle**.

Generate Bundle ? | X

The diagnostic bundle is ready to download.

Generating bundle on 1 VM

☒ **Generate** Generate

Select the number of days you want to generate the server bundle.

☐ Since Beginning

☒ 1 Days ▾

☒ Completed (1 / 1)  
[Show more](#)

☒ **Download** Download

Bundle will be downloaded in a zipped folder. You can share the downloaded archive to the support team for assistance.

API Equivalent Close

**Figure 57: Downloading a Diagnostic Bundle**

4. Select the time period for which you want to generate the server bundle and click **Generate**.  
Wait until NDB completes the generation of the diagnostics bundle and click **Download**.  
You can then share the diagnostics bundle with Nutanix Support.

## Downloading the Diagnostics Bundle using the NDB CLI

### About this task

Perform the following procedure to generate and download a diagnostic bundle using the NDB CLI.

### Procedure

1. SSH into the NDB server using the following default login credentials.  
Username: *era*  
Password: *Nutanix.1*

2. Launch the NDB CLI and run the following command.

```
era > diagnostic_bundle VM Options
```

Replace VM Options with one of the following.

- local. Generates diagnostic bundle for the current VM
- all. Generates diagnostic bundle for all VMs and the NDB Server.
- era\_server. Generates diagnostic bundle for the NDB Server.
- VMs. Generates diagnostic bundle for specific VMs.

## Operation Logs

NDB allows you to view the log files for operations such as provisioning, registration, cloning, refresh, upgrade, and patching of databases. For each operation, NDB provides a list of log files classified according to the NDB component or server for which the message is logged. You can view the log messages in the UI or download the log files to your system to view them separately. The **Operations** page provides an intuitive representation of each operation, the subtasks involved in the operation, and their statuses.

Operation logs can help you identify and troubleshoot issues for an operation. You can also download and share the logs with Nutanix Support for help with investigating an issue further.

## Viewing Operation Logs

You can view and download the log files for each operation in NDB.

About this task

Perform the following procedure to access and view the logs for an operation.

Procedure

1. From the main menu, select **Operations**.
2. Select the desired filters for the operation from the left panel.

**Note:** NDB displays the **Current/Recent** operations by default. Select **Scheduled** to display the scheduled tasks.

3. From the list of operations that appears, select the operation you want to view.

Operations

1 Total Operations

Current/Recent

Scheduled

Filter

Nutanix Cluster

All

From Last

1 month

Status

All

Name

All

Owner

All

Show System Operations

Refresh Data

every 5 seconds

Stop Resubmit

Name Entity Nutanix Cluster Owner Status Start Time

Refresh Clone

Default\_706\_2022\_Mar\_08\_07\_52\_45

taurus03

admin

70%

2022-03-08 10:10:33

< 1 >

Refresh Clone - 70%

Generate Bundle Show Logs

Duration: 1 minute 23 seconds End Time: 2022-03-08 10:12:07 Operation ID: 39802bade-279b-4a4b-b491-83c4da0327a2

Operation Name: Refresh SQL Server Database Group from Snapshot d:\Ndbase\6274-4278-WinR-394a4bdc-72

Message: manual failure injection

Status: Failed

| Name                                | Status | Start Time          | End Time            | Duration   |
|-------------------------------------|--------|---------------------|---------------------|------------|
| Pre-processing                      | ✓      | 2022-03-08 10:10:38 | 2022-03-08 10:10:38 | < 1 second |
| Shutdown Databases and Detach disks | ✓      | 2022-03-08 10:10:40 | 2022-03-08 10:10:56 | 26 seconds |
| Restore Databases from Snapshot     | ✓      | 2022-03-08 10:10:52 | 2022-03-08 10:10:55 | 23 seconds |
| Recover Databases                   | ✓      | 2022-03-08 10:10:55 | 2022-03-08 10:10:56 | 1 second   |
| Discover Clone Databases            | ✗      | 2022-03-08 10:10:56 | 2022-03-08 10:10:56 | < 1 second |
| Rollback                            | ✓      | 2022-03-08 10:10:56 | 2022-03-08 10:12:07 | 22 seconds |

Figure 58: Operation Selection

Detailed information on the subtasks involved in the operation appears. This includes the name of the subtask, status, start time, end time, and duration.

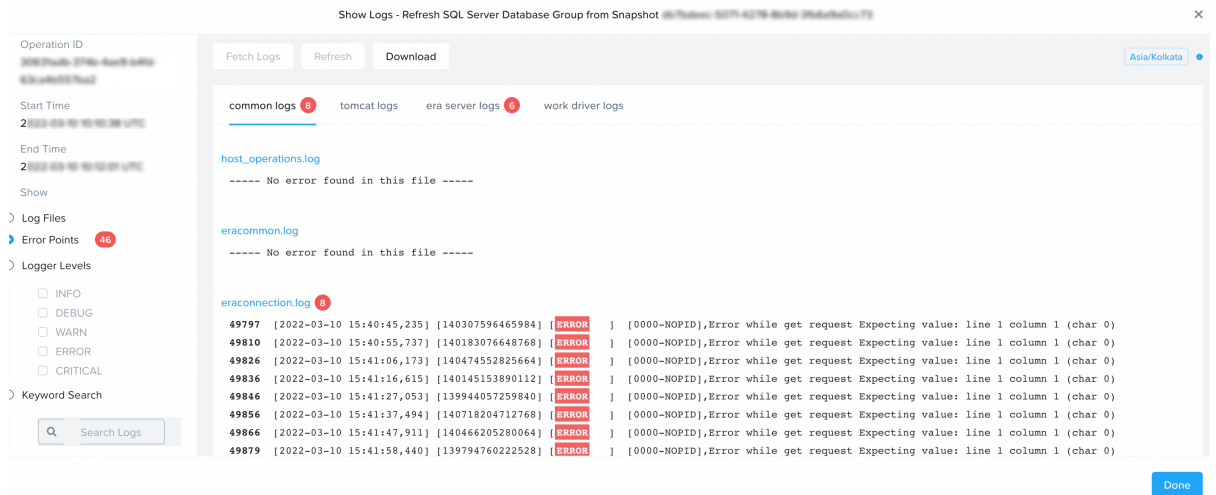
4. If you want to generate and download the diagnostic bundle, select **Generate Bundle**.

a. In the **Generate Bundle** window that appears, click **Generate**.

b. Wait for the generation to finish and click **Download**.

The diagnostic bundle is downloaded onto your system in a zipped folder.

- To view the logs, click **Show Logs**.



**Figure 59: Operation Logs**

**Note:** If log files do not appear, click **Fetch Logs** and click **Refresh**.

The Show Logs page appears with the logs in separate tabs such as common logs, Tomcat logs, NDB server logs, and work driver logs.

- If you want to download the log files, click **Download**.  
The log files are downloaded onto your system in a zipped folder.
- If you want to view a specific log file, click the log file.  
The log information appears on a new page.

## Show Logs Page

The Show Logs page displays the logs and the associated error messages for an operation. You can apply filters to view the logs in different ways for different entities and can also search for specific log messages.



**Figure 60: Show Logs Page**

The left panel allows you to configure the following filters for viewing the log messages:

- **Show Menu.** The show menu provides you with an option to view the logs in three different ways: Log Files, Error Points, and Logger Levels.  
Log Files: Lists the log files only.  
Error Points: Lists the log files along with the related error messages. Error Points is selected by default and this view provides insight into a failed operation.  
Logger Levels: Allows you to display the log messages based on the severity levels. The levels that you can select are: **INFO**, **DEBUG**, **WARN**, **ERROR**, and **CRITICAL**.
- **Keyword Search.** The keyword search enables you to find log messages by entering specific keywords that appear in the log messages.
- **Servers in Operation.** The Servers in Operation menu allows you to select the server you want to view the log files for. For example, for a database provisioning operation, you can view the logs for either the NDB server or the database server VM.

## Show Logs Tabs

The Show Logs page organizes the log files and messages into various tabs. Within a tab, you can click on a log file to view the logged information in a new window. The tabs that appear for an operation are dependent on the type of operation and the server you choose to view the logs for.

**Table 19: Show Logs Tabs**

| Name                  | Description                                                                                                                                                                          |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Common logs           | Displays the common log files for host operations, NDB connections, and NDB.                                                                                                         |
| Operation logs        | Displays logs and errors associated with the operation.                                                                                                                              |
| Operation report      | Displays information regarding an operation in an organized and intuitive manner. This includes the operation name, operation status, failed time, failed step, and operation error. |
| Work driver logs      | Displays the logs related to the work driver.                                                                                                                                        |
| Tomcat logs           | Displays the logs for the Tomcat server.                                                                                                                                             |
| NDB server logs       | Displays the logs for the NDB server.                                                                                                                                                |
| Zipped operation logs | Displays zipped operation log files. These log files generally include rollback bundles from failed operations.                                                                      |

## Maintenance Window

A maintenance window allows you to set a schedule that is used to automate repeated maintenance tasks such as OS patching and database patching. NDB allows you to create a maintenance window and then associate the maintenance window with a list of database server VMs or an instance. You can create a weekly or monthly schedule for a maintenance window.

### Maintenance Window States

A maintenance window has three states:

- **Active.** The state when a maintenance window is associated with a database server or an instance but is not yet scheduled.
- **Scheduled.** The state when the process is scheduled for the upcoming task such as operating system patching or database patching.
- **Running.** The state when the task (operating system patching or database patching) starts. It remains in this state until the task completes.

## Creating a Maintenance Window

This procedure shows how to create a maintenance window.

About this task

Only a Super Administrator, Infra Administrator, or Database Infra Administrator has the privilege to create a maintenance window.

Procedure

1. From the Dashboard, select **Policies > Maintenance Window**.
2. In the **Maintenance Window** page, click **Create** and do the following in the indicated fields.

The screenshot shows a 'Create Maintenance Window' dialog box. It contains the following fields and values:

- Maintenance Window Name:** Weekly-OS-Patch
- Description:** Patching the OS per week
- Schedule:** Weekly
- Repeat On:** Friday (selected)
- Time:** 15:28:20
- Duration (hours):** 2
- Will end current day:** 17:28:20 IST u

At the bottom right, there are three buttons: 'Cancel', 'API Equivalent', and 'Create'.

**Figure 61: Creating a Maintenance Window**

- a. **Maintenance Window Name.** Type a name for the maintenance window.
- b. **Description.** Type a description for the maintenance window.
- c. **Schedule.** Select either a Weekly or Monthly schedule.
  - If you select Weekly, select the day and time when the maintenance window triggers.
  - If you select Monthly, select a week (1st, 2nd, 3rd, or 4th ), day of the week, and a time when the maintenance window triggers.

### 3. Click **Create**.

The newly created maintenance window is listed in the Maintenance Window page. You can associate the maintenance window to any database instance or a database server VM from the database server list page.

## Associating a Maintenance Window

### About this task

This procedure shows how to associate a maintenance window with a database instance, database server VM, or a cluster.

### Procedure

1. From the main drop-down, click **Database Server VMs > List**.
2. Select the database instance, database server VM, or a cluster and then click **Actions > Maintenance Window**. You can select multiple database instances, database server VMs or clusters.
3. In the **Maintenance** page, do the following in the indicated fields.
  - a. Select **Operating System Patching** or **Database Patching** or both.
  - b. Select the maintenance window to associate from the **Maintenance Window** drop-down.
  - c. (Optional) Click **Advance Options** and add **Pre** and **Post** commands.
  - d. Click **Associate**.

## Operating System Patching

NDB supports patching of database server VM operating systems for Oracle, PostgreSQL, and MongoDB. NDB can apply on-demand and scheduled operating system patches across database server VMs running Linux operating system. This feature enables you to keep the underlying OS up to date with the latest security patches.

To patch the Linux OS, NDB issues the update command for the package manager native to each OS. For example, if you use RHEL 7.x, NDB runs YUM and if you use Centos 8.x, NDB runs DNF. For OS patching, the software profile used to create the database server VM must include the proper software repository configuration such that a YUM update (in the case of Centos 7) performs the required actions. You can manage what updates are available to a database server VM by curating the contents of the repository available to the VM.

If the database is in an HA configuration, NDB performs a rolling upgrade during OS patching. The system patches the secondary database server VM first, which then becomes the new primary. NDB patches the old primary database server VM soon after. If a database server VM is a part of an HA cluster, then each database server VM has the same maintenance window. The nodes inherit the maintenance window from the cluster.

You can perform operating system patching in two ways:

- **Manual.** For more information, see *Patching an Operating System*.
- **Automated.** You can enable automatic OS patching while provisioning or registering a database and database server VM.

### Rollback

You can roll back a database server VM if the OS patch operation fails.

- **Automatic rollback.** NDB triggers automatic rollback if the OS patch operation fails. Automatic rollback is only available for OSES using the YUM package manager to apply the OS patches.

- **Manual rollback.** NDB takes a database snapshot before starting a patching operation. If the OS patch operation fails, you must manually restore using this snapshot.

If the NDB OS patch operation is successful but you encounter a problem with the applied patches, you need to roll back the applied patches manually from within the database server.

**Note:** You must contact Nutanix Support before performing a manual rollback using the snapshot.

## Patching an Operating System

Before you begin

- You must configure the repository (from which the package managers such as YUM or APT fetch the updates) in the gold image VM used to create or update the software profile. Not configuring the repositories can cause unexpected behaviors.
- Ensure that no database-related software is included in the software repository to prevent any conflicts with database installation.


About this task

This procedure shows how to patch a Linux operating system.

Procedure

1. From the main menu, select **Database Server VMs**.
2. Go to **List** and click the database server VM that you want to update.

The database server VM summary page appears.



**Figure 62: Database Server VMs page displaying a PostgreSQL Server VM**

3. Go to **Actions** and click **Patch Now**.

4. (Optional) Select **Advanced Options** and do the following:
  - a. **Pre Operating System Patching Command.** Type a complete OS command that you want to run before patching the OS.
  - b. **Post Operating System Patching Command.** Type a complete OS command that you want to run after patching the OS.
5. Click **Update**.

NDB applies all patches available in the configured repository.

# DATABASE MANAGEMENT

---

## Updating a Database

The operation to update a database allows you to update the name, description and the associated tag values of the database.

About this task

Perform the following procedure to update a database.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Sources**, select the database that you want to update, and click **Update**.  
The **Update Database** window appears.
3. Do the following in the indicated fields:
  - a. **Name**. Update the name of the database.
  - b. **Description**. Update the description of the database.
  - c. **Tags**. Displays the associated tags and tag values of the database. You can update the associated tag values. See [NDB Tags Management](#) on page 113 for more information.
  - d. Click **Update**.

NDB displays a message indicating that the operation to update the database is successfully completed.

## Removing a Database

The operation to remove a database deletes the database and the time machine that is associated with the database. NDB allows you to delete only greenfield databases.

About this task

Perform the following procedure to remove a database.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Sources**, select the database that you want to remove, and click **Remove**. The **Remove Database** window appears.

3. Do the following in the indicated fields:

- a. **Unregister the database from NDB.** By default, this option is selected.

**Note:** NDB deregisters the database when you perform remove operation.

- b. **Delete source database from the VM.** Select this check box if you want to delete the source database from the VM.

**Note:** This option is not available for brownfield databases.

- c. **Delete the database's Time Machine (snapshots/logs) from the NDB.** Select this check box to delete the time machine that is associated with the database that you want to delete.
- d. Under **Confirm this request to remove the database, <name of the database>, by providing the name of the database**, type the name of the database that you want to remove
- e. Click **Remove**.

NDB displays a message indicating that the operation to remove a database has started. Click the message to monitor the progress of the operation.

Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

**Note:** When the database remove operation is started, NDB checks if the data drives associated with the database being deleted are shared with other application or clones. If drives are not shared with other applications or clones on that host, NDB deletes the data drives.

## Scaling a Database

NDB provides you the flexibility to extend or scale up the storage capacity of the NDB-managed databases. Scaling of a database functions differently for each database engine. NDB expands the data area and log area for SQL Server databases and PostgreSQL instances. The data area and fast recovery area expands for Oracle databases.

See [Scaling an Oracle Database](#), [Scaling a SQL Server Database](#), and [Scaling a PostgreSQL Database](#) for more information.

## Updating a Database Server VM or Database Server Cluster

This task describes the steps to update the name, description, and NDB password of a database server VM or cluster.

About this task

Perform the following procedure to update the name, description, and password for a database server VM or cluster.

Procedure

1. In the drop-down list of the main menu, select **Databases Server VMs**.
2. Go to **List** and select the database server VM or cluster that you want to update.

3. Click **Actions** and select **Update**.

The **Update Database Server VM** or **Update Database Server Cluster** window appears.

Update Database Server VM

Name

ndb\_postgres\_db\_VM

Description

dbserver for ndb\_postgres\_db

☒ Update in Nutanix Cluster

☐ Update Credentials

| Type           | User Name | Password |  |
|----------------|-----------|----------|--|
| NDB Drive User | era       |          |  |

Credentials updates will only be applied to NDB. Changes on the DB Server VM should be made outside of NDB.

API Equivalent

Cancel

Update

**Figure 63: Update Database Server VM**

4. Do the following in the indicated fields:

- Name.** Enter a name for the database server VM or cluster.
- Description.** Enter a description for the database server VM or cluster.
- For database server VMs, check the **Update in Nutanix Cluster** box to update the name and description of the database server VM in the Nutanix cluster.
- For database server VMs, select **Update Credentials** and enter the new password for the NDB drive user in the **Password** column.

**Note:** This password change applies only to NDB. To make the changes on the database server VM, update the password outside NDB.

e. Click **Update**.

NDB displays a message indicating that the operation to update the database server VM or cluster is complete.

## Updating the IP Address of a Registered Database Server VM

### About this task

You can manually update the IP address of a database server VM if the IP address has changed after you register the database server VM in NDB.

Perform the following procedure to update the IP address of a registered database server VM.

**Note:** Perform the following steps only on source database server VMs and not on clone database server VMs.

## Procedure

1. Launch NDB CLI.

```
era >
```

2. Get database server VM details.

```
era > dbserver list
```

3. Update the database server VM.

```
era > dbserver update id=dbserver_id associate_new_ip ip_address=new_ip_address
```

Replace *dbserver\_id* with the ID of the database server VM.

Replace *new\_ip\_address* with the IP address of the database server VM.

**Note:** You must assign unique IP addresses for database server VMs.

## Removing a Database Server VM or Database Server Cluster

### About this task

Perform the following procedure to remove a database server VM or a database server cluster. NDB allows you to delete only greenfield database server VMs.

### Before you begin

- Ensure that you remove the database before removing a database server VM. For more information, see [Removing a Database](#) on page 128.

## Procedure

1. In the drop-down list of the main menu, select **Databases Server VMs**.
2. Go to **List**, select the database server VM or database server cluster that you want to remove, and click **Remove**.  
(For database server VM) The **Remove Database Server VM** window appears.  
(For database server cluster) The **Remove Server Cluster** window appears.

3. Do the following in the indicated fields:

- a. (For database server VM) **Unregister the database server VM from NDB.** By default, this option is selected.

(For database server cluster) **Unregister server cluster from NDB.** By default, this option is selected.

**Note:** NDB de-registers the database server VM or database server cluster when you perform remove operation.

- b. **Delete the VM and associated storage.** Select this check box if you want to delete the VM and associated storage.

**Note:** This option is disabled for a brownfield database server.

- c. Type the name of the database server VM or database server cluster that you want to remove as a confirmation.

- d. Click **Remove**.

NDB displays a message indicating that the operation to remove a database server VM or a database server cluster has started. Click the message to monitor the progress of the operation.

Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

## Vertical Scaling of Database Server VMs

This task describes the CLI steps to perform vertical scaling of database server VMs.

About this task

You can scale up and scale down the CPU and memory of database server VMs through the CLI and API by updating the compute profile of the VM. You can scale database server VMs that host a single instance database or are part of a cluster. If the database server VM is part of a cluster, it is scaled independently of other VMs in the cluster, resulting in an asymmetric cluster after scaling.

Before you begin

You must have an existing compute profile with a different configuration of CPUs, cores, and RAM or create a new compute profile. For more information, see the [Creating a Compute Profile](#).

**Note:**

- The database server VM may power off while updating the VM configuration through Prism (cold-add).
- Generally, vCPUs and memory can be added on a running VM (hot-add). But if needed, NDB implicitly powers off the database server VM during scaling.
- The underlying hypervisor impacts the hot-add capabilities.

### Procedure

1. Enter the NDB CLI.

```
era >
```

2. Update the compute profile of the database server VM.

```
era > dbserver update id=dbserver_id compute_profile_id=computer_profile_id
```

Replace *dbserver\_id* with the ID of the database server VM.

Replace *computer\_profile\_id* with the compute profile ID of the database server VM.

When the operation is complete, the database server VM summary page shows the new compute profile.

# CLONE MANAGEMENT

---

NDB allows you to create and refresh clones to a point in time either by using transactional logs or by using snapshots. You can clone and refresh your clones to a point in time within minutes, thereby making the clone and refresh operations fast. NDB internally leverages the Nutanix time-efficient snapshots to create the clones. These database clones are highly space-efficient, consuming near-zero bytes by having only the further writes constituting to its size. NDB thereby lowers the cost of managing multiple copies of your databases significantly. You can update, refresh, or remove a database clone any time after its creation.

NDB allows you to perform log catch-up operation so that you can clone the source database at a point in time (that is, select a time). Also, the database is not quiesced for the time the log catch-up operation is performed. You can also manually perform a log catch-up operation to copy the current transaction logs of the source database to NDB.

See [Oracle Database Clone](#), [SQL Server Database Clone](#), [PostgreSQL Database Clone](#), [MySQL Database Clone](#), [MariaDB Database Clone](#), and [MongoDB Database Clone](#) for information on creating database clones.

## Updating Database Clone

About this task

Perform the following procedure to update a database clone. You can update the name, description and the associated tag values of the database clone.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones**, select the database clone that you want to update, and click **Update**.  
The **Update Clone** window appears.
3. Do the following in the indicated fields:
  - a. **Name**. Update the name of the database clone.
  - b. **Description**. Update the description of the database clone.
  - c. **Tags**. Displays the associated tags and tag values of the database clone. You can update the associated tag values. See [NDB Tags Management](#) on page 113 for more information.
  - d. Click **Update**.

NDB displays a message indicating that the operation to update the database clone is successfully completed.

## Refreshing Database Clones (Manual)

About this task

You can refresh your cloned databases to the latest state of your source databases. You refresh database clones either to a point in time (by using transaction logs) or by using snapshots. For refresh operations, pre-clone and post-clone scripts should be present on the target database.

Perform the following procedure to refresh a cloned database.

**Note:** does not support the refresh operation on database clones that are shared through an entity-sharing policy.

Procedure

1. In the drop-down list of the main menu, select **Databases**.

2. Go to **Clones**, select the cloned database you want to refresh, and click **Refresh**.  
The **Refresh Database** window appears.
3. Select the **Nutanix cluster** on which you want to refresh the database.

**Note:**

- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
- The **Nutanix Cluster** drop-down will only list the clusters associated with the selected time machine.

4. Select the time to which you want to refresh the database. Do the following to select the time.
  - a. On the **Month** tab, select a day of the month.
  - b. Click the **Day** tab to select the clone time to which you want to clone the database.
  - c. Do one of the following:
    - » **Point in time.** Select this option if you want to clone the source database to a point in time. If you select this option, the database is cloned by using the transactional logs.
    - » **Snapshot.** Select this option if you want to clone the source database by using the available snapshots.
  - d. Click **Next**.
5. Click **Refresh** to successfully refresh the database.

## Creating a Clone Refresh Schedule

### About this task

Refresh schedules update the clone with the most recent data available. You can define a frequency and time at which NDB automatically refreshes the database clone. Perform the following procedure to define a clone refresh schedule.

**Note:** Clone refresh schedules can also be defined when you create a clone. For more information, see [Oracle Database Clone](#), [SQL Server Database Clone](#), [PostgreSQL Database Clone](#), and [MariaDB Database Clone](#).

### Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones** and click the clone for which you want to define a refresh schedule.
3. From the **Schedule** drop-down list, select **Refresh > Create**. The **Create Refresh Schedule** window appears. Do the following in the indicated fields:
  - a. **Refresh Data Every (days).** Enter the number of days after which the clone should get refreshed.  
For example, if you have entered seven in this field, then NDB refreshes the clone after every seven days.
  - b. **Refresh Time.** Enter the time at which clone refresh operation should start.

4. Click **Create**.

A message appears indicating that the refresh operation has been created. You can click the message and then go to **Scheduled** on the **Operations** page to view all the scheduled tasks.

To change the frequency and time slot of the schedule, you can select **Refresh > Update** from the **Schedule** drop-down list, and then provide the new frequency and time slot.

## Deleting a Clone Refresh Schedule

About this task

You can cancel the schedule for clone refresh operation if you do not want to perform the clone refresh at the scheduled time. Perform the following procedure to remove a clone refresh schedule.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones** and click the clone for which you want to remove the refresh schedule.
3. From the **Schedule** drop-down list, select **Refresh > Remove**.

4. Click **Yes** to confirm.

This step removes the clone refresh schedule from NDB. However, you can always create a new schedule if needed. For more information, see [Creating a Clone Refresh Schedule](#) on page 135.

Alternatively, you can go to **Operations > Scheduled**, select the schedule, and then click **Delete** to remove the schedule.

## Removing a Database Clone (Manual)

Remove clone operation de-registers the database clone from NDB server and deletes the corresponding production database.

About this task

Perform the following procedure to remove a database clone from NDB.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones**, select the database clone that you want to remove, and click **Delete**. The **Remove Clone** window appears.

3. Do the following in the indicated fields:

- a. **Unregister the database clone from NDB.** By default, this option is selected.
- b. **Delete the database clone from VM.** Select this check box to remove the source database clone from the VM.
- c. Under **Confirm this request to remove the clone, <database clone name>, by providing the name of the clone**, type the name of the database clone that you want to remove.
- d. Click **Remove**.

NDB updates the status of the database clone and displays a message indicating that the operation to remove a database clone has started. Click the message to monitor the progress of the operation.

Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation.

**Note:** When clone delete operation is started, NDB checks if the data drives associated with the database clone being deleted are shared with other application or clones. If drives are not shared with other applications or clones on that host, NDB deletes the data drives.

## Creating a Clone Removal Schedule

About this task

You can define a clone removal schedule at which NDB automatically deletes the database clone on the scheduled date. Perform the following procedure to create a database clone removal schedule.

**Note:** Clone removal schedules can also be defined when you create a clone. For more information, see [MySQL Database Clone](#) and [Creating Snapshots \(Manual\)](#) on page 145.

Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones** and click the clone for which you want to define a removal schedule.
3. From the **Schedule** drop-down list, select **Removal > Create**. The **Create Removal Schedule** window appears. Do the following in the indicated fields:
  - a. **Date of Removal.** Select a date on which you want NDB to perform automatic clone removal operation.
  - b. **Time.** Select a time at which you want NDB to start clone removal operation.
  - c. **Database Server Time Zone.** Select a time zone for the database server.
  - d. Under **Items to remove**, select **Delete the database clone from the VM** if you want NDB to delete the respective clone from the VM.  
By default, NDB de-registers the database clone from the VM when this operation is triggered.
  - e. Under **Alerts**, select **Create an alert before removal** if you want to get notified before a database clone is removed and then enter the number of days before which you should get notified.  
For example, if you enter seven in this field, NDB notifies you seven days before a database clone is removed. A message appears indicating that the removal schedule has been created. You can click the message and then go to **Scheduled** on the **Operations** page to view all the scheduled tasks.

To modify the removal schedule, you can select **Removal > Update** from the **Schedule** drop-down list and then update the indicated fields.

## Deleting a Clone Removal Schedule

### About this task

You can cancel the schedule for clone removal operation if you do not want NDB to remove the clone at the defined schedule. You can then remove the clone manually (For more information, see [Removing a Database Clone \(Manual\)](#) on page 136).

Perform the following procedure to remove a clone refresh schedule.

### Procedure

1. In the drop-down list of the main menu, select **Databases**.
2. Go to **Clones** and click the clone for which you want to remove the removal schedule.
3. From the **Schedule** drop-down list, select **Removal** > **Remove**.
4. Click **Yes** to confirm.

This step removes the clone removal schedule from NDB. However, you can always create a new schedule if needed. For more information, see [Creating a Clone Removal Schedule](#) on page 137.

Alternatively, you can go to **Operations** > **Scheduled**, select the schedule, and then click **Delete** to remove the schedule.

# NDB TIME MACHINE MANAGEMENT

NDB time machines capture and maintain snapshots and transaction logs of your source databases as defined in the Data Access Management policies (DAM policies) and SLA schedules. A time machine is created for every source database that you register with NDB. You can choose the time machine availability across registered Nutanix clusters either during the process of database provisioning or by manually configuring the DAM policy in the **Time Machines** home page.

In the drop-down list of the main menu, select **Time Machines** and click the name of the time machine to open the home page of that time machine.

If a time machine is associated with a tag, **Time Machines** screen displays a search bar. You can click the search bar and select the tag to view the list of time machines associated with the selected tag.

## Time Machine Behaviour and Functionality

### Time Machine Behaviour

The following time machine behaviours are observed based on the status of the time machine:

| Time Machine Status | Time Machine Behaviour                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| READY               | When the time machine is in a READY state, all the snapshot and log catch-up operations are performed as configured or scheduled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| PAUSED              | <p>When the time machine is paused or in a PAUSED state, all the automatic snapshot and log catch-up (automatic and manual) activities are halted.</p> <p>Manual snapshots cannot be created when the time machine is in a PAUSED state.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| FROZEN              | <p>When a database is de-registered without removing the associated time machine, both the time machine and the database are set as FROZEN state.</p> <p>The following functionalities of a frozen time machine are locked:</p> <ul style="list-style-type: none"><li>• Automatic and manual snapshot creation.</li><li>• Automatic and manual log-catchup creation.</li><li>• Pause and Resume operations.</li></ul> <p>New clones can be created and existing clones can be refreshed using a frozen time machine.</p> <p>A frozen time machine can be deleted at any point in time. The clones created using a frozen time machine do not get deleted automatically when a frozen time machine gets deleted.</p> |

**Note:** NDB time machine cannot be paused, updated or deleted when the update time machine operation is in progress.

## Time Machine Functionality

### Snapshots

The time machine captures and replicates snapshots of the source database across multiple clusters (as defined in the DAM policy) at the time and frequency you specify in the schedule. Snapshots are used to clone the source database to a point in time when the snapshots were taken. If you want to take a snapshot of the current state of the source database, you can take a manual snapshot.

See [Snapshot Management](#) on page 144 for more information.

### Log Catch-up Operation

A log catch-up operation copies transaction logs from the source database based on the schedule you specify during the registration of the source database.

See [Database Log Management](#) on page 151 for more information.

### Schedule

You specify the schedule to take snapshots and copy transaction logs when you register the source database. The schedule includes details such as the number of snapshots that must be taken in a day, time the first snapshot must be taken, frequency of performing a log catch-up operation (copying transaction logs), and the time the daily, weekly, monthly, and quarterly snapshots must be taken. The time machine takes snapshots and performs the log catch-up operation at the time and frequency you specify in the schedule. You specify the schedule at the time of provisioning or registering a database.

### SLAs

SLAs are data retention policies that allow you to specify how long the daily, weekly, monthly, and quarterly snapshots are retained in NDB.

See [SLA Management](#) on page 153 for more information.

### Time Machine Heal Snapshot

The Time Machine heal snapshot is a database snapshot that NDB dispatches to remove gaps in the Time Machine's capability.

See [Time Machine Heal Snapshot](#) for more information.

### Time Machine Size

The size of the time machine refers to the storage space used by a specific time machine. It is calculated as the total space consumed by the snapshots and the associated transaction logs copied to NDB.

## Updating a Database Time Machine

### About this task

Perform the following procedure to update a database time machine.

### Procedure

1. In the drop-down list of the main menu, select **Time Machines** and click the time machine of the source database you want to update.

The Time Machines home page for the selected time machine appears.

2. Select the Nutanix cluster on which you want to perform the update.

**Note:**

- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
- Only the clusters associated with the selected time machine are listed.

3. Click **Actions** and select **Update**.

The **Update Time Machine** window appears.

4. Do the following in the indicated fields.

- a. **Name.** Update the name for the time machine.
- b. **Description.** Update the description for the time machine.
- c. **SLA.** Update the time machine SLA from the drop-down list. An SLA is a snapshot retention policy that indicates how long snapshots are retained in NDB. For more information, see [SLA Management](#) on page 153.

NDB has the following built-in SLAs.

| Name                                         | Databases | Continuous | Daily   | Weekly   | Monthly   | Quarterly   |
|----------------------------------------------|-----------|------------|---------|----------|-----------|-------------|
| <input type="radio"/> DEFAULT_OOB_GOLD_SLA   | 0         | 30 days    | 90 days | 16 weeks | 12 Months | 35 Quarters |
| <input type="radio"/> DEFAULT_OOB_SILVER_SLA | 0         | 14 days    | 60 days | 12 weeks | 12 Months | 0 Quarters  |
| <input type="radio"/> DEFAULT_OOB_BRONZE_SLA | 5         | 7 days     | 30 days | 8 weeks  | 6 Months  | 0 Quarters  |
| <input type="radio"/> DEFAULT_OOB_BRASS_SLA  | 1         | 0 days     | 7 days  | 0 weeks  | 0 Months  | 0 Quarters  |
| <input type="radio"/> NONE                   | 0         | 0 days     | 0 days  | 0 weeks  | 0 Months  | 0 Quarters  |

**Figure 64: SLAs**

**Note:** Modifying an SLA causes the time machine to set the retention policy of the new SLA. If you switch from a higher SLA to a lower SLA, snapshots and logs older than the new SLA retention levels are deleted. For example, if you change the SLA from DEFAULT\_OOB\_GOLD\_SLA to DEFAULT\_OOB\_BRONZE\_SLA, 30 days of daily snapshots are retained and older snapshots are deleted.

- d. Select the Nutanix clusters on which you want the logs and snapshots to be available.

**Note:** The number of Nutanix clusters displayed is based on the number of source clusters associated with the selected time machine.

5. Under **Schedule**, update the schedule to take and retain the snapshots.

6. Click **Update**.


NDB updates the selected time machine. An operation is initiated depending on the database engine and the SLA update (non-PITR SLA to a PITR SLA and vice versa).

A message appears indicating that the operation to update the time machine has started. Click the message to monitor the progress of the operation. Alternatively, select **Operations** from the drop-down list of the main menu to monitor the progress of the operation.

## Data Access Management

NDB multi-cluster allows you to manage time machine data availability across all the registered Nutanix clusters in NDB. Using the Data Access Management policies (DAM policies), you can add, update, and remove the time machine data accessibility across the multiple registered clusters in NDB.

Click the **Diagram** tab in the **Data Access Management** page to display a diagram of the registered Nutanix cluster with the source databases, the associated time machines, and the time machine data availability across clusters. You can click each node displayed under the **Diagram** tab to view the respective details.



**Figure 65: Data Access Management: Diagram View**

## Adding Time Machine Data Access to a Nutanix Cluster

Adding time machine data access to a Nutanix cluster associates the time machine of the source database with that cluster. This allows you perform all snapshot operations (snapshot creation, replication, deletion and so on) on the associated Nutanix clusters.

Before you begin

Ensure that you have enabled multi-cluster in NDB and have registered the required Nutanix clusters.

### Caution:

- If the database server VMs or the source databases contain disks from storage containers other than the NDB server storage container or the storage container used by NDB for provisioning, then you must configure the storage mappings between the source and target clusters for such containers before performing any replication operation.


About this task

Perform the following procedure to add time machine data access to a Nutanix cluster.

### Procedure

1. In the drop-down list of the main menu, select **Time Machines**.  
This page displays a list of time machines available in NDB.
2. Click the time machine of the source database for which you want to add data access.  
The overview page of the time machine is displayed.

3. Go to **Data Access Management**, select the **Table** view and click **Add**.  
The **Add Data Access to Nutanix Cluster** window appears.
4. Do the following in the indicated fields.
  - a. **Nutanix Cluster**. Select a Nutanix cluster.
  - b. **DAM Policy**. Select an SLA from the drop-down list.
  - c. Select the Nutanix clusters on which you want to replicate the snapshot.



**Add Data Access to Nutanix Cluster** ? | X

Nutanix Cluster

Select a Cluster

DAM Policy

DEFAULT\_OOB\_BRASS\_SLA

API Equivalent Cancel Add

**Figure 66: Adding Data Access**

5. Click **Add** to add data access to the Nutanix cluster.

## Updating Time Machine Data Access to a Nutanix Cluster

Before you begin

Ensure that you have enabled multi-cluster in NDB and have registered the required Nutanix clusters.

About this task

Perform the following procedure to update time machine data access to a Nutanix cluster.


Procedure

1. In the drop-down list of the main menu, select **Time Machines**.  
This page displays a list of time machines available in NDB.
2. Click the time machine of the source database for which you want to add data access.  
The overview page of the time machine appears.
3. Go to **Data Access Management**, select the **Table** view, and select the Nutanix cluster that you want to update.

4. Click **Update**. The **Update Nutanix Cluster's Data Access** window appears.

**Note:** Follow the [Updating a Database Time Machine](#) on page 140 procedure for the Nutanix clusters that are not allowed to be updated in the Data Access Management page.

5. Do the following in the indicated fields.
  - a. **Nutanix Cluster**. Displays the name of the selected cluster.
  - b. **DAM Policy**. Update the SLA from the drop-down list.



The screenshot shows a dialog box titled "Update Nutanix Cluster's Data Access". It features two dropdown menus. The first, labeled "Nutanix Cluster", has "ta03" selected. The second, labeled "DAM Policy", has "DEFAULT\_OOB\_BRASS\_SLA" selected. Below the dropdowns are three buttons: "API Equivalent", "Cancel", and "Update".

**Figure 67: Updating Data Access**

6. Click **Update**.

NDB updates the cluster data access according to the **DAM Policy** selected. An operation initiates depending on the database engine and the SLA update (non-PITR SLA to a PITR SLA and vice versa).

A message appears indicating that the operation to update the time machine has started. Click the message to monitor the progress of the operation. Alternatively, select **Operations** from the drop-down list of the main menu to monitor the progress of the operation.

## Snapshot Management

NDB time machine allows you to capture and replicate snapshots of the source database across multiple clusters (as defined in the DAM policy) at the time and frequency specified in the schedule. Snapshots are used to clone the source database to a point in time when the snapshots were taken. If you want to take a snapshot of the current state of the source database, you can take a manual snapshot.

Snapshots store the state of a database at various points in time (as specified in the schedule). Therefore, if you want to clone a database by using snapshots, your choice of state for the database clone is limited to the points in time at which the snapshots are taken. The database is quiesced for the time the snapshot is being taken. Therefore, you would want to take only a minimum number of snapshots in a day.

**Note:** You must not restore database server VMs using OS snapshots taken on Prism or otherwise, as this causes metadata inconsistency on NDB.

## Creating Snapshots (Manual)

About this task

If you want to take a snapshot of the current state of the source database, you can take a manual snapshot.

Perform the following procedure to take a manual snapshot.

Procedure

1. You can create a manual snapshot by doing any one of the following.

- » In the drop-down list of the main menu, go to **Databases > Sources** and select the database for which you want to create a manual snapshot.
- » In the drop-down list of the main menu, select **Time Machines** and click the time machine of the source database for which you want to create a manual snapshot.

The overview page of the time machine is displayed.

2. Click **Snapshot** or Go to **Snapshots** and click **Create**.

The **Create Snapshot** window appears.

3. Do the following in the indicated fields.

- a. **Snapshot Name.** Type a name for the snapshot.
- b. **Removal Schedule.** Select this check box if you want to define a snapshot removal schedule and then enter the number of days after which the snapshot should be removed.
- c. Select the Nutanix cluster on which you want to replicate the snapshot.

**Note:**

- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
- Only the clusters associated with the selected time machine is listed.

Snapshot Name

☐ Removal Schedule

Nutanix clusters where this snapshot should be made available.

| <input type="checkbox"/> Name           | Description |
|-----------------------------------------|-------------|
| <input checked="" type="checkbox"/> tm4 | tm4         |
| <input type="checkbox"/> tm4            | tm4         |

API Equivalent Cancel Create

**Figure 68: Creating Manual Snapshot**

4. Click **Create** to create a snapshot.

A message appears indicating that the operation to create snapshot has started. Click the message to monitor the progress of the operation.

Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation. On the same page, click **Stop** to stop the snapshot operation.

## Replicating Snapshots (Manual)

### About this task

You can replicate existing snapshots of a time machine to the other Nutanix clusters associated with that time machine.

Perform the following procedure to replicate a snapshot.

## Procedure

1. In the drop-down list of the main menu, select **Time Machines**.  
This page displays a list of time machines available in NDB.
2. Click the time machine of the source database for which you want to replicate a snapshot.  
The overview page of the time machine is displayed.
3. Go to **Snapshots** and select a Nutanix cluster.  
This page displays a list of snapshots available in the selected Nutanix cluster.
4. Select an existing snapshot and click **Replicate**. The **Replicate Snapshot** window appears.
5. Do the following in the indicated fields.
  - a. Select the Nutanix clusters on which you want to replicate the snapshot.

### Note:

- NDB multi-cluster must be enabled to select different Nutanix clusters. See [Nutanix Cluster Management](#) on page 91 for more information.
- Only the clusters associated with the selected time machine is listed. See [Data Access Management](#) on page 141 for more information.

- b. **Removal Schedule.** Select this check box if you want to define a snapshot removal schedule and then enter the number of days after which the snapshot should be removed.

Replicate Snapshot

?

×

Nutanix clusters where this snapshot can be replicated.

| Name ↕                               | Description ↕ |
|--------------------------------------|---------------|
| <input checked="" type="radio"/> tm4 |               |

Remove in (days)

☒ Removal Schedule

–

7

+

Nutanix clusters where this snapshot already exists.

| Name ↕     | Description ↕ |
|------------|---------------|
| EraCluster |               |

API Equivalent

Cancel

Replicate

Figure 69: Replicating Snapshot

6. Click **Replicate** to successfully replicate a snapshot.

A message appears indicating that the operation to replicate snapshot has started. Click the message to monitor the progress of the operation.

Alternatively, select **Operations** in the drop-down list of the main menu to monitor the progress of the operation. On the same page, click **Stop** to stop the snapshot operation.

## Sanitized Snapshots

Data sanitization refers to the secure and permanent erasure of sensitive data from a database. With NDB, you can create sanitized snapshots of databases to create sanitized database clones. You can achieve this in two ways:

- Synchronous approach: Creating a sanitized snapshot using create or refresh clone workflows where you perform sanitization using post scripts.
- Asynchronous approach: Creating a sanitized snapshot on an existing clone where you perform sanitization using a third-party tool or a custom script.

### Note:

- NDB does not support sanitized snapshots for the Oracle database engine.
- A clone created from a sanitized snapshot is called a sanitized clone. You can refresh a sanitized clone only with a sanitized snapshot.
- You cannot create a sanitized snapshot from a sanitized clone.

### Creating Sanitized Database Snapshots Synchronously

This topic describes the steps to create a sanitized database snapshot in a synchronous manner.

Before you begin

You must have a script that sanitizes the database.

### Procedure

1. To create a sanitized snapshot by using the synchronous approach, do the following:
  - a. Create the input file for clone creation with sanitized snapshot:

```
clone generate_input_file time_machine_id=<time_machine_id>
engine=<database_engine> is_clustered_clone=false create_dbserver=true
vm_name_prefix=<prefix> compute_profile=<compute_profile>
network_profile=<network_profile> nx_cluster_id=<cluster_id>
db_parameter=<db_parameter_profile> snapshot_id=<id> capture_sanitised_snapshot
sanitised_snapshot_name=<name> delete_sanitised_snapshot_in_days=<no_of_days>
output_file=<output_file.json>
```

The following are the parameters of the CLI command:

- `clone`. Specify this parameter for clone operations.
- `generate_input_file`. Specify this parameter to generate input file.
- `time_machine_id`. Specify the time machine ID.
- `engine`. Specify the database engine.
- `is_clustered_clone`. Specify if the clone is clustered.

- `create_dbserver`. Specify if you want to create a database server. This value can be set to `true` or `false` as required.
  - `vm_name_prefix`. Specify the VM name prefix.
  - `compute_profile`. Specify the compute profile.
  - `network_profile`. Specify the network profile.
  - `nx_cluster_id`. Specify the Nutanix cluster ID.
  - `db_parameter`. Specify the database parameter profile.
  - `snapshot_id`. Specify the snapshot ID.
  - `capture_sanitised_snapshot`. Specify this parameter to capture sanitised snapshot.
  - `sanitised_snapshot_name`. Specify a sanitised snapshot name.
  - `delete_sanitised_snapshot_in_days`. Specify the number of days after which NDB deletes the sanitised snapshot.
  - `output_file`. Specify the name of the output file.
- b. Specify the complete path to the sanitization script as the value of the `post_clone_cmd` parameter in the JSON file. The script must be present on the clone VM.
- c. Create the clone using the command:

```
era > clone create info_file=<output_file.json> time_machine_id=<time_machine_id>
snapshot_id=<snapshot_id> name=<clone_name>
```

- `clone`. Specify this parameter for clone operations.
- `create`. Specify this parameter to create a clone.
- `info_file`. Specify the input file.
- `time_machine_id`. Specify the time machine ID.
- `snapshot_id`. Specify the snapshot ID.
- `name`. Specify the name of the clone.

The create clone operation starts. You can monitor the status of the operation from the **Operations** menu. NDB runs the specified script that sanitizes the database in the 'Execute Post-Scripts' sub-operation. The **Time Machines** page displays the sanitized snapshots.

2. (Optional) Alternatively, you can create a sanitized snapshot during a clone refresh operation as follows:

a. Create the input file for clone refresh with sanitized snapshot:

```
clone refresh generate_input_file name=<clone_name> snapshot_id=<id>
capture_sanitised_snapshot sanitised_snapshot_name=<name>
delete_sanitised_snapshot_in_days=<no_of_days> output_file=<output_file>
```

The following are the parameters of the CLI command:

- `clone`. Specify this parameter for clone operations.
- `refresh`. Specify this parameter to refresh a clone.
- `generate_input_file`. Specify this parameter to generate input file.
- `name`. Specify the name of the clone.
- `snapshot_id`. Specify the snapshot ID.
- `capture_sanitised_snapshot`. Specify this parameter to capture sanitised snapshot.
- `sanitised_snapshot_name`. Specify a sanitised snapshot name.
- `delete_sanitised_snapshot_in_days`. Specify the number of days after which NDB deletes the sanitised snapshot.
- `output_file`. Specify the name of the output file.

b. Perform clone refresh operation by using the command:

```
era > clone refresh info_file=<output_file.json>
```

- `clone`. Specify this parameter for clone operations.
- `refresh`. Specify this parameter to refresh a clone.
- `info_file`. Specify the input file.

The clone refresh operation starts. You can monitor the status of the operation from the **Operations** menu. The **Time Machines** page displays the sanitized snapshots.

### Creating Sanitized Database Snapshots Asynchronously

This topic describes the steps to create a sanitized database snapshot in an asynchronous manner.

Before you begin

You must have a sanitized database clone to create the snapshot from.

#### Procedure

To create a sanitized snapshot by using the asynchronous approach, use the following command:

```
clone snapshot id=<clone_id> snapshot_name=<name> sanitised=true
delete_sanitised_snapshot_in_days=<no_of_days>
```

The following are the parameters of the CLI command:

- `clone`. Specify this parameter for clone operations.
- `snapshot`. Specify this parameter to create a snapshot for a clone.
- `id`. Specify the ID of the clone to take the snapshot.
- `snapshot_name`. Specify the snapshot name.

- `sanitised`. specify whether the clone is sanitized.
- `delete_sanitised_snapshot_in_days`. Specify the number of days after which NDB deletes the sanitised snapshot.

This command creates a sanitized snapshot from the existing database clone.

## Database Log Management

A log catch-up operation copies database logs (transaction logs, redo logs, wal logs) from the source database to NDB. The logs get copied based on the schedule specified during the registration of the source database. The frequency of the log catches up operation is higher than the number of snapshots taken in a day. By default, NDB performs a log catch-up operation every 30 minutes. Therefore, the log catch-up operation enables you to clone the source database at a point in time (that is, select a time). Also, the database is not quiesced for the time the log catch-up operation is performed. You can also manually perform a log catch-up operation to copy the current transactional logs of the source database to NDB.

With NDB log replication, when a log catch-up operation is performed on the source database or the NDB active node, the logs are replicated to all the associated clusters with PITR-based SLAs.

Example: Consider a time machine that is associated with clusters C1, C2, and C3. All the clusters are configured with PITR-based SLAs and cluster C1 is the NDB active node. With NDB log replication, when a log catch-up operation is performed on cluster C1, the logs are also replicated to clusters C2 and C3.

**Note:** To avoid log catch-up issues while using Volume Group, enable data services IP address (DSIP) for VG connections.

See [Updating Time Machine Data Access to a Nutanix Cluster](#) on page 143 for information on updating the time machine DAM policy.

## Database Log Backup Management in NDB

NDB allows you to create and refresh zero-byte database clones at any point in time. NDB captures the database transaction logs using the Log Catch-up operation, which can be configured to run every 15 minutes to 2 hours.

### Oracle

For Oracle databases, NDB switches the current online redo logs and then copies all the archived logs generated since the last log catch-up to the permanent log storage location (also known as log drive) of NDB. Note that NDB does not remove any logs from their original location after copying. The management and retention of those logs are according to the policies set by the database administrator.

### Microsoft SQL Server

NDB uses the SQL Server transaction log backups for Microsoft SQL Server. Whenever NDB takes a log backup of the managed SQL Server database, the database logs get truncated. NDB requires exclusive access to the log backups of the managed database.

### PostgreSQL

For PostgreSQL instances, NDB uses the archive command to perform log catch-up. PostgreSQL maintains a write-ahead log (WAL) in the `pg_wal/` subdirectory of the cluster data directory. During a log catch-up operation, NDB switches the current WAL segment and then copies all the archived WAL segments generated since the last log catch-up to the permanent log storage location (also known as log drive) of NDB.

For Highly Available PostgreSQL Instance, NDB configures an additional archive command to archive WAL segments to an archive location. These WAL segments can be used for replication or recovery during failover. As part of the snapshot operation NDB purges older WAL segments in this archive location based on an instance property defined in NDB called `archive_wal_expire_days`. By default the value of this property is -1 which indicates that automatic purging of older WAL segments is disabled. All WAL segments older than `archive_wal_expire_days` is purged in the next snapshot operation.

**Note:** You can modify the `archive_wal_expire_days` property using NDB API or CLI.

#### MariaDB and MySQL

For MariaDB and MySQL instances, NDB copies the binary log files generated since the last Log Catch-up from the configured `bin_log` location to the permanent log storage location (also known as log drive) of NDB. Note that binary logging must be enabled for the instance to be registered in NDB.

#### MongoDB

NDB uses the transaction log backups for MongoDB. Whenever NDB takes a log backup of the managed MongoDB database, the database logs get truncated. NDB requires exclusive access to the log backups of the managed database.

**Note:** NDB automatically expands the log drive size by a default value of 20 percent when the space usage reaches a default threshold of 80 percent.

**Table 20: Working with Third-Party Backup Solutions**

| Database Engine      | Recommendation for working with third-party backup solutions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle               | Third-party backups vendors can continue to backup the archived logs from the database log destination. If the backup product is configured to delete the logs after a backup, NDB will still be able to capture them, considering the log catch-up begins before the logs get backed up.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Microsoft SQL Server | If the database is managed by another third-party backup vendor, you can choose to grant exclusive log access to NDB while registering the database to NDB. If you do not grant exclusive log access to NDB, the Point-in-Time (PIT) recovery will not be supported, but snapshot-based recovery of the databases is supported.<br><br><b>Note:</b> Nutanix does not recommend third-party log backup solutions for databases managed by NDB. This may result in external transaction log truncation which could cause log sequence number (LSN) mismatch issues and trigger multiple heal snapshots for the same time machine for SQL Server. This may also result in gaps in the point in time from which you can perform a restore operation. |
| PostgreSQL           | Since PostgreSQL supports complex archive commands, third-party backups vendors can continue to use the archiving feature of PostgreSQL to perform the backup of WAL segments.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| MariaDB and MySQL    | Third-party backups vendors can continue to backup the binary logs from the binary log location. If the backup product is configured to delete the logs after the backup, NDB will still be able to capture them, considering the log catch-up begins before the logs get purged.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## Performing a Log Catch-up Operation (Manual)

### About this task

You can manually perform a log catch-up operation to copy the current transactional logs of the source database to NDB.

Perform the following procedure to perform a manual log catch-up operation.

### Procedure

1. In the drop-down list of the main menu, select **Time Machines**.
2. Select the time machine of the source database, click **Actions** and select **Log Catch Up**.
3. In the confirmation message that appears, click **Yes**.

## SLA Management

SLAs are data retention policies that allow you to specify how long the daily, weekly, monthly, and quarterly snapshots are retained in NDB.

SLAs include the following details.

### Daily snapshots

Duration in days for which a daily snapshot must be retained in NDB.

When you register a source database with NDB, you specify a time at which the daily snapshot must be taken. The snapshot taken at this time becomes your daily snapshot. The daily snapshot is the only snapshot that is retained for that day.

For example, if you specify that the daily snapshot must be taken at 11:00 AM and in the SLA you specify that daily snapshots must be retained for 30 days, the snapshot that is taken at 11:00 AM everyday is retained as daily snapshot and this snapshot is retained for 30 days.

### Weekly snapshots

Duration in weeks for which a weekly snapshot must be retained in NDB.

When you register a source database with NDB, you specify a day of the week on which the weekly snapshot must be taken. The snapshot taken on this day of the week becomes your weekly snapshot. The weekly snapshot is the only snapshot that is retained for that week along with the daily snapshot of each day of the week.

**Note:** The daily snapshots are deleted if their duration expires.

For example, if you specify that the weekly snapshot must be taken on Monday and in the SLA you specify that weekly snapshots must be retained for 8 weeks, the snapshot that is taken on Monday of every week is retained as a weekly snapshot and this snapshot is retained for 8 weeks.

### Monthly snapshots

Duration in months for which a monthly snapshot must be retained in NDB.

When you register a source database with NDB, you specify a day of the month on which the monthly snapshot must be taken. The snapshot taken on this day of the month becomes your monthly snapshot. The monthly snapshot is the only snapshot that is retained for that month along with the weekly and daily snapshots of that month.

**Note:** The weekly and daily snapshots are deleted if their duration expires.

For example, if you specify that the monthly snapshot must be taken on the 15th of every month and in the SLA you specify that monthly snapshots must be retained for 6 months, the snapshot that is taken on the 15th of every month is retained as a monthly snapshot and this snapshot is retained for 6 months.

### Quarterly snapshots

Duration in number of quarters for which a quarterly snapshot must be retained in NDB.

The snapshot taken on the first day of the first month of the quarter is retained as a quarterly snapshot in NDB. When you register a source database with NDB, in the schedule if you select **Jan, Apr, Jul, Oct** from the drop-down list, snapshots taken on January 1, April 1, July 1, and October 1 are retained as quarterly snapshots. The quarterly snapshot is the only snapshot that is retained for that quarter along with the monthly, weekly, and daily snapshots of that month.

**Note:** The monthly, weekly, and daily snapshots are deleted if their duration expires.

For example, if you specify in the SLA that the quarterly snapshots must be retained for 10 quarters, this quarterly snapshot is retained for 10 quarters.

### Continuous

Duration in days for which transaction logs are retained in NDB.

For example, if you set this value to 10 days, transaction logs of the last 10 days are retained from the time the first daily snapshot is taken.

## Built-in SLAs

NDB has five built-in SLAs. In addition, you can create custom SLAs according to your requirements. You can also modify the configured SLAs for the time machines.

**Table 21: Built-in SLA values**

| Built-in SLA           | SLA Type     | Continuous Retention (in days) | Daily Retention (in days) | Weekly Retention (in weeks) | Monthly Retention (in months) | Quarterly Retention (in quarters) |
|------------------------|--------------|--------------------------------|---------------------------|-----------------------------|-------------------------------|-----------------------------------|
| DEFAULT_OOB_GOLD_SLA   | PITR SLA     | 30                             | 90                        | 16                          | 12                            | 35                                |
| DEFAULT_OOB_SILVER_SLA | PITR SLA     | 14                             | 60                        | 12                          | 12                            | 0                                 |
| DEFAULT_OOB_BRONZE_SLA | PITR SLA     | 7                              | 30                        | 8                           | 6                             | 0                                 |
| DEFAULT_OOB_BRASS_SLA  | Non-PITR SLA | 0                              | 7                         | 0                           | 0                             | 0                                 |
| NONE                   | Non-PITR SLA | 0                              | 0                         | 0                           | 0                             | 0                                 |

## Creating Custom SLAs

### About this task

If you do not want to use or update the built-in SLAs, you can create a new SLA with the values that suit your requirement.

Perform the following to create a new SLA.

### Procedure

1. In the drop-down list of the main menu, select **SLAs**.
2. Click **Create**.  
The **Create SLA** window appears.
3. Do the following in the indicated fields.
  - a. **Name**. Type a name for the SLA.
  - b. **Description**. Type a description for the SLA.
  - c. **Continuous log retention (days)**. Select the duration in days for which the transaction logs are retained in NDB.
  - d. **Daily snapshot retention (days)**. Select the duration in days for which a daily snapshot must be retained in NDB.
  - e. **Weekly snapshot retention (weeks)**. Select the duration in weeks for which a weekly snapshot must be retained in NDB.
  - f. **Monthly snapshot retention (months)**. Select the duration in months for which a monthly snapshot must be retained in NDB.
  - g. **Quarterly snapshot retention (quarters)**. Select the duration in number of quarters for which a quarterly snapshot must be retained in NDB.

- Click **Create** to create a new SLA.

? | ✕

**Name**

**Description**

**Frequency and retention**

Select how long each type of log/snapshot will be stored

Continuous log retention (days)

–
30
+

Daily snapshot retention (days)

–
7
+

Weekly snapshot retention (weeks)

–
2
+

Monthly snapshot retention (months)

–
2
+

Quarterly snapshot retention (quarters)

–
1
+

Cancel

Create

**Figure 70: Create Custom SLA**

## Comparing SLAs

SLAs are compared based on the frequency of snapshots and database logs captured. Nutanix does not recommend comparing SLAs based on their retention period.

Example: The following tables list the retention periods and the frequency of data captured by custom SLAs.

**Table 22: Custom SLAs: Retention Period**

| SLA           | Retention Period               |                           |                             |                               |                                   |
|---------------|--------------------------------|---------------------------|-----------------------------|-------------------------------|-----------------------------------|
|               | Continuous Retention (in days) | Daily Retention (in days) | Weekly Retention (in weeks) | Monthly Retention (in months) | Quarterly Retention (in quarters) |
| PITR_SLA_1    | 40                             | 30                        | 10                          | 6                             | 2                                 |
| PITR_SLA_2    | 1                              | 0                         | 0                           | 0                             | 0                                 |
| DAILY_SLA_1   | 0                              | 50                        | 20                          | 10                            | 1                                 |
| DAILY_SLA_2   | 0                              | 1                         | 2                           | 1                             | 1                                 |
| WEEKLY_SLA_1  | 0                              | 0                         | 4                           | 2                             | 1                                 |
| WEEKLY_SLA_2  | 0                              | 0                         | 8                           | 0                             | 0                                 |
| MONTHLY_SLA_0 | 0                              | 0                         | 0                           | 2                             | 1                                 |

| SLA             | Retention Period               |                           |                             |                               |                                   |
|-----------------|--------------------------------|---------------------------|-----------------------------|-------------------------------|-----------------------------------|
|                 | Continuous Retention (in days) | Daily Retention (in days) | Weekly Retention (in weeks) | Monthly Retention (in months) | Quarterly Retention (in quarters) |
| MONTHLY_SLA_0   |                                | 0                         | 0                           | 4                             | 2                                 |
| QUARTERLY_SLA01 |                                | 0                         | 0                           | 0                             | 2                                 |
| QUARTERLY_SLA02 |                                | 0                         | 0                           | 0                             | 4                                 |
| NONE_SLA        | 0                              | 0                         | 0                           | 0                             | 0                                 |

**Table 23: Custom SLAs: Snapshots and Database Logs Captured**

| SLA                | Frequency of Data Captured |      |           |           |            |            |            |            |
|--------------------|----------------------------|------|-----------|-----------|------------|------------|------------|------------|
|                    | Daily                      |      | Weekly    |           | Monthly    |            | Quarterly  |            |
|                    | Snapshots                  | Logs | Snapshots | Logs      | Snapshots  | Logs       | Snapshots  | Logs       |
| PITR_SLA_1 or more | 1 or more                  |      | 7 or more | 7 or more | 30 or more | 30 or more | 90 or more | 90 or more |
| PITR_SLA_2 or more | 1 or more                  |      | 7 or more | 7 or more | 30 or more | 30 or more | 90 or more | 90 or more |
| DAILY_SLA1 or more | 0                          |      | 7 or more | 0         | 30 or more | 0          | 90 or more | 0          |
| DAILY_SLA2 or more | 0                          |      | 7 or more | 0         | 30 or more | 0          | 90 or more | 0          |
| WEEKLY_SLA_1       | 0                          |      | 1 or more | 0         | 4 or more  | 0          | 12 or more | 0          |
| WEEKLY_SLA_2       | 0                          |      | 1 or more | 0         | 4 or more  | 0          | 12 or more | 0          |
| MONTHLY_SLA_1      | 0                          |      | 0         | 0         | 1 or more  | 0          | 3 or more  | 0          |
| MONTHLY_SLA_2      | 0                          |      | 0         | 0         | 1 or more  | 0          | 3 or more  | 0          |
| QUARTERLY_SLA_1    | 0                          |      | 0         | 0         | 0          | 0          | 1 or more  | 0          |
| QUARTERLY_SLA_2    | 0                          |      | 0         | 0         | 0          | 0          | 1 or more  | 0          |
| NONE_SLA0          | 0                          |      | 0         | 0         | 0          | 0          | 0          | 0          |

Following are the conclusions when the SLAs are compared based on the frequency of data captured.

|                                             |
|---------------------------------------------|
| PITR_SLA_1 is equal to PITR_SLA_2           |
| DAILY_SLA_1 is equal to DAILY_SLA_2         |
| WEEKLY_SLA_1 is equal to WEEKLY_SLA_2       |
| MONTHLY_SLA_1 is equal to MONTHLY_SLA_2     |
| QUARTERLY_SLA_1 is equal to QUARTERLY_SLA_2 |
| PITR_SLA is higher than DAILY_SLA           |

|                                                 |
|-------------------------------------------------|
| DAILY_SLA is higher than WEEKLY_SLA             |
| WEEKLY_SLA is higher than MONTHLY_SLA           |
| MONTHLY_SLA is higher than QUATERLY_SLA         |
| NONE_SLA does not capture any snapshots or logs |

## Authorizing Database Server VMs

### About this task

If you want to clone the source database to a database server VM that you have already registered with NDB, you must associate the time machine of the source database with the target database server VM.

Perform the following procedure to authorize database server VMs.

### Procedure

1. In the drop-down list of the main menu, click **Time Machines**.
2. Select the time machine of the source database, click **Actions** and select **Authorize Database Server VMs**. The **Authorize Database Server VMs** window appears
3. Select the target database server VM, move the database server VM to the **Authorized Database Server VMs** list, and click **Update**.

## Deauthorizing Database Server VMs

### About this task

Perform the following procedure to deauthorize database server VMs.

### Procedure

1. In the drop-down list of the main menu, click **Time Machines**.
2. Select the time machine of the source database, click **Actions**, and select **Authorize Database Server VMs**. The **Authorize Database Server VMs** window appears.
3. Select the database server VM from the **Authorized Database Server VMs** list and move it to the **Database Server VMs** list.
4. Click **Update**.

## Time Machine Heal Snapshot

A time machine captures and maintains snapshots and transaction logs of your source databases as defined in the Data Access Management policies (DAM policies) and SLA schedules. NDB uses snapshots and transaction logs to clone the source database to a point in time where the continuous region is available.

However, there are instances when the transaction log catchup operation succeeds with the gaps in the logs and the continuity of the time machine recovery window is broken. The logs collected after the broken time event cannot be consumed for cloning. Therefore, time machine waits for the next database snapshot to occur so that the subsequent logs can be used along with the latest snapshot to clone the source database. In these scenarios, a time machine heal snapshot is required to minimize the gap as much as possible.

The time machine heal snapshot is a database snapshot that NDB dispatches to remove gaps in the time machine's capability. There are two ways to heal a broken time machine recovery window:

- **Manual Heal.** Manually monitor time machine capability for gaps and request for time machine heal snapshot using the below CLI:

```
era > time_machine heal id=<tm_id>|name=<tm_name>
```

You can use either *tm\_id* or *tm\_name* variable in the earlier command.

Replace *tm\_id* with time machine ID number.

Replace *tm\_name* with time machine name.

- **Automatic Heal.** By default, the time machine auto-heal feature is enabled. You can disable auto-heal by setting the `enable_time_machine_auto_correct_engine` flag to false in the **ndb-config** file. Time Machine auto-heal is an engine that runs on the NDB server every 5 minutes and checks for different cases that require a heal. If needed, the NDB server triggers heal snapshots.

# ALERT MANAGEMENT

NDB generates an alert if any operation fails. Every operation such as provisioning a database or creating a clone has an alert policy associated with it. If an operation fails and the policy associated with the operation is enabled, NDB generates an alert. All alerts are displayed in the alerts dashboard.

You can configure NDB to send alert notifications to an email recipient.

# Alerts Dashboard

The alerts dashboard displays all the alerts with details such as the entity, severity level, status, owner, last modified time, the number of occurrences, and the last occurred time. The dashboard groups duplicate alerts and displays the number of occurrences for the same alert.

[illegible]

### Figure 71: Alerts Dashboard

You can do the following from the alerts dashboard.

### Set Status to Resolved

Resolve the alerts if the issue is resolved.

### Set Status to Acknowledged

Acknowledge the alerts if you are aware of the issue.

## Alert Policies Management

NDB has an alert policy defined for every operation that you perform. NDB generates an alert for an operation only if the alert policy associated with that operation is enabled. All alert policies are enabled by default and the severity level of the alert policies are set to **Critical**. You can disable an alert policy if you do not want NDB to generate an alert for an operation even if that operation fails. You can also change the severity level of the alert policies.

## Disabling or Enabling Alert Policies

### About this task

All alert policies are enabled by default. You can disable an alert policy if you do not want NDB to generate an alert for an operation even if that operation fails. Enable an alert policy if you have previously disabled that policy.

Perform the following procedure to disable or enable alert policies.

### Procedure

1. From the drop-down list of the main menu, select **Alerts**.
2. Go to **Policies**.
3. Select the policy and click **Set Disabled** or **Set Enabled**.

## Setting Alert Retention

### About this task

After you have resolved an alert, you can set the duration (in days) for which NDB should retain that alert.

### Procedure

1. From the drop-down list of the main menu, select **Alerts**.
2. Go to **Policies**.
3. Click **Update** next to **Alert Retention**.

The screenshot shows the NDB Alerts page. The 'Alerts' menu is selected in the top navigation bar. On the left sidebar, the 'Policies' link is highlighted. The main content area displays a table of alert policies. At the bottom of the page, the 'Alert Retention' section is visible, with an 'Update' button highlighted by a green box. Below the button, it states 'Resolved Alerts will be deleted after 180 day(s)'.


| Policy Name                                             | Severity Assigned to Alert | Enabled | Operation               | Last Modified       |
|---------------------------------------------------------|----------------------------|---------|-------------------------|---------------------|
| <input type="radio"/> Scheduled tasks are not executing | Warning                    | ✓       |                         | 2022-07-05 21:01:45 |
| <input type="radio"/> vCenter Snapshots Present         | Critical                   | ✓       |                         | 2022-01-25 13:57:15 |
| <input type="radio"/> SQL Database Version Changed      | Critical                   | ✓       |                         | 2022-01-25 13:57:15 |
| <input type="radio"/> Dissociating Software Profile     | Critical                   | ✓       |                         | 2021-12-27 07:14:08 |
| <input type="radio"/> Database Version Changed          | Warning                    | ✓       |                         | 2021-11-09 19:07:44 |
| <input type="radio"/> Database is down                  | Critical                   | ✓       |                         | 2021-09-18 20:37:40 |
| <input type="radio"/> DBServer is unhealthy             | Critical                   | ✓       |                         | 2021-09-18 20:37:40 |
| <input type="radio"/> Storage Space Critical            | Critical                   | ✓       |                         | 2021-09-18 20:37:40 |
| <input type="radio"/> Add Database                      | Critical                   | ✓       | Add Database            | 2021-09-18 20:37:39 |
| <input type="radio"/> DB Server TimeZone Changed        | Critical                   | ✓       |                         | 2021-09-18 20:37:39 |
| <input type="radio"/> Extend Database Storage           | Warning                    | ✓       | Extend Database Storage | 2021-09-18 20:37:39 |

Figure 72: Update Alert Retention

The **Alert Retention** window appears.

4. Increase or decrease the number of days you want NDB to retain alerts after you have resolved the alerts, and click **Update**.

**Note:** The alert retention setting is a global setting and not a per alert policy setting.

The image shows a dialog box titled "Alert Retention" with a question mark icon and a close button (X). Below the title bar, it says "After alert status is set to 'Resolved', delete after :". There is a numeric input field containing "180" with minus and plus buttons on either side, followed by the text "days.". At the bottom right, there are two buttons: "Cancel" and "Update".

Alert Retention ? X

After alert status is set to 'Resolved', delete after :

- 180 + days.

Cancel Update

**Figure 73: Alert Retention Window**

## Resolving Alerts

NDB allows you to resolve multiple alerts based on the alert policy or the entity for which the alert occurred. You can resolve all alerts related to an alert policy, a database, a database server VM, or a time machine with a single click.

### Resolving All Alerts for an Alert Policy

This task describes how to resolve all NDB alerts related to an alert policy.

About this task

Perform the following steps to resolve all alerts related to an alert policy in NDB.

Procedure

1. From the drop-down list of the main menu, select **Alerts**.
2. Go to **Policies** and select an alert policy.
3. Click **Manage Alerts**.  
The **Manage Alerts** page appears.
4. Click **Resolve All Alerts**.  
All the alerts related to the alert policy are resolved.

### Resolving All Alerts for a Database or Database Server VM

This task describes how to resolve all NDB alerts related to a database or database server VM.

About this task

Perform the following steps to resolve all alerts related to a database or database server VM in NDB.

Procedure

1. From the drop-down list of the main menu, select **Databases** or **Database Server VMs**.
2. For databases, select **Sources** or **Clones**. For database server VMs, select **List**.

3. Select a database or database server VM from the list.  
The summary page for the database or database server VM appears. The **Alerts** tile displays the five most recent alerts.
4. Click **View All Alerts** on the **Alerts** tile.
5. Click **Resolve All Alerts**.  
All the alerts related to the database or database server VM are resolved.

### Resolving All Alerts for a Time Machine

This task describes how to resolve all NDB alerts related to a time machine.

About this task

Perform the following steps to resolve all alerts related to a time machine in NDB.

Procedure

1. From the drop-down list of the main menu, select **Time Machines**.
2. Select a time machine from the list.  
The **Time Machines** home page for the selected time machine appears.
3. From the left menu, click **Alerts**.
4. Click **Resolve All Alerts**.  
All the alerts related to the time machine are resolved.

## Alert Notifications

You can configure NDB to send alert notifications to an email recipient.

### Sending Alert Notifications to an Email Recipient

You can opt to receive notifications in your email when an operation fails or completes successfully. NDB sends these notifications only for entities owned by you.

About this task

Configure a SMTP server to send notifications. You can also enable NDB APIs to receive the notifications on slack.

Perform the following procedure to configure the server and send alert notifications to an email recipient.

Procedure

1. From the drop-down list of the main menu, select **Administration > NDB Service**.  
This step displays information of the NDB Service VMs, Pulse status, NDB Service details, and notifications.

2. Go to **Configure Notifications** widget and click **Update**. The **Update SMTP Server Configuration** window appears.

The screenshot shows the 'Update SMTP Server Configuration' window. It has a title bar with a question mark and a close button. The main area contains several input fields and a security section. The 'Address' field contains 'mail.mycompany.com'. The 'Port' field is a spinner box set to '587'. The 'Security' section has three radio buttons: 'TLS' (selected), 'SSL', and 'None'. The 'Username' field contains 'john.doe@mycompany.com'. The 'Password' field is empty. The 'Sender's Email' field contains 'john.doe@mycompany.com'. The 'Test Message Recipient' field contains 'era\_admin@nutanix.com'. At the bottom, there are four buttons: 'API Equivalent', 'Test', 'Cancel', and 'Update'.

**Figure 74: Updating SMTP Server Configuration**

3. Do the following in the indicated fields.**Address**. Type a domain address.
  - a. **Address**. Type the company email address of the account that you want NDB to use for sending notifications.
  - b. **Port**. Select the port for the SMTP server.
  - c. **Security**. Select either **TLS**, **SSL**, or **None**.

**Note:** Select **None** to configure unsecured SMTP servers for mail notifications.

  - d. **Username**. Type the username of the account that you want NDB to use to access the server.
  - e. **Password**. Type the password of the username entered.
  - f. **Sender's Email**. Type the email address of the account that you want NDB to use for sending notifications.
  - g. **Test Message Recipient**. Type the email address of the recipient to whom you want to send notifications.
4. Click **Test** to send a test email message to a recipient email address.
5. Click **Update** to update the configuration details.

**Note:** The **Update** button is enabled only if the test email is sent successfully.

6. You can select the operations for which you want to receive email notifications. To select the operations, do the following:

a. At the top-right corner of the screen, select **Profile** from the user drop-down list.

b. Under **Receive email notifications for**, select the following:

» **Alerts.** Select this option to receive notifications for all the operations.

» **Failed Operations.** Select this option to receive notifications for all the failed operations.

**Note:** All the failed operations are not alerts.

» **Successful Operations.** Select this option to receive notifications for the following successful operations: provision a database, clone a database, patch a database server, patch a database server cluster, refresh a clone, upgrade a database server VM, and upgrade database server cluster.

**Note:**

If you choose all the three email notification options, then you will receive mails from the operations (Alerts, Failed and Successful operations) that satisfies the conditions.

If you want NDB to stop sending alert notifications, you can delete SMTP server configuration. To delete the configuration, click **Delete** and then click **Yes**.

# NDB SECURITY

This section caters to architects and database administrators who seek to understand NDB from a security perspective. This section includes information on the following:

- NDB management plane security
- NDB management plane communication
- NDB management plane with database virtual machines (VMs)
- NDB management plane network placement

**Note:** This section covers only the most common security topics.

## NDB Architecture

The NDB management plane operates on one or more VMs running in a Nutanix cluster with at least three Nutanix nodes (physical servers). The default NDB management plane deployment in a single Nutanix cluster consists of one VM that runs front-end services (API, NDB agent, web service) and the back-end service (database or repository).

The following figure presents a logical implementation of NDB DBaaS, including the NDB management VM, the logical components, databases, and one Nutanix cluster.




Figure 75: NDB Architecture


## Multi cluster NDB Architecture

NDB follows the Nutanix scale-out philosophy and can connect to more than one Nutanix cluster. To simplify administration and improve database workload availability, you can place:

- Database server VMs in different Nutanix clusters to provide a clustered database.
- Database snapshots in multiple Nutanix clusters (not only where the database is hosted).
- Database log files in multiple Nutanix clusters (not only where the database is hosted).

In a multi cluster NDB environment, the NDB management plane requires one NDB management agent VM (running only the agent service) in each additional Nutanix cluster it manages.

The following figure expands on the previous one to include the three Nutanix clusters and three NDB management agent VMs.



**Figure 76: NDB Multi-cluster Architecture**

## NDB High-availability Architecture


To increase availability at the NDB management plane layer, you can deploy it in a high-availability configuration. The consumers (users and third-party systems) access a virtual IP address that directs the session to the VM, which acts as the service leader. The following figure describes the logical components of a highly available NDB architecture.

This setup has an additional component: an NDB proxy or API gateway.

- Database server VMs in different Nutanix clusters to provide a clustered database.
- Database snapshots in multiple Nutanix clusters (not only where the database is hosted).
- Database log files in multiple Nutanix clusters (not only where the database is hosted).

In a multi cluster NDB environment, the NDB management plane requires one NDB management agent VM (running only the agent service) in each additional Nutanix cluster it manages.

The following figure expands on the previous one to include the three Nutanix clusters and three NDB management agent VMs.



**Figure 77: NDB High-availability Architecture**

## NDB Security Implementation

Security is a broad topic and there are multiple areas to account for when you design and implement a solution. In this section, we provide information on NDB management plane security that you should use when you make design decisions and operate your DBaaS solution.

NDB development follows the Nutanix security development life cycle (SecDL), which integrates security into every step of product development, rather than applying it as an afterthought.

**Note:** Security for database engines and database server VMs managed by NDB is beyond the scope of this document.

## Account Management

This section describes the accounts that the NDB solution requires, including the accounts for the NDB management plane VMs and the database server VMs and databases managed by NDB. This section also describes the ways to manage these accounts.

### NDB User Account

Use this account when you consume NDB through API and web UI. This account is either an NDB management plane local or active directory-based account.

### NDB Prism Element Account

NDB requires a Nutanix cluster (Prism Element) account with at least the cluster administrator role. You configure this user during initial NDB setup and when you add a new Nutanix cluster to NDB. If you change the user credentials in the Nutanix cluster, you must change them in NDB.

For information on how to change the Nutanix Prism Element user password, see [Updating a Nutanix Cluster](#) on page 97.

## NDB Management Plane VMs Accounts

NDB management plane VMs come with the following built-in Linux OS accounts:

- era
- postgres
- root

Use the era account to access the NDB management plane via SSH and gain root access using sudo. Nutanix recommends that you change the era user password to improve security. Use the following era-server CLI commands to manage the era user password. Depending on your NDB implementation you may have to use one, two, or all three.

```
era-server > security manage_passwords era_server username=era
era-server > security manage_passwords era_agent username=era
era-server > security manage_passwords era_pg_ha_vms username=era
```

The CLI prompts you to enter the new password after the command runs. After you change the password, you don't need to do anything else.

The root and postgres passwords are locked by default on NDB server and NDB agent VMs versions 2.2.0.1 and later. The accounts password might be unlocked on Postgres high availability VMs if you enabled high availability on a version before 2.0, and you might need to update the password for these accounts using the appropriate CLI command.

## NDB Management Plane Repository Account

The NDB management plane uses a PostgreSQL database as a repository. NDB in a single VM NDB management plane deployment uses the PostgreSQL user postgres. NDB uses both the PostgreSQL users (postgres and era) in an NDB management plane high availability deployment.

Use the following commands to update the required passwords:

```
era-server > security manage_passwords era_repository username=postgres
era-server > security manage_passwords era_repository username=era
```

The update operation changes the user passwords in the PostgreSQL database, configures files in an encrypted format, and updates corresponding metadata entries.

## Database Server VM Nutanix Database Service Agent Account

Each database server VM is managed through an NDB agent. The NDB management plane creates a user account for the agent that connects the database server VM to NDB and to authenticate the database server VM to NDB. The database server VM doesn't use a regular password; instead, it creates an access key, including a key ID and a key value, for the user. The database server VM then uses the username plus the access key when it calls the NDB API during NDB authentication.

**Note:** The access key is protected using symmetric encryption and stored in the NDB repository.

## Database Server VM User Account

One user account is required in every database server VM managed by NDB. A database server VM OS account runs an NDB worker service (Windows Server) or process (Linux) during the life cycle of the database server VM. The account is local for all database engines apart from Microsoft SQL Server which runs on Windows Server, in which case it can be local or domain-based.

The service is responsible for picking up actions scheduled by the NDB management plane and running them. These actions include all day-two operations.

The password or SSH public key for this account is provided during database provisioning.

The NDB worker service account is the only user that can invoke NDB CLI commands.

## Linux-Based Database Server VMs

For Linux-based database server VMs, the NDB worker service account must have root permission user rights through sudo. You use this account during registration, provisioning, and various day-two operations.

## Windows Server-Based Database Server VMs

For Windows Server-based database server VMs, the NDB worker service account must have administrator rights and can be local or domain based. You use this account during registration, provisioning, and various day-two operations.

## Nutanix Database Service Worker Service Account Password

There may be situations where you need to change the NDB worker service account password. Follow the appropriate steps to update the password after you've changed it in the database server VM or Active Directory domain.

- If you use Linux, update the password in the NDB repository by updating the database server VM through the API, CLI, or UI.

- API:

```
https://<server_ip>/era/v0.9/dbservers/<dbserver_id>/creds
Method: POST
Payload:
{
 "hostId": dbserver_id,
 "username": ,
 "password": password,
 "label" :
}
```

- CLI:

```
era > dbserver update id=database_server_id creds username=user_name set_password
```

- UI:

To update the password from the UI, see [Updating a Database Server VM or Database Server Cluster](#) on page 129.

- If you use Windows Server, update the password in the NDB repository through the API, CLI or UI:

- API:

```
https://<server_ip>/era/v0.9/dbservers/<dbserver_id>/creds
Method: POST
Payload:
{
 "hostId": dbserver_id,
 "username": ,
 "password": password,
 "label" : admin_creds
}
```

- CLI:

```
era > dbserver update id=database_server_id creds username=user_name set_password
label=admin_creds
```

- UI:

If the database server VM is part of a windows domain, update the corresponding user password in the NDB windows domain profile.

Update Windows Domain Profile

?

×

Profile Name

NTNXLAB

Profile Description

ntnxlab.local domain profile

Domain to Join (FQDN)

ntnxlab.local

☐ Publish

Organizational Units

☐ Custom OU Path for Server Clusters 

i

☐ Custom OU Path for Database Servers VMs 

i

Domain Account with Permission to Join Computer to the Domain

☐ Update Credentials

Username 

i

administrator@ntnxlab.local

Password

SQL Service Startup Account

☐ Update Credentials

☐ group Managed Service Account (gMSA)

User Name

administrator@ntnxlab.local

Password

NDB Worker Service Account

☐ Update Credentials

User Name 

i

administrator@ntnxlab.local


Password

Cancel

Update

**Figure 78: Update Windows Domain Profile**

If the database server VM isn't part of a Windows Domain, navigate to the NDB Worker Log On tab of the Windows Server service and enter the new password.



**Figure 79: NDB Worker Properties**

If you use the NDB cloning ability to create a database and database server VM from snapshots taken before the password change, the procedure fails. Instead, use NDB cloning to create a database in an existing database server.

NDB maintenance and day-two operations fail if you don't follow the appropriate password change procedure.

### **Database or Database Instance Account**

Because NDB provides life-cycle management for the database or database instance, it needs a database or database instance account with the highest database permissions (user postgres, root, or sa, depending on your database engine).

You can define the password for these accounts during database provisioning. The following sections describes the steps to update the database instance passwords for various database engines.

### **Updating Passwords for PostgreSQL, MongoDB, MariaDB, and MySQL**

Perform the following steps to update the database instance account passwords for PostgreSQL, MongoDB, MariaDB, and MySQL.

- Update the user password in the database instance. NDB cannot perform this operation.

**Note:** For MariaDB, change the passwords for both root@localhost and root@<hostname>.

- Update the NDB metadata to reflect the database instance password change by using the CLI or API.

Get the ID of the database instance:

```
era > database list
```

```
Databases :
+-----+-----+-----+-----+
Type	Id	Name	Clustered
a97612cf-481e-436c-aa2e-0fa036eccebb	mongodb_clustered	Yes	
mongodb_database	2f090ad5-ee66-4	DELETED_FROM_CL	
	8ef-a7b8-9b9bfc	USTER	
	577e23		
+-----+-----+-----+-----+
Total records: 1
```

To update the NDB metadata from the CLI, use the following command:

```
era > database update creds id=database_id username=user_name set_password
```

To update the NDB metadata from the API, use the following API request.

```
URL: https://<era-server-ip>/era/v0.9/databases/<database-id>
Method: PATCH
Payload:
{
 "resetCredential": True,
 "credential": {
 "userName": "<username>",
 "password": "<password>"
 }
}
```

**Note:** For MongoDB, MariaDB, and MySQL:

- The clone operation on new and already existing VMs works fine if the user chooses a snapshot or PITR which is taken after changing the database instance password.
- The clone operation on new and already existing VMs will fail if the user chooses a snapshot or PITR from a time before the password is changed.

## Updating Passwords for SQL Server

If the SQL Server instance is using Mixed authentication, perform the following steps to update the database instance account password.

To update the password for a SQL Server instance using Windows authentication, see [Database Server VM User Account](#) on page 169.

- Update the user password in the database instance. NDB cannot perform this operation.

- Update the NDB metadata to reflect the database instance password change by using the CLI or API.

Get the ID of the database server VM by listing the database server VMs:

```
era > dbserver list
```

Use the following command to update the metadata for the SQL service account.

```
era > dbserver update id=database_server_id creds username=user_name set_password
label=service_account_creds
```

Use the following command to update the metadata for the SQL Server sysadmin account password.

```
era > dbserver update id=database_server_id creds username=user_name set_password
label=sysadmin_creds
```

If the database server VM is using a Windows domain profile, use the following API request to update the metadata. This API request ensures that the change applies to all the VMs that are associated with the Windows domain profile.

```
https://<server_ip>/era/v0.9/profiles/<profile_id>/versions/<profile_version_id>
Method: PUT
Payload:
{"name": "<profile_name>",
 "description": "", "published": <true/false>,
 "Properties":
 [{"name": "DB_SERVER_OU_PATH", "value": "", "secure": false,
 "description": "Custom OU path for database server VMs"},
 {"name": "DOMAIN_NAME", "value": "<domain_name>", "secure": false, "description": "Name of
 the Windows domain (in FQDN format)"},
 {"name": "DOMAIN_USER_NAME", "value": "", "secure": false, "description": "Username with
 permission to join computer to domain"},
 {"name": "DOMAIN_USER_PASSWORD", "value": "", "secure": false, "description": "Password for
 the username with permission to join computer to domain"},
 {"name": "CLUSTER_OU_PATH", "value": "", "secure": false, "description": "Custom OU path for
 server clusters"},
 {"name": "SQL_SERVICE_ACCOUNT_USER", "value": "", "secure": false, "description": "Sql
 service account username"},
 {"name": "SQL_SERVICE_ACCOUNT_PASSWORD", "value": "", "secure": false, "description": "Sql
 service account password"},
 {"name": "ERA_WORKER_SERVICE_USER", "value": "", "secure": false, "description": "Era worker
 service account username"},
 {"name": "ERA_WORKER_SERVICE_PASSWORD", "value": "", "secure": false, "description": "Era
 worker service account password"},
 {"name": "RESTART_SERVICE", "value": false, "secure": false, "description": "Restart sql
 service on the dbservers"},
 {"name": "UPDATE_CREDENTIALS_IN_DBSERVERS", "value": "true", "secure": false, "description": "Update
 the credentials in all the dbservers"},
 {"name": "ALLOW_ERA_WORKER_SERVICE_ACCOUNT_OVERRIDE", "value": false, "secure": false, "description": "
 override of era worker service account in provisioning workflows"},
 {"name": "IS_GMSA_SQL_SERVICE_ACCOUNT", "value": "false", "secure": false, "description": "The
 sqlservice startup account is a GMSA account."},
 {"name": "ALLOW_SERVICE_ACCOUNT_OVERRIDE", "value": false, "secure": false, "description": "Allow
 override of sql service account in provisioning workflows"}]}
```

If the database server VM is not using a Windows domain profile, use the following API request to update the metadata:

```
https://<server_ip>/era/v0.9/dbservers/<dbserver_id>/creds
Method: POST
Payload:
{
 "hostId": dbserver_id,
 "username": username,
 "password": password,
}
```

```
"label" :label
}
```

Replace `label` in the above API request with `service_account_creds` or `sysadmin_creds` as applicable.

## NDB GUI and CLI

This topic describes the procedure to reset the NDB passwords for the GUI and CLI.

Before you begin

You must have SSH access to NDB.

### Procedure

1. To reset the NDB GUI password, do the following:
  - a. SSH to NDB.
  - b. Execute the following command to enter the era-server CLI:

```
era-server
```

- c. Execute the following command to reset the password:

```
era-server> security password reset
```

- d. Enter the new password at the prompt.

2. To reset the NDB CLI password, do the following:

- a. SSH to NDB.
  - b. Execute the following command to reset the password:

```
sudo passwd era
```

- c. Enter the new password for the era user at the prompt.

**Note:** You must also reset the passwords for NDB agent, NDB HA VMs, and NDB repository based on the implementation. For more information see, [NDB Management Plane VMs Accounts](#) on page 169.

## Authentication

There are two options for authenticating users who consume the NDB service via API and web UI:

- NDB local user authentication.
- Active Directory user- and group-based authentication.

The initial active directory authentication configuration requires an active directory service account to verify connectivity between NDB and active directory, search for users and groups, and add users and groups to the NDB repository. The active directory service account isn't used for NDB user authentication. Instead, NDB uses the user authentication credentials, username, and password to perform a bind against active directory. active directory groups and users added to NDB through active directory groups are refreshed every 24 hours, and you can use the UI to manually refresh when needed.

Both local and active directory users can use authentication tokens when they consume the NDB API. They can request the tokens through the API using their username and password.

NDB stores the names of local and active directory-based users and groups in the NDB repository. For the local accounts, NDB stores the users and passwords in the NDB repository, and the passwords are stored using symmetric encryption.

## Network Placement

NDB is placed on a network with the following capabilities, considering company network definitions and regulations for making services available to consumers:

- Allows consumers to access NDB through HTTPS (port 443). Consumers can be users and other systems that consume NDB, such as Nutanix Calm or ServiceNow.
- Allows communication between NDB and Nutanix clusters.
- Allows communication between NDB and the database server VMs it manages.
- Allows communication between NDB administrators and NDB through SSH (port 22) and HTTPS (port 443).
- Allows communication between NDB and Active Directory (if using Active Directory).
- Allows communication from NDB to NTP server.
- Allows NDB to send telemetry data to [insights.nutanix.com](https://insights.nutanix.com).
- Allows NDB to communicate with the Nutanix portal to download upgrades.
- Allows communication between NDB components (applies to NDB multi cluster and NDB high availability deployments).
- Allows NDB and vCenter Server to communicate (if using ESXi).

For network port specifications, see [NDB Network Requirements](#).

## Role-based Access Control

NDB employs role-based access control (RBAC) to govern access to the NDB system. RBAC allows system administrators to restrict user access and limit the operations they can perform. There are four predefined roles in NDB and an option to create custom roles based on requirements and the available user rights.

For more information, see [Role-based Access Control](#) on page 74.

## Patch Management

NDB is delivered as a virtual appliance. Don't manually upgrade the software included in the appliance or install additional software. The only exception is if Nutanix support recommends and performs such an action.

The supported method for upgrading NDB is through the NDB UI, by navigating to the Administration section, then selecting Maintenance.

**Note:** There is no mechanism in the NDB appliance that actively prevents a manual software upgrade or software installation.

## SSL Server Certificate Authentication

The NDB management plane generates a self-signed certificate during deployment. Change this certificate to a valid certificate to increase security and ensure a trustworthy connection between NDB and the client. For more information, see [Configuring an SSL Certificate](#) on page 38.

### Extending the Self-signed Certificate

Follow the steps below to extend the expiry of a self-signed certificate.


Launch the NDB server CLI:

```
era-server
```

Run the following command to extend the self-signed certificate.

```
era-server > security ssl extend_self_signed_certificate
```

You can use the certificate viewer utility in your browser to view the validity period of the self-signed certificate.



**Figure 80: Self-Signed Certificate Expiry**

## Entity Sharing

The user who created an entity in NDB (such as a database, database server VM, or time machine) owns and manages it. You can share the following entities with a user or an active directory group:

- Databases
- Database clones
- Database server VMs
- Database server clusters
- Snapshots
- Time Machine

Depending on use case, you can share the entity with different access levels and user rights according to:

- View access
- Use access
- Manage access
- Full access

NDB doesn't support all access and user rights levels for all entities.

In addition to shared entities, NDB has the following features:

- Any NDB user can consume all published profiles.
- NDB users with the Super Admin role can manage all databases, database server VMs, SLAs, and time machines.

Entity ownership changes when a user or an active directory group is deleted but the entity still exists in NDB. Specify the new owner during the user delete process.

For more information on entity sharing policies, see [NDB Entity Sharing](#) on page 85.

# NDB UTILITIES

## Database Server VM Registration Pre-requirement Checks

To register a database server VM to NDB, the database server VM must satisfy all the required software and configuration dependencies. When you are registering a database server VM, NDB performs a few checks to ensure if database server VM meets all the requirements. If database server VM fails to meet all the required dependencies, NDB displays the following validation report.

| Server Response : |                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------|
| Message           | Details                                                                                     |
| Error Code        | ERA-HST-0000003                                                                             |
| Message           | Not all the dependencies/configurations required by Era are present on the VM (10.46.18.64) |
| Reason            | Not all the dependencies/configurations required by Era are present on the VM (10.46.18.64) |
| Remedy            | Please make sure all the dependencies/configurations required by Era are present on the VM  |

== Below dependencies and configurations are required to onboard the DBServer with Era Service (please refer report below):

| Missing Software Dependencies : |                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------|
| Software Name                   | Required Action                                                                  |
| lvcreate                        | Make sure the LVM and related utilities are installed on the target VM           |
| gcc                             | You can try installing gcc using 'sudo yum install gcc -y' on the target VM      |
| unzip                           | You can try installing gcc using 'sudo yum install unzip -y' on the target VM    |
| readline                        | You can try installing gcc using 'sudo yum install readline -y' on the target VM |

| Missing Configurations : |                                                                |
|--------------------------|----------------------------------------------------------------|
| Configuration            | Required Action                                                |
| crontab                  | Make sure crontab is enabled for the ERA user on the target VM |

**Figure 81: Pre-requirement Validation Report Sample Output**

(Optional) On NDB CLI, you can run the following command to ensure if database server VM meets all the requirements to be registered with NDB.

```
era > dbserver add ip=<ip_address> nx_cluster_name=<nx_cluster_name>
info_file=<info_file>
```

**Note:** Ensure that *tar* is installed on the database server VM.

## Downloading Pre-requirement Script (CLI)

You can download and run the pre-requirement script to register a database server to NDB service.

About this task

To download the pre-requirement script using CLI, perform the following procedure.

Procedure

1. Launch NDB CLI.
2. Download the pre-requirement script by using the following command:

```
era > era_util download type=pre_checks platform=<platform_type>
output_file_path=<path_to_output_file>
```

For example, use the following command for downloading the Linux pre-requirement script:

```
era > era_util download type=pre_checks platform=linux output_file_path=/home/era/
linux_prereqs.sh
```

```
** The pre-requirement script file for linux has been downloaded at /home/era/
linux_prereqs.sh
```

Use the following command for downloading the Windows pre-requirement script:

```
era > era_util download type=pre_checks platform=windows output_file_path=/home/era/
windows_prereqs.ps1
** The pre-requirement script file for linux has been downloaded at /home/era/
windows_prereqs.ps1
```

You can also download the pre-requirement script without specifying the `output_file_path` parameter. For example, use the following command for downloading the Linux pre-requirement script:

```
era > era_util download type=pre_checks platform=linux
** The pre-requirement script file for linux has been downloaded at /tmp/
era_linux_prechecks.sh
```

Use the following command for downloading the Windows pre-requirement script:

```
era > era_util download type=pre_checks platform=windows
** The pre-requirement script file for windows has been downloaded at /tmp/
era_windows_prechecks.ps1
```

## Downloading Pre-requirement Script (GUI)

You can download and run the pre-requirement script to register a database server to NDB service.

About this task

To download the pre-requirement script from the user interface, perform the following procedure.

Procedure

1. At the top-right corner of the screen, select **Downloads** from the user drop-down list.
2. Under **Check Database Server**, select either **Download for Linux** or **Download for Windows**.

This step downloads the script file to your local machine.

3. (Optional) Under **Check Database Server**, select **Download Sudo Logger for Linux**.

This step downloads the script file to your local machine. Run this script along with the previous script to log all sudo commands used in the Linux check.

## Running the Pre-requirement Script (Linux)

About this task

The pre-requirement script helps you validate a database server VM before registration. This script generates a validation report listing all the dependencies along with their installation status. To print NDB pre-requirement validation report in Linux, perform the following procedure.

Procedure

1. Copy the pre-requirement file to the target VM or copy the contents of the file and paste them in a script inside the VM.

2. Run the script by using the command:

```
$./era_linux_prechecks.sh [-t|--database_type <database_type>] [-c|--cluster_ip <cluster_ip>] [-p|--cluster_port] [-d|--detailed] [-s|--restricted_sudo]
```

Replace the variables with the following:

- Replace *database\_type* with the database name such as *oracle\_database*, *postgres\_database*, *mariadb\_database*, or *mysql\_database*.
- Replace *cluster\_ip* with the IP address of the cluster.

**Note:** In the above command, *-t* is a required argument (*database\_type*) while *-c*, *-p* and *-d* are the optional argument. It is recommended that you provide *-c* (*cluster\_ip*) to validate whether Prism APIs are accessible from the database server. If *-c* is not provided, then at the end, a warning message Cluster API was not provided. Couldn't go ahead with the Prism API connectivity check. Please ensure Prism APIs are callable from the host. is displayed.

This step displays the following pre-requirements validation report.

```

|-----|
Era Pre-requirements Validation Report
General Checks:
1) Username : era
2) Package manager : yum
2) Database type : oracle_database

Era Configuration Dependencies:
1) User has sudo access : YES
2) User has sudo with NOPASS access : YES
3) Crontab configured for user : YES
4) Secure paths configured in /etc/sudoers file : YES

Era Software Dependencies:
1) GCC : NO
2) readline : YES
3) crontab : YES
4) lvcreate : NO
5) lvscan : YES
6) lvsdisplay : YES
7) vgcreate : YES
8) vgsan : YES
9) vgsdisplay : YES
10) pvcreate : YES
11) pvscan : YES
12) pvdisplay : YES

Summary:
This machine does not satisfy all of the dependencies required by Era.
It can not be onboarded to Era unless all of these are satisfied.

Note: You can run this script using '-detailed' option to know the complete details
|-----|
```

**Figure 82: NDB Pre-requirements Validation Report Sample Output**

3. Install the missing dependencies, if any, to onboard the database server VM to NDB.

## Downloading the Uninstall Script (GUI)

If cleanup operations fail and thus NDB fails to remove its components from the database server, you can download the uninstall script. You need to copy and execute this script on your registered database server to uninstall all NDB components.

About this task

To download the uninstall script from the user interface, perform the following procedure.

Procedure

1. At the top-right corner of the screen, select **Downloads** from the user or admin drop-down list.
2. Under **NDB Uninstall**, select either **Download for Linux** or **Download for Windows**.  
This step downloads the script file to your local machine.

# NDB UNINSTALLATION

---

## Uninstalling NDB

About this task

Perform the following procedure to uninstall and remove NDB.

Procedure

1. In the drop-down list of the main menu, select **Databases**. Go to **Sources**, select all the databases, and click **Remove** to remove all the databases from NDB.  
Clicking **Remove** starts the database removal operation. You can view the progress of the operation on the **Operations** page.
2. In the drop-down list of the main menu, select **Databases**. Go to **Clones**, select all the database clones, and click **Delete** to remove all the database clones from NDB.
3. In the drop-down list of the main menu, select **Time Machines**. Select the time machines for all the source databases and click **Delete** to remove all the database clones.
4. In the drop-down list of the main menu, select **Database Server VMs**. Go to **List**, select all the database server VMs, and click **Remove** to remove all the database server VMs from NDB.  
If a dialog box appears, select the check box to delete all the VMs from the cluster if the VMs are not needed.
5. In the drop-down list of the main menu, select **Profiles**. Go to **Softwares**, select all the software profiles, and click **Delete** to remove all the profiles from NDB.  
Wait for all the database and database server VM removal operations to finish.

**Note:** You cannot delete the out-of-the-box (OOB) profiles.

6. From Prism, delete the NDB VM from the cluster.

**Note:**

- Delete all the NDB VMs if the deployment is an HA configuration.
- Delete the NDB agent VMs if multi-cluster is enabled.

# NDB SHUTDOWN

---

## Shutting Down NDB

This topic describes the steps to safely shut down NDB HA and non-HA configurations.

### Procedure

1. For an NDB non-HA configuration, you can shut down the setup in a single step by running the script:

```
python3 /home/era/era_server_config/server_scripts/shutdown_server.py
```

**Note:** After you run the script, NDB waits for the running operations to complete and does not accept the submission of new operations.

The NDB non-HA setup shuts down.

2. For an NDB HA setup, power off the API server nodes first:
  - a. Go to **Administration**, Select **NDB Service**, and identify the standby nodes and the leader nodes. Alternately, you can issue the following command to identify the nodes:

```
era-server > high-availability list
```

- b. Shut down the standby nodes by running the following script:

```
python3 /home/era/era_server_config/server_scripts/shutdown_server.py
```

- c. Shut down the leader node by running the following script:

```
python3 /home/era/era_server_config/server_scripts/shutdown_server.py
```

3. Power off HA proxy nodes by using the command:

```
shutdown -P now
```

4. Power off repository nodes by following the steps below:

- a. Check Patroni status:

```
systemctl status patroni
```

- b. Check the status of repository nodes:

```
patronictl -c /etc/patroni/patroni.yml list
```

- c. Stop Patroni on standby nodes by using the command:

```
sudo systemctl stop patroni
```

- d. Stop Patroni on leader nodes by using the command

```
sudo systemctl stop patroni
```

- e. Power off the repository nodes in any order by using the command:

```
shutdown -P now
```

The NDB HA setup shuts down.

# COPYRIGHT

---

Copyright 2024 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.