

# Hitachi Solution for Databases

## Non-Virtualized PostgreSQL Database on Hitachi Advanced Server HA820 G3 using VSP One Block 28

---

Reference Architecture Guide

© 2024 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., Hitachi Vantara, Ltd., or Hitachi Vantara Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

**Notice:** Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

**Notice on Export Controls.** The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html>.

---

## Reference Architecture Guide

Use this reference architecture guide to understand how Hitachi Solution for Databases provides a high-performance, low latency, integrated, converged solution for PostgreSQL Database using Hitachi Virtual Storage Platform One Block.

The Hitachi Unified Compute Platform (UCP) solution for PostgreSQL databases is an engineered, pre-tested, qualified system that delivers unparalleled high performance using a highly reliable and scalable architecture with lower latency in demanding and dynamic PostgreSQL environments.

This Hitachi solution for databases is built on Hitachi Advanced Server and Hitachi Virtual Storage Platform One Block storage systems. Hitachi Advanced Server is equipped with powerful two-socket 5th generation Intel Xeon Emerald Rapids Processors featured with reliability, high availability, and scalability for processing all kinds of workloads.

VSP One Block is the latest storage offering from Hitachi that offers efficient storage for critical OLTP/OLAP databases which fosters database performance because of its unique capabilities, such as best security, scalability, data compression, easy management, energy efficiency, and compatibility with hybrid cloud.

The VSP One Block storage system with NVMe SSDs is used in this reference architecture to run a dedicated non virtualized PostgreSQL database on RedHat Linux operating system.

To configure the solution environment, the following components are used:

- Hitachi Virtual Storage Platform One Block 28 (VSP One Block 28) is used as a storage resource.
- Hitachi Advanced Server HA820 G3 (HA820 G3) with Intel® Xeon® 8568Y+ Emerald Rapids Processor 48c CPUs for storage and computing resources.

This solution provides the flexibility to select storage and compute resources based on unique requirements. Deploy small databases as well as very large databases, depending on resource availability.

This document is for the following audiences:

- Database administrators
- Storage administrators
- System administrators
- IT professionals responsible for planning and deploying a PostgreSQL Database solution

To use this document, you need familiarity with the following:

- Hitachi Virtual Storage Platform One Block
- Hitachi Advanced Server HA820 G3 servers
- Storage Area Networks

- PostgreSQL DB
- RedHat Enterprise Linux
- HammerDB



**Note:** Testing of this configuration was in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

## Solution overview

The Hitachi solution for PostgreSQL databases on Hitachi Unified Compute Platform (UCP) is an engineered, pre-tested, and qualified system to provide high performance and high reliability in demanding and dynamic PostgreSQL environments.

This Hitachi UCP solution for databases demonstrates a reference architecture for running a PostgreSQL database using Hitachi Advanced Server with Hitachi Virtual Storage Platform One Block. It leverages the latest capabilities of data storage and integrations to create, protect, and manage on-premises database clusters.

When we think of open-source RDBMS in the market, PostgreSQL has the most suitable, robust, and consistently proven track record database. It is a well-known and highly demanding database for its performance capabilities, and ability to handle complex queries and transactional workloads in diverse conditions.

To underline PostgreSQL database performance using VSP One Block storage, we used the open-source HammerDB benchmarking tool, popular among industry standards. HammerDB is derived from TPC specifications. It executes two types of workloads, namely transactional and analytical, to measure performance.

### Business benefits

The following are some of the benefits of this reference architecture:

- Achieve high PostgreSQL Database performance with VSP One Block 28 storage systems.
- Provide a solution for customers who are looking for low I/O latency, high throughput, and minimal response time for PostgreSQL databases.

### High-level infrastructure

Hitachi Solution for Databases with PostgreSQL includes the following components:

- Hitachi Advanced Server HA820 G3 servers
- Hitachi Virtual Storage Platform One Block 28
- Brocade G720 32 Gbps SAN infrastructure
- Cisco 10/25 GbE LAN infrastructure

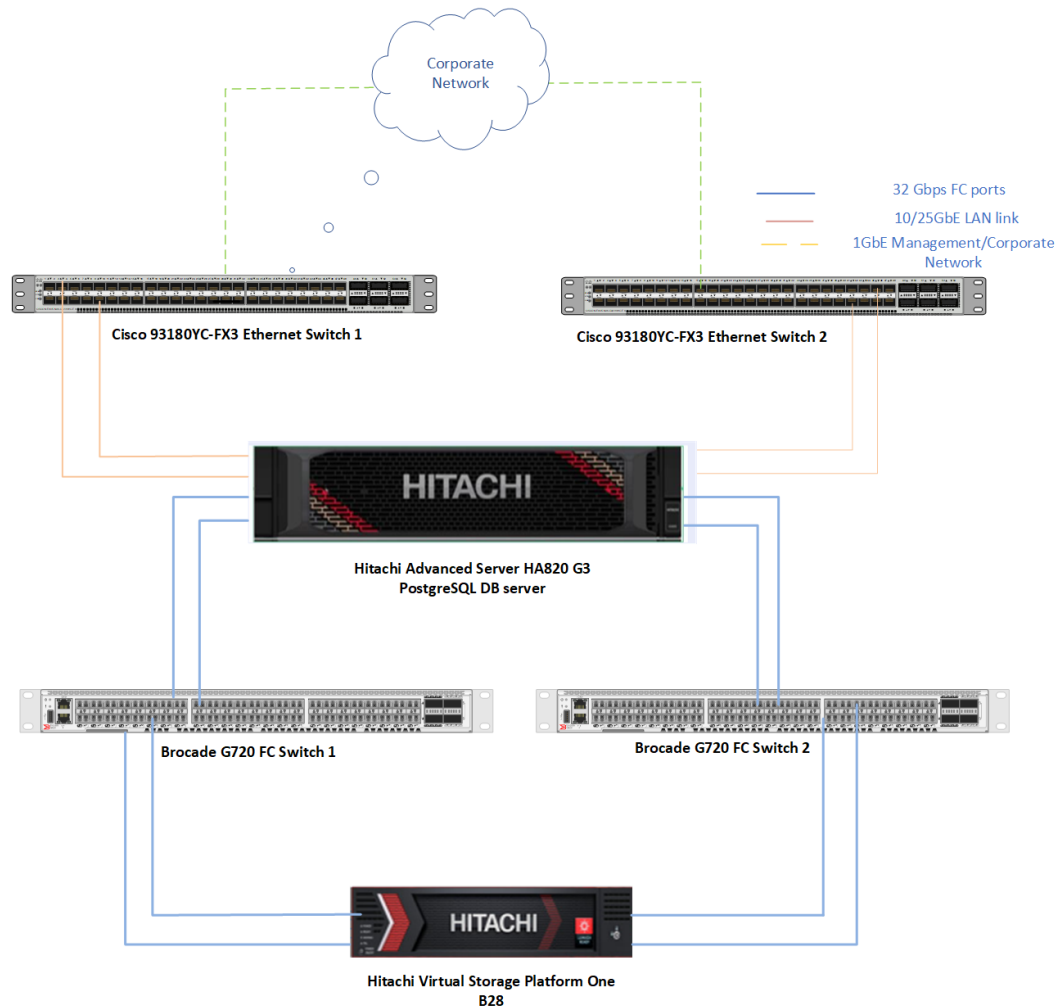
The configuration of Hitachi Virtual Storage Platform One Block 28 and Hitachi Advanced Server HA820 G3 has the following characteristics:

- Fully redundant hardware
- Dual fabric connectivity between hosts and storage



**Note:** While this reference architecture uses the VSP One Block 28, the solution supports the entire Hitachi Vantara storage portfolio.

The following illustration shows the high-level infrastructure network for this solution.



To avoid any performance impact to the production database, Hitachi Vantara recommends using a configuration with the following characteristics:

- A dedicated storage system for the production database
- A dedicated storage system for storing backup data, if needed

The uplink speed to the corporate network depends on the customer environment and requirements. The Cisco Nexus 93180YC-FX3 switch used in this reference architecture can support uplink speeds of 25 GbE, or 100 GbE if higher bandwidth is required.

## Key solution components

The key solution components for this solution are listed in the following tables. Detailed component information is provided in [Product descriptions \(on page 16\)](#).

The following table lists the hardware components used in this solution.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	VSP One Block 28	Two Controllers 4 × 24 Gbps Fibre Channel ports 1024 GB cache memory 24 × 3.8 TB NVMe SSD	A3-02-00-40/03	1
Hitachi Vantara	HA820 G3	2 × Intel Emerald Rapids Processors 8568Y+ CPU @ 2300 MHz 64 GB × 32 DDR5 (2048 GB) RDIMM	iLO 6: 1.58 Mar 22 2024 System ROM: U54 v2.16 (03/01/2024)	2
		2 × SN1700E 64 GB 2p FC HBA	Firmware: 14.2.589.5 Driver: lpfc Driver version: 14.0.0.4	
		2 × Dual Port 25 GbE NIC Intel E810 PCIe cards	Driver: ice Driver version: 5.15 Firmware: 4.30	
Brocade	G720 Fibre Channel switches	24 × 48 port Fibre Channel switches 32 Gbps SFPs	Fabric OS: v9.1.1c	2
Cisco	Cisco Nexus C93180YC-FX3	48 × 10/25 GbE ports 6 × 40/100 Gbps Quad SFP (QSFP28) ports	NXOS: version 9.3.8	2



**Note:** The solution was tested with PCIe and OCP Mezzanine NIC cards. Using all PCIe cards is recommended for consistency and better NIC bonding options. SATADOM, SAN boot, or local boot can be used for the boot option. Customers can choose larger capacity SSDs to fit their business requirements.

The following table lists software components for compute nodes.

Software	Version	Function
RedHat Enterprise Linux	8.6 x86_64 or later	Operating system for compute node
PostgreSQL Database	15.7	Database software
HammerDB	4.11	Benchmarking tool



**Note:** This solution was tested with RedHat Enterprise Linux 8.6 x86\_64, but it can support later RHEL 8.x versions.

## Solution design

This section describes the reference architecture environment to implement Hitachi Unified Compute Platform (UCP) Solution for PostgreSQL Database on bare metal hosts. The environment uses Hitachi Virtual Storage Platform One Block 28 as a storage resource.

The infrastructure configuration includes the following:

- PostgreSQL compute node — There is a host configured as a PostgreSQL database server.
- Storage System — There are vVols mapped to each port that are presented to the server as LUNs.
- SAN Connection — There are SAN connections to connect the Fibre Channel HBA ports to the storage through Brocade G720 switches.

## Storage architecture

This section describes the storage architecture for this solution.

### Storage configuration

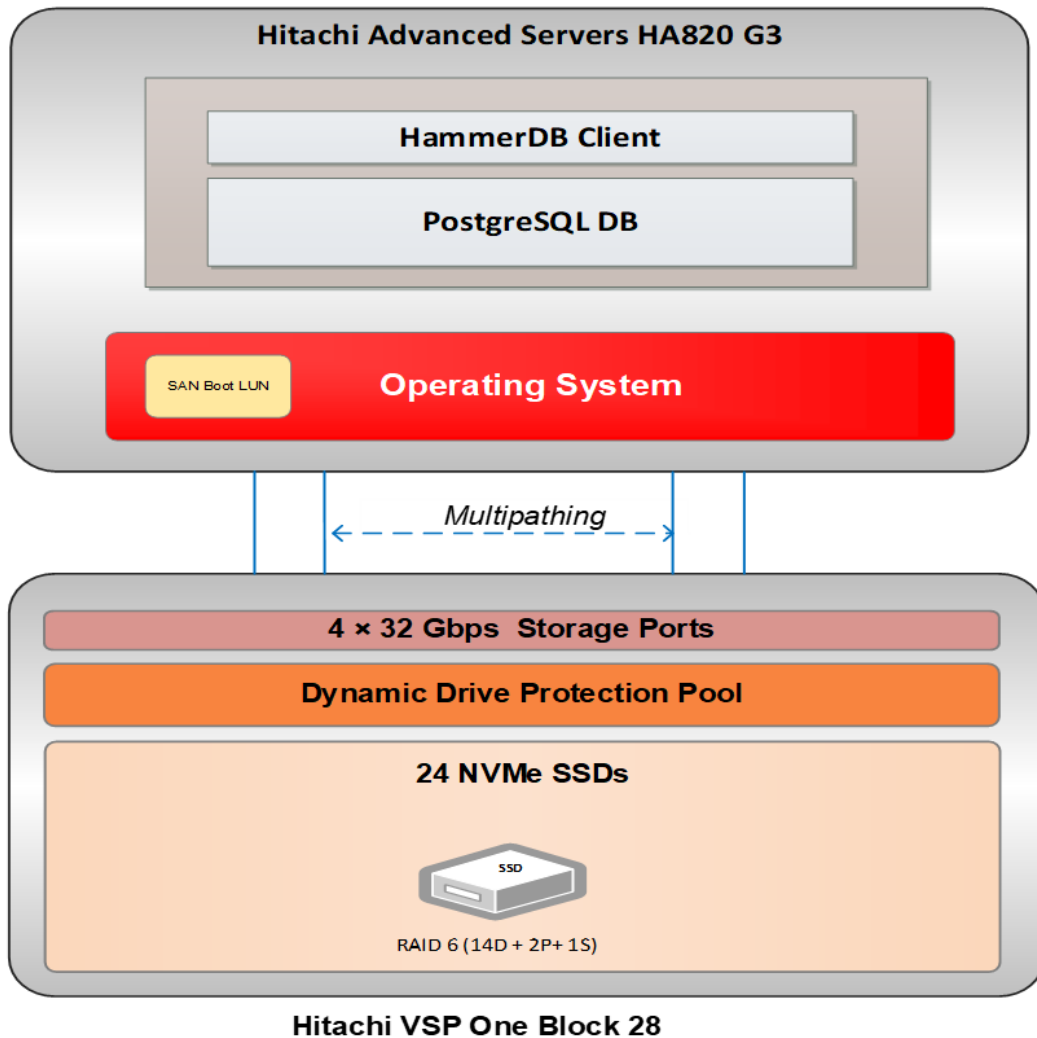
The storage configuration takes into consideration Hitachi Virtual Storage Platform One Block 28 for the design and deployment of database storage. Hitachi introduced VSP One Block storage which eliminates complexity in storage management and makes data storage easy, simple, and secure.

VSP One Block comes with many advanced functionalities, which reduce space usage and increase capacity. It also introduces Dynamic Drive Protection Pool (DDP) for the first time, an implementation in which parity data is stored on the drives instead of on dedicated parity drives. VSP One Block supports dual-parity DDP groups (6D+2P, 14D+2P).

For more details about VSP One Block see [VSP One Block Overview of Provisioning](#).

In this solution, we used RAID 6 (6D+2P+1Spare) configuration for the storage systems with 24 NVMe SSDs with a single DDP pool. Each SSD Drive capacity is 3.8 TB with DDP enabled, with a total usable capacity of 68 TB.

The following illustration shows the high-level storage configuration for this solution.



The following table shows the VSP One Block 28 storage pool configuration used for this solution.

Pool ID	PostgreSQL-Pool
Pool Type	Dynamic Drive Protection Pool
RAID Group	1-1
RAID Level	RAID 6 (14D+2P+1 Spare)
Drive Type	3.8 TB NVMe SSD
Number of Drives	24
Number of LDEVs	24
LDEV Size	3.8 GB
Pool Capacity	68 TB



In this solution, servers use SAN boot with RAID 6 protection. The following table shows the VSP One Block 28 logical storage configuration used in this solution.

LVM disk group	Total number of dynamic provisioning volumes	Dynamic provisioning volume sizes (GB)	Purpose	Storage Ports
DATA_PG	15	200 GB	OLTP, OLAP Application data	1A, 3A, 1B, 3B
N/A	1	300 GB	SAN boot for OS Volumes	

### Database layout

The database layout design uses recommended practices from Hitachi Vantara for VSP One Block 28 for small random I/O traffic such as OLTP and OLAP transactions.

Base the storage design for database layout needs on the requirements of the specific application implementation. The design can vary greatly from one implementation to another, based on the RAID configuration type and number of drives used in the implementation.

The components in this solution have the flexibility to be used in various deployment scenarios to provide the right balance between performance and ease of management for a given scenario.

## Database configurations

**Data and Indexes** — Create multiple LUNs of the required size on VSP One Block 28 storage and map them to the host. While testing we created and assigned a LUN for data. Best practice is to choose a small size LUN so that IO can be spread among different disks to get maximum performance. In this solution, we configured 3 TB of data. To get an optimal storage response, LUNs are configured as 200 GB each.

After LUN mapping to the host, create a physical volume and volume group, and then mount them using a logical volume (LV). Mount the LV on the file system, which can be further used to store application data, WAL files, and all configuration files related to the PostgreSQL Database.

### PostgreSQL initial parameters

The following table shows the PostgreSQL Database settings.

Environment	Value
shared_buffers	128 GB
work_mem	4 GB
max_worker_processes	96

Environment	Value
max_parallel_workers	64
max_parallel_workers_per_gather	16
max_parallel_maintenance_workers	8
effective_cache_size	32 GB
parallel_setup_cost	0
parallel_tuple_cost	0
force_parallel_mode	on

### PostgreSQL disk mappings

The following table shows the details of the disk mappings from the LUNs to the disk groups for PostgreSQL Database size 3 TB.

LV Disk Group	Disks	Mpath disks	LUN Details	Purpose
DATA_PG	DATA1- DATA15	/dev/mapper/xvd[a- o]1	15 × 200 GB	PostgreSQL database application data

### PostgreSQL server configuration

The following table lists the operating system configurations for the PostgreSQL server. We installed the Red Hat Enterprise Linux 8.6 kernel 4.18.0-372.9.1.el8.x86\_64 operating system on the server. Later we configured native multipathing with the following Hitachi configuration. Multipathing is configured to avoid path failure to the storage system.

Server Configuration	Server OS Setting Details
/etc/multipath.conf	<pre> defaults { user_friendly_names yes find_multipaths yes } devices { device { vendor HITACHI product Hi-SDS path_grouping_policy group_by_prio prio alua path_checker readsector0 no_path_retry 6 } } </pre>

Server Configuration	Server OS Setting Details
	<pre> multipaths { multipath { wwid xxxxxxxxxxx0000001 alias data01} } </pre>
Swap space	32 GB

## Server and application architecture

This reference architecture uses one Hitachi Advanced Server HA820 G3 server with Intel Xeon 5th Generation Emerald Rapids Processors. This provides the compute power for the PostgreSQL database to manage complex database queries and a large volume of transaction processing in parallel.

The following table lists the details of the server configuration for this solution.

Server Make and Model	Server Host Name	Role	CPU Type	CPU Core	RAM
Hitachi Advanced Server HA820 G3	Compute node	RedHat Linux 8.6	2 × Intel (R) Xeon(R) 8568Y+ Emerald Rapids Processors 48 CPUs Cores @2.3 GHz 64-bit	96 (2 × 48)	2048 GB (64 GB × 32)

## SAN architecture

The solution uses switched fabric SAN connections. The SAN connectivity is detailed in this section.

Map the provisioned LDEVs to multiple ports on Hitachi Virtual Storage Platform One Block 28. These LDEV port assignments provide multiple paths to the storage system from the host for high availability. This reference architecture uses two dual port SN1700E 64GB 2p FC HBAs per Advanced Server HA820 G3.

### Compute servers

- 4 SAN switch connections are used for VSP One Block 28 Fibre Channel ports.
- 4 SAN switch connections are used for server HBA ports.

The following table lists details of the Fibre Channel switch connect configuration on VSP One Block 28 ports.

Server	HBA	Host Group Name	Host Name	Switch Zone	Storage System	Storage Port	Brocade G720 Switch
HA820 G3 Server	HBA1	CN31	CN31_HBA1_1	CN31_HBA1_1_ASE42_43_1A	VSP One Block 28	1A	29
	HBA2	CN31	CN31_HBA1_2	CN31_HBA1_2_ASE42_43_2A		2A	30
	HBA3	CN31	CN31_HBA2_1	CN31_HBA2_1_ASE42_43_3A		3A	29
	HBA4	CN31	CN31_HBA2_2	CN31_HBA2_2_ASE42_43_4A		4A	30



**Note:** In a production environment, use separate storage ports for the management servers to avoid impact on database performance. Shared storage ports can be used; however, port utilization should be monitored to avoid performance issues in high performance environments.



**Note:** Port assignment is preliminary and can change depending on hardware availability. The final configuration might be different based on the lab equipment that is acquired and performance tuning that is performed during testing.

## Network configuration

This architecture uses the following separate network:

- Public Network — This network provides client connections to PostgreSQL Database and other applications.

### Network configuration

The following table shows management IP addresses for servers, Cisco network switches, Brocade SAN switches, and VSP One Block 28.

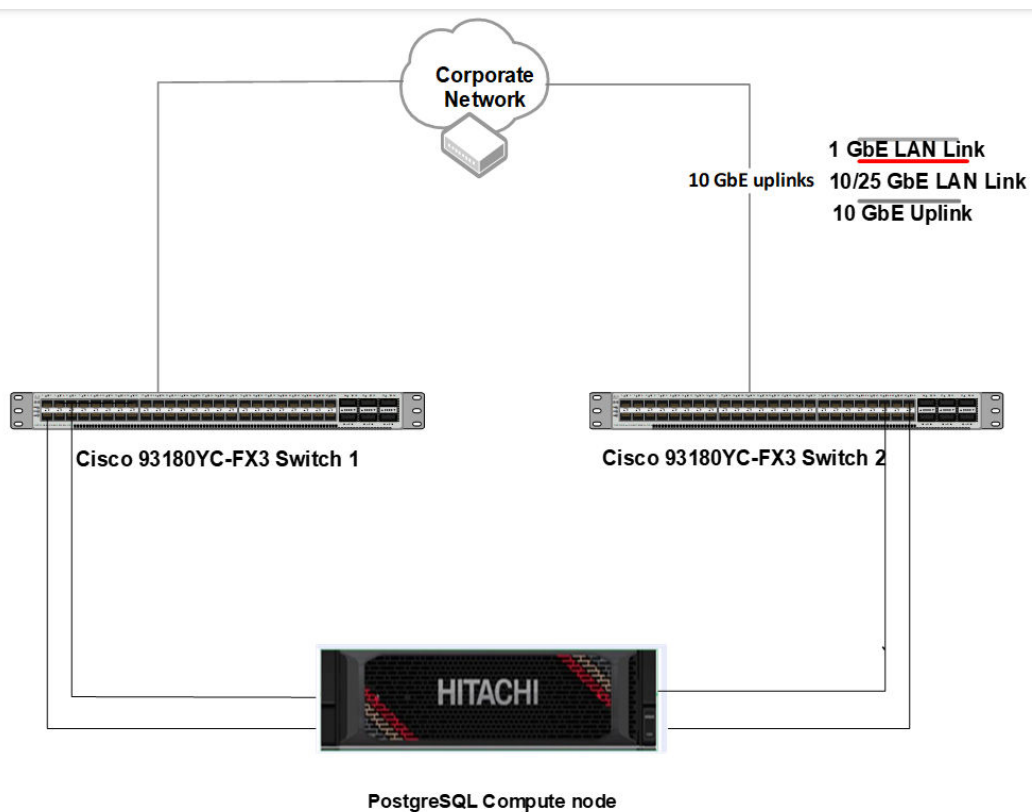
Name	IP Address	Notes
PostgreSQL Database Server	10.76.32.152	Database Compute node
VSP One Block 28	10.76.32.34	Storage System
Cisco Management Switch	192.168.242.11	Network Switch
Brocade G720 Switch-1	10.76.32.29	Fibre Channel Switch 1
Brocade G720 Switch-2	10.76.32.30	Fibre Channel Switch 2

Note that Hitachi Vantara recommends using pairs of 25 Gbps NICs for the public network to improve bandwidth. Observe these guidelines when configuring private and public networks in your environment:

- For each server in the configuration, use at least two identical, high-bandwidth, low-latency NICs for the interconnection.
- Use NIC bonding to provide failover and load balancing of interconnections within a server.
- Set all NICs to full duplex mode.
- Use at least two public NICs for client connections to the application and the database.

## Physical network configuration

The following figure shows the IP network switch connection.



The following table lists the HA820 G3 network configuration for this solution.

Server	NIC Ports	Subnet	NIC	IP Address	Network	Bandwidth (Gbps)	Cisco Nexus 93180YC-FX3 Switch	
							Switch Number	Port
HA820 G3 Server	NIC 1-Port 0	242	Ens2f0	192.168.242.204	Public	25	1	41
	NIC 1-Port 1					25	2	
	iLO-Dedicated NIC	32	-	10.76.32.xx	Management	1		



**Note:** When creating NIC bonding pairs, ports should be used on different cards to avoid single points of failure (SPoF).

The following table lists the network configuration for servers and VSP One Block 28.

Name	IP Address
PostgreSQL DB RedHat Linux host 1	192.168.242.xx
Management Server 1	192.168.242.xx
VSP One Block 28	192.168.242.xx
VSP One Block CTL1	192.168.242.xx
VSP One Block CTL2	192.168.242.xx

The following table lists the network configuration for switches used in this solution.

SwitchType	Model	Switch Name	IP Address for MGMT port
Cisco 10G/25GbE Network Switch	Cisco Nexus C93180YC-FX3	Cisco C93180YC-FX3-1	192.168.242.xx
Cisco 10G/25GbE Network Switch	Cisco Nexus C93180YC-FX3	Cisco C93180YC-FX3-2	192.168.242.xx
Brocade Fibre Channel SAN Switch	G720	SAN-switch1	192.168.242.xx

SwitchType	Model	Switch Name	IP Address for MGMT port
Brocade Fibre Channel SAN Switch	G720	SAN-Switch 2	192.168.242.xx

## Engineering validation

This section summarizes the key lab verification tests performed on Hitachi UCP Solution for Databases - PostgreSQL using Hitachi Virtual Storage Platform One Block 28 and Hitachi Advanced Server HA820 G3.

### Database configuration

The following table lists parameter details for a PostgreSQL database on a bare metal host.

PostgreSQL Database Parameter	Value
Compatible	15.7
PostgreSQL Database size	3 TB
Database storage type	SAN
Database fill factor	70%

### Test environment

The following table lists configuration details for VSP One Block 28 testing.

Item	Value
Operating System	RHEL 8.6 4.18.0-372.9.1.el8.x86_64
Workload Type	OLTP/OLAP
Database Size	3 TB
Number of vCPUs	192
Host Network	2 × 25 Gbps NIC Bonding

## Test methodology

The test results are demonstrated using HammerDB. It is a benchmarking tool, popular among industry standards. HammerDB is derived from TPC specifications. For measuring performance it executes two types of workloads, namely transactional and decision support. TPROC-C (OLTP) is used for executing transactional workload and TPROC-H (OLAP) is used for executing decision support workload test cases.

- Performance verification (quality assurance)
- Evaluation of different infrastructure products, technologies, and solutions (price/performance comparison)
- Performance optimization (improvement in efficiency)

This provides transparency and comparability in price versus performance considerations for a PostgreSQL infrastructure. HammerDB version 4.11 is used to validate this solution.

## Conclusion

We performed various database validation tests for this unified converged system consisting of a software and hardware stack configured on Hitachi Advanced Server HA820 G3 as a compute node and Hitachi Virtual Storage Platform One Block 28 storage. We found that all database functionalities for PostgreSQL database ran efficiently without any issues. We validated several database operations for various workloads using the HammerDB performance benchmarking tools and found that all types of workloads executed smoothly with optimal performance results.

This solution was also tested by running multiple workloads in parallel for a longer duration. We received good results for OLTP/OLAP transactions. Therefore, Hitachi Vantara hardware provides robust and high-performance infrastructure components, including storage systems, processors, and networking equipment, which ensures that database operates on a reliable and efficient foundation, delivering consistent performance even under heavy workloads.

Hitachi Virtual Storage Platform One Block 28 storage system and Hitachi Advanced Server HA820 G3 give power to compute nodes in non-virtualized environment to perform heavy transactions and workload in diverse conditions. At the same time, the server is stable in memory and CPU resources utilization.



**Note:** Contact the Hitachi sales and engineering team for more details about performance results and best configuration practices. [https://support.hitachivantara.com/en\\_us/contact-us.html](https://support.hitachivantara.com/en_us/contact-us.html).

## Product descriptions

These products are used in this reference architecture.

### Hitachi Virtual Storage Platform One Block



The Hitachi Virtual Storage Platform One Block series simplifies system setup and management through Hitachi Clear Sight and VSP One Block Administrator. Dynamic Drive Protection reduces RAID complexity, and always-on compression and deduplication enhance simplicity.

Dynamic Carbon Reduction optimizes energy usage by switching CPUs to ECO mode during low activity. Adaptive Data Reduction (ADR) is always on, enhancing efficiency and reducing the overall CO2 footprint.

Thin Image Advanced (TIA) integrates with major snapshot ecosystems, prioritizing security by defending against threats and ensuring data confidentiality. CyberArk Privileged Access Manager plugins enhance block storage system security by prioritizing data confidentiality, ensuring compliance, and actively defending against security threats.

Hitachi Virtual Storage Platform One Block includes the following 3 dedicated models:

- VSP One Block 24 – 256 GB Cache + SW Advanced Data Reduction (ADR) + 24 cores
- VSP One Block 26 – 768 GB Cache + 2x Compression Accelerator Module (CAM) + 24 cores
- VSP One Block 28 – 1 TB Cache + 4x CAM + 64 cores

All have the same capacity (72 NVMe flash drives, the appliance, and 2 × media trays) and they support Fibre Channel, iSCSI, and NVMe TCP connectivity. The new capabilities remove complexity like data reduction is always on, Dynamic Drive Protection removes complicated RAID setup and Dynamic Carbon Reduction delivers real world reduction in power consumption. In addition, the models are FIPS compliant.

In short, the Hitachi Virtual Storage Platform One Block series combines simplicity, sustainability, and robust security features to optimize system management, energy efficiency, and data protection.

## Hitachi Advanced Server HA820 G3

Hitachi Advanced Server HA820 is a high-performance two-socket rackmount server designed for optimal performance and power efficiency. This allows owners to upgrade computing performance without overextending power consumption and offers non-latency support to environments that require the maximum memory capacity. Hitachi Advanced Server HA820 G3 provides flexible I/O scalability for today's diverse data center application requirements.

Optimized for performance, high density, and power efficiency in a dual-processor server, HA820 G3 delivers a balance of compute and storage capacity. These rack mounted servers have the flexibility to power a wide range of solutions and applications.

The highly scalable memory supports up to 8 TB RAM using 32 slots of 2300 MHz DDR5 RDIMM. HA820 G3 is powered by the Intel Xeon Emerald Rapids scalable processor family for complex and demanding workloads. Flexible OCP and PCIe I/O expansion card options are available.

## PostgreSQL

PostgreSQL is the most demanding and powerful open-source RDBMS worldwide. It has a consistently proven track of innovation, robustness, and advanced features. With the evolution of time, they continue to add new features and functionalities in upcoming releases.

To help with database issues they have a wide community to support users. See <https://www.postgresql.org/community/> for details.

## HammerDB

HammerDB is the most popular, advanced, open-source, easy to use performance benchmarking tool available in the market. Source code for benchmarking workloads is based on the TPC specification. HammerDB can be used to run OLTP and OLAP workloads. In HammerDB, an OLAP workload is called TPROC-H which is derived from the TPC-H specification and an OLTP workload is called TPROC-C which is derived from the TPC-C specification.

TPROC-H Benchmarking is a decision support benchmark. It consists of a business-oriented standard set of 22 complex queries running against a database. The queries in this benchmark are designed in a way that it scans a large volume data with a high degree of complexity such as table joins, data sorting, and use of arithmetic functions. To increase the load on a system for capacity testing we run workloads using concurrent users or run workloads with additional parallel threads.

TPROC-C Benchmarking is an online transaction processing benchmark. It consists of queries or transactions running to the database. The transactions running in this workload have a combination of select, insert, update or delete records on tables. We can execute transactions from multiple users and the result should be minimal response time. We get two values— one is new orders per minute (NOPM) and the other is transactions per minute (TPM).

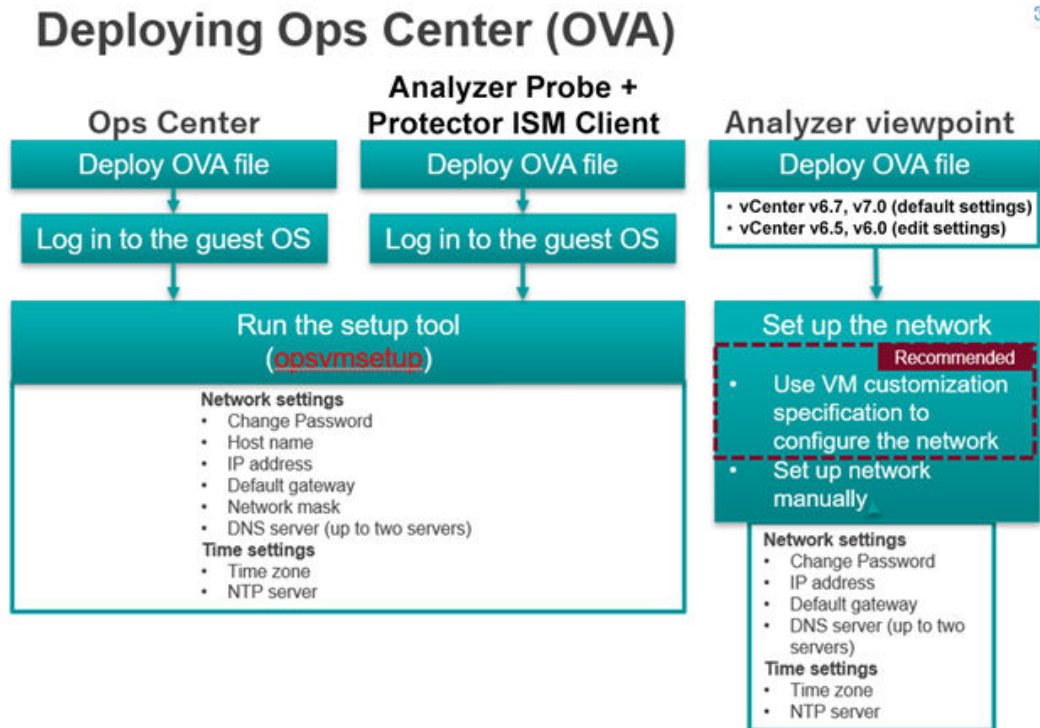
For more information, see <https://www.hammerdb.com/>.



**Note:** HammerDB Benchmark workloads are derived from the TPC specification, but results should not be compared with official TPC benchmark results.

## Hitachi Ops Center

Hitachi Ops Center is an integrated suite of applications that enable you to optimize your data center operations through integrated configuration, analytics, automation, and copy data management. These features allow you to administer, automate, optimize, and protect your Hitachi storage infrastructure.



See the *Ops Center Administrator Getting Started Guide* at <https://docs.hitachivantara.com/r/en-us/ops-center-administrator/11.0.x/mk-99adm000> for information about Hitachi Ops Center system requirements.

## Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

Brocade Fibre Channel switches deliver industry-leading performance with seventh generation 64Gb/sec Fibre Channel interfaces, simplifying scale-out network architectures. Get the high-performance, availability, ease of management, and support for the next generation of Hitachi Virtual Storage Platform storage systems on a solid storage network foundation that can grow as your need grows.

## Cisco Nexus switches

The Cisco Nexus switch product line provides a series of solutions that make it easier to connect and manage disparate data center resources with software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data, and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

**Hitachi Vantara**

Corporate Headquarters  
2535 Augustine Drive  
Santa Clara, CA 95054 USA



[HitachiVantara.com/contact](https://HitachiVantara.com/contact)