

Nutanix Database Service - NDB Architecture


[PDF generated May 06 2026. For all recent updates please see the Nutanix Bible releases notes located at https://nutanixbible.com/release_notes.html. Disclaimer: Downloaded PDFs may not always contain the latest information.]

NDB Architecture

NDB is a software solution that can be deployed on Nutanix clusters running AHV or ESXi. NDB can be deployed with a single control plane VM, across multiple Nutanix clusters, with high availability (HA) configurations, and multi-cluster with HA configurations. However NDB is installed and configured, the database administration functions - provisioning, cloning, copy data management, patching, and backup - remain the same. In an HA deployment, the API, GUI, and metadata repository are resilient and will stay online during upgrades or a failure of a single VM. You can learn more about the configuration of the NDB control plane from the [next page](#) in this guide.

NDB Single Control Plane VM Deployment

NDB can be deployed as a single control plane VM; this install is ideal for situations where you do not need data resiliency for NDB's API, GUI or metadata. However, the Nutanix Cloud Platform hosting the VM is highly available and resilient. During upgrades, NDB functionality will be offline during parts of the upgrade process.



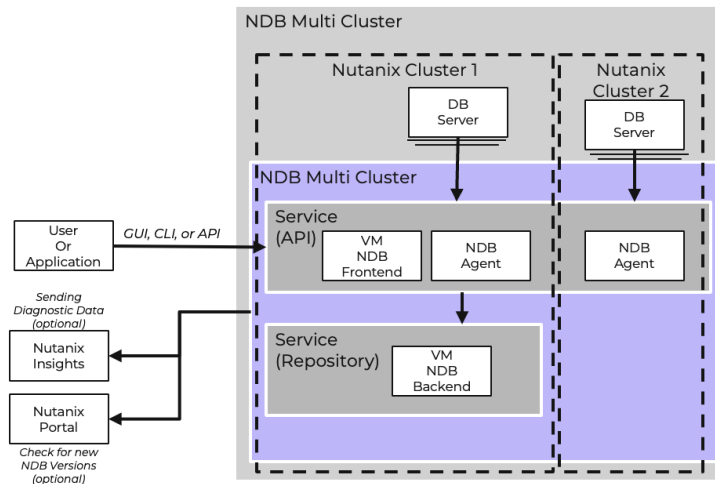
NDB Single Control Plane VM Components

- NDB Frontend
 - The frontend server provides the web interface and serves as an API endpoint.
- Nutanix Database Agent
 - The Agent is a service that schedules and monitors the NDB-related operations on database server VMs.
- Nutanix Database Service Backend The backend is the repository database that manages metadata such as deployed databases, database VMs configured, time machines, snapshots, log catch-up operations, retention policies, etc.

NDB Multi-Cluster Deployment

If you have multiple Nutanix clusters, you can deploy NDB across two or more clusters, which can address several use cases, including the following.

- Create and manage databases across multiple Nutanix clusters for high availability and disaster recovery (for example, Microsoft SQL Server availability groups or PostgreSQL high availability).
- Manage time machine data availability across all registered Nutanix clusters using data access management policies.
- Administer your production and nonproduction database instances across geographically distributed Nutanix clusters from a single panel.
- Clone databases for test and development from source databases across multiple sites.



NDB High Availability (HA) Deployment

To ensure high availability with auto-failover for the NDB control plane, Nutanix requires three Nutanix clusters, running AHV or ESXi, 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 25ms. Network latency greater than 25ms between selected clusters is not supported. 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.

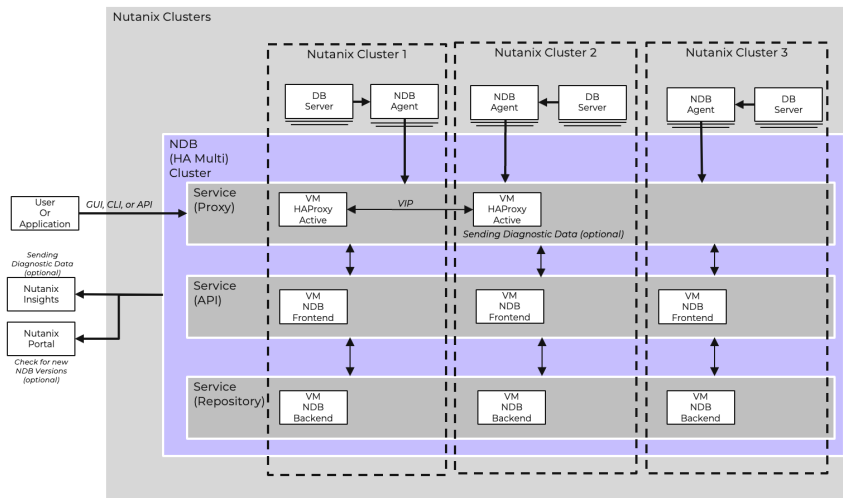
Nutanix Database Service High Availability Components

- API server
 - The stateless REST API gateway for the NDB control plane is used for the API frontend and the user interface.
- Repository
 - A PostgreSQL database runs the data persistence layer of the NDB control plane and stores NDB metadata. In an HA configuration, NDB uses a Patroni failover manager with three PostgreSQL replicas (primary, synchronous, and asynchronous).
- HAProxy
 - A high availability load balancer and proxy server that routes connectivity to API and repository servers.
- VIP
 - A floating virtual IP address (VIP) spans HAProxy nodes to provide high availability to user connections. All front-end user connections go through the VIP. NDB uses Keepalived as the routing software to float a VIP across HAProxy nodes.
- NDB agent
 - A VM in the NDB control plane management stack is created when you enable NDB for multiple clusters. All remote Nutanix clusters added to NDB have an instance of the NDB agent that manages, schedules, and runs database operations for that cluster.

- NDB database agent

- An NDB service runs on a database server VM. It schedules and monitors NDB-related operations on database server VMs.

When running HA on multiple clusters, the VLAN for the HA VMs needs to be shared on both clusters, or a stretched VLAN between the HA Proxy's servers is required to place them on different Nutanix Clusters.



IP Address Requirements and VM Distribution for Different NDB Topologies

The [NDB Administration Guide](#) has the most up-to-date requirements for IP address and VM distribution to fit the needs of your Nutanix environment. The NDB HA reference architecture link below is also a good guide to see what best works in your environment.

Pro tip

* In multi-cluster topologies:

* NDB server: In a non-HA configuration, the NDB server hosts all NDB service control plane components. 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.

High Availability Reference Architecture

Follow this link to find a [reference architecture for NDB HA](#). This reference architecture is a comprehensive guide to configuring and deploying NDB High Availability (HA) with Nutanix AHV across multiple Nutanix clusters.