

# ClustrixDB Reference Server Configurations

## AWS and On-Premises Deployments

- [Deployment in Amazon AWS](#)
- [Deployment in Google Cloud Platform](#)
- [Bare Metal Deployments](#)

ClustrixDB is designed to achieve high levels of transactional scale with high availability. Its ability to deliver both of these is directly linked to the capabilities of the server hardware used. This document provides server recommendations for production ClustrixDB clusters running in AWS and on bare metal.

One of the key benefits of the ClustrixDB is that it runs on a variety of server platforms. If you are considering different server specifications than outlined here, we are happy to provide guidance.

### Minimum System Requirements

- Minimum of 3 servers required for a ClustrixDB cluster
- Operating System: RHEL or CentOS 7.4+
- We recommend between 8 and 32 CPU cores per server
- SSDs for database storage
- 64GB of RAM

## Deployment in Amazon AWS

This section contains the AWS EC2 instance configurations recommended for production ClustrixDB clusters running on AWS EC2 instances.

### Recommended Instance Types

- Recommended:
  - I3 (1900GB NVMe SSD storage): 2xlarge (8-cores), or 4xlarge (16-cores)
- Previous Generation:
  - C3 (160GB SSD storage): 2xlarge (8-cores), or 4xlarge (16-cores)
  - I2 (1600GB SSD storage): 2xlarge (8-cores), or 4xlarge (16-cores)
- All instances in a ClustrixDB cluster must be the same type

### Disk Storage

- Only Ephemeral SSD may be used for database data files
  - EBS attached storage is NOT supported for volumes containing database data files
- EBS volumes may be used for volumes containing log files, or files managed outside of ClustrixDB

### Networking

- Enhanced Networking must be enabled
- ClustrixDB instances must be in a Virtual Private Cloud (VPC)
- ClustrixDB instances within an Availability Zone should be in the same Placement Group for best performance
- All ClustrixDB instances must be in a Security Group that allows all TCP/UDP traffic within the Security Group (i.e. between ClustrixDB instances)
- ClustrixDB clusters deployed across multiple Availability Zones must be within the same region and have <2ms network latency between all nodes

### Load Balancer

- ClustrixDB leverages a load balancer to present a single IP to the application servers, and spread DB connections equally across all ClustrixDB instances
- Using AWS's Elastic Load Balancer (ELB) is recommended due to its ease of operation
  - The Network Load Balancer offers the best performance, but requires considerations regarding how the Security Groups are configured
  - The Classic Load Balancer is easier to configure with Security Groups, but tends to add latency to traffic passing through it
- Alternatively, a software load balancer (e.g. HAProxy) running on a separate EC2 instance, or directly on the application server is also suitable
- The ELB should be created as an Internal Load Balancer:
  - <http://docs.aws.amazon.com/ElasticLoadBalancing/latest/DeveloperGuide/elb-internal-load-balancers.html>

## Deployment in Google Cloud Platform

This section contains the Google Cloud Platform (GCP) configurations recommended for ClustrixDB clusters running on Google Compute Engine (GCE) virtual machines.

## Recommended Virtual Machines Types

- Recommended:
  - n1-highmem-16
  - n1-standard-16
- Not Recommended:
  - n1-highcpu-16 as it has the same CPU power as the highmem and standard, yet less RAM

Note: Currently Clustrix recommends using 16-CPU VMs, and not using the 8-CPU VMs in GCE

## Disk Storage

- SSD Persistent Disk
- Size of at least 1TB per disk to achieve the maximum of 10,000 IOPS per disk
- Multiple 1TB disks may be used for larger storage requirements per VM

## Networking

- The VPC Network: Firewall rules must be configured to allow all TCP/UDP traffic between all ClustrixDB VMs
- ClustrixDB clusters deployed across multiple Zones must be within the same region and have <2ms network latency between all nodes

## Load Balancing

- ClustrixDB leverages a load balancer to present a single IP to the application servers, and spread DB connections equally across all ClustrixDB instances
- Google Load Balancer:
  - Select: TCP Load Balancing
  - Select: Internal facing
  - Select: Multiple Regions (or not sure yet)
- Alternatively, a software load balancer (e.g. HAProxy) running on a separate GCE VM instance, or directly on the application server is also suitable

## Bare Metal Deployments

This section contains the server specifications recommended for production ClustrixDB clusters running on-premises on physical servers.

### General ClustrixDB Server Recommendations:

- CPUs:
  - 8-32 physical CPU cores per node
    - 16 or 20 physical cores recommended
  - Hyperthreading enabled
- RAM:
  - 64GB or more
- Disks:
  - SATA or NVMe SSDs (for DB data) configured in RAID-0
  - SATA HDD (for OS, logs, etc)
- Network:
  - Separate front and back end networks
  - Back end network should be at least 10Gbps
  - 10-Gigabit Ethernet switch with enough ports to connect each ClustrixDB server (or 2x switches for redundancy)

### Other Recommendations:

**SSD Quality:** Solid state drives (SSDs) are available in various levels of quality. Clustrix recommends using SSDs that have the following characteristics:

- Enterprise-class SSDs (not consumer-class available from Amazon.com)
- Contains a "Power Loss Protection" feature
- Recommend using SSDs manufactured by Intel (every Intel SSD we've seen has Power Loss Protection)

## Example ClustrixDB Server Configuration (Dell)

This is an example server specification using a Dell configuration. Other server vendors including HP and IBM offer servers with very similar specifications.

### Dell PowerEdge R630 Server, No TPM

- Chassis with up to 8, 2.5" Hard Drives, up to 2 PCIe Slots (With Optional Riser)
- CPUs
  - Intel® Xeon® E5-2667 v4 3.2GHz,25M Cache,9.60GT/s QPI,Turbo,HT,8C/16T (135W) Max Mem 2400MHz
  - Upgrade to Two Intel® Xeon® E5-2667 v4 3.2GHz,25M Cache,9.60GT/s QPI,Turbo,HT,8C/16T (135W)
- Memory
  - 4x 16GB RDIMM, 2400MT/s, Dual Rank, x8 Data Width
- RAID Controller
  - PERC S130 RAID Controller
- Disks
  - 2x 960GB Solid State Drive SATA Mix Use MLC 6Gpbs 2.5in Hot-plug Drive, SM863a
  - 2x 1TB 7.2K RPM SATA 6Gbps 2.5in Hot-plug Hard Drive
- Network
  - Embedded NIC: 4x 1Gb, 2x 1Gb + 2x 10Gb, 4x 10Gb
- Other Notes:
  - No Operating System
  - Unconfigured RAID (deployment will use software-level RAID-0)