Brochure

Empower data-first modernization for mission-critical workloads

Hewlett Packard Enterprise

intel

Key workloads

- SAP HANA, Oracle, and Microsoft SQL Server databases
- Enterprise resource planning (ERP)
- Supply chain management
- Core banking
- Electronic healthcare records (EHR)
- Real-time data analytics
- Fraud management
- Electronic trading
- Cybersecurity
- Pre and post data processing for AI workloads
- Smart manufacturing
- UNIX® to Linux migrations

HPE Compute Scale-up Server 3200

Business and IT leaders are prioritizing data-first modernization to keep pace with digital transformation and changing business requirements. Data drives competitive advantage, and in a rapidly changing environment, businesses need to respond with speed to build new capabilities, and move faster than their competition. At the same time, when it comes to mission-critical workloads, enterprises need a solid, secure and reliable technology foundation.

At enterprises who are still running legacy and multi-generational IT, data is often spread across silos. Some key data may by trapped in legacy systems that are difficult to modernize. These legacy systems are also difficult to maintain, because a growing number of personnel with the skills to run them are reaching retirement. Many businesses are therefore highly cautious about making changes to mission-critical systems that work, and which already deliver the uptime levels they need. Often, they find the challenge overwhelming. However, by postponing modernization, businesses fall further behind their data-driven, cloud-enabled competitors.

Hewlett Packard Enterprise has developed a solution to empower data-first modernization for your mission-critical applications. The HPE Compute Scale-up Server 3200 is a powerful, single-system solution that uses a modular, scale-up architecture to reliably handle the workloads at the core of your business. These workloads span the full data continuum, including **systems of record** that process transactions, **systems of intelligence** that analyze data, and **systems of insight** that are fueled by that intelligence and augmented with AI to drive the business forward. The Intel® powered HPE Compute Scale-up Server 3200 offers a unique combination of scalability, reliability, and security to power critical databases and applications, including SAP HANA®, Oracle®, and Microsoft SQL Server, as well as Electronic Healthcare Records, Core Banking systems, Enterprise Resource Planning and more. It can comb through massive IoT and AI datasets at the edge or in the core to deliver real-time data analytics and tackle complex HPC problems holistically.

High velocity data transactions

Enterprise resource planning (ERP)
Core banking
Supply chain management
Electronic healthcare records (EHR)

Faster business insights

Pre and post data processing for AI
Smart manufacturing
Fraud management
Enterprise and edge AI

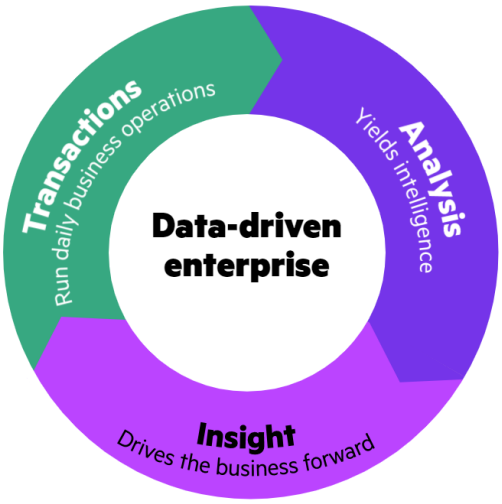

Real-time data analytics

Graph analytics
Risk analytics
Trading
Cybersecurity

Chart showing the data continuum of transactions, analytics and insights

HPE Compute Scale-up Server 3200 At-a-glance

- 4-16 processors in a single system
- Scale up in 4-socket increments
- Powered with 4th Gen Intel Xeon Scalable processors
- Modular scale-up 5U building block/chassis (4 sockets per chassis)
- 128GB - 32TB of shared DDR5 memory
- Choice of 6, 12, or 16 PCIe 5.0 slots per 4-socket chassis
- Extreme RAS and superior security
- As-a-service consumption with HPE GreenLake

Picture of the HPE Compute Scale-up Server 3200 4-socket chassis

Outpace evolving data demands with optimum flexibility

HPE Compute Scale-up Server 3200 utilizes the HPE Superdome Flex modular architecture. Customers can start with four 4th Gen Intel® Xeon® Scalable processors in a 5U chassis and seamlessly scale up to 16 sockets, in 4-socket increments. The server provides a single pool of shared memory and compute resources, even at maximum scale. This scale-up architecture reduces server sprawl and is easier to manage compared to scale-out deployments. By handling everything in one system, scale-up architectures also eliminate internode latency and the need to break up and reassemble large datasets. This can enable faster time to insight than is possible with scale-out deployments.

Offering up to 16 sockets and up to 960 processor cores, HPE Compute Scale-up Server 3200 offers plenty of headroom to grow. Customers have a cost-efficient entry point for mission-critical workloads at four sockets and can scale up to 16 at their own pace, meaning there is no need to overprovision.

Process and analyze ever-growing data at extreme speed

HPE Compute Scale-up Server 3200 provides the performance levels required to run the most demanding workloads. At its core are 4th generation Intel Xeon Scalable processors that support up to 60 cores per processor[1], more than twice as many as the prior generation. These processors feature a range of built-in accelerators designed to boost performance across the fastest-growing workloads, including AI and data analytics.

To support growth of in-memory databases and data-intensive workloads, HPE Compute Scale-up Server 3200 offers large, shared memory capacity ranging from 128 GB to 32 TB. It utilizes ultra-fast DDR5 DIMMs with speeds up to 4800 MT/s and 16 DIMM slots per processor, delivering twice the memory bandwidth[2] compared to prior generations.

The server's flexible and well-balanced I/O subsystem is based on PCIe 5.0, delivering four times more I/O bandwidth[3] than the prior model. Customers can choose from several I/O configurations to suit specific workload requirements:

- **"Max Drive" for database workloads:** 6 x I/O slots and 24 x EDSFF drives per 4-socket chassis.
- **"Max I/O" for I/O-intensive workloads:** 16 x I/O slots and up to 10 x2.5" drives per 4-socket chassis.
- **"Max GPU" for HPC use cases:** 12 x I/O slots and up to 10 x 2.5" drives per 4-socket chassis. These include 4 x full-height and double-width GPU slots and 4 x low profile x16 PCIe card slots.
- Customers can also select an expansion chassis with no I/O.

[1] 3rd Gen Intel Xeon Scalable processors deliver up to 28 cores per processor
[2] Source: HPE R&D, June 2023
[3] Compared to PCIe 3.0, utilized by HPE Superdome Flex 280

HPE Compute Scale-up Server 3200 is built on a technology foundation that delivers low latency and high bandwidth across every component – from processors, to memory, to I/O, to the fabric that connects each 4-socket chassis together. An innovative fabric link design and new external node controller (XNC) double the effective fabric bandwidth[4] compared to previous-generation servers and dramatically improve system serviceability.

HPE Compute Scale-up Server 3200 delivers system administration, control, and platform management using the scalable, modular HPE Superdome Flex family architecture. A Remote Management Processor (RMP) is embedded in each 4-socket chassis. One RMP is designated as the Rack Management Controller (RMC), which provides a single management interface regardless of whether the system has one, two, three or four 4-socket chassis. The management system embeds some of the key reliability capabilities of the server and improves resiliency with a redundant internal management LAN that protects against switch and cable failures. The RMC has a standards-based Redfish API for scripting and automation, as well as an HPE web GUI for common tasks such as system inventory, health, configuration, RMC security, and LAN settings. It also provides a powerful command line interface (CLI) for easy access to all RMC functions, providing potential scripting and power-user convenience. The Redfish API can be used in many ways, including directly in simple scripts to obtain inventory and monitoring information, as well as using tools and for compatibility. HPE OneView can also be used to manage multiple HPE systems concurrently.

Protect mission-critical workloads with proven reliability and trusted security

Because downtime events for mission critical environments can cost from thousands to hundreds of thousands of dollars per hour, choosing a reliable infrastructure foundation to run core workloads is essential. HPE Compute Scale-up Server 3200 is built on the proven reliability, availability, and serviceability (RAS) framework of the HPE Superdome Flex family, designed to enable the highest service levels for critical workloads with superior capabilities across every component of the IT stack. The server implements the full RAS functionality provided by the 4th Generation Intel® Xeon® Scalable processors to enhance memory reliability, reduce memory outage rates, and more.

Some of the server's unique RAS capabilities include:

A unique **firmware-first** approach to problem diagnosis. Firmware with detailed knowledge of HPE Compute Scale-up Server 3200 is first on the scene when problems occur, to quickly analyze and fix them. This approach results in appropriate actions being taken on the platform before any interruption can occur at the OS and application layers.

An **analysis engine** embedded in the management system offers outstanding predictive fault handling and initiates self-repair without operator assistance. It constantly monitors all system hardware, analyzes log and telemetry data, and determines corrective actions for high system uptime, often performing them without any user intervention. The analysis engine is designed to be combined with HPE Insight Remote Support or HPE OneView to connect back to HPE for the fastest-possible service response times. In the event of a fatal error, an error logging service initiates out-of-band error collection, creates logs, and alerts an offline system analyzer that provides post-failure analysis for complex error conditions.

Automatic de-configuration capabilities respond to errors by de-configuring the failing components, then resetting and rebooting the system. De-configuration scenarios are available at the core, I/O, DIMM, socket and fabric levels.

[4] Source: HPE R&D, May 2023

With the increasing pace of security breaches and attacks, businesses must ensure every component of their IT environment is resilient and protected against threats. This is especially significant when it comes to mission-critical workloads and data. The HPE Compute Scale-up Server 3200 is built on a security strategy designed to reduce threat exposure to vulnerabilities, including those found in common server firmware. Starting with an end-to-end secure supply chain foundation, the server features a comprehensive set of security capabilities including:

- Best-in-class **HPE Silicon Root of Trust**, which is implemented directly in specialized, HPE-controlled hardware. This detects potentially compromised firmware and prevents its execution.
- **Isolated management network:** Separate core and management LAN interface ports prevent the host operating system from accessing the management network directly; stopping it from getting to other management processors through the RMC.
- Hardware that prevents the host from updating system and management firmware directly, eliminating a major source of vulnerabilities.
- **Bundled and authenticated** firmware updates.
- **Automatic and secure recovery** of any compromised firmware.
- **Dual Trusted Platform Modules** in every chassis, one for the system and one for the management sub-system, soldered to each server in a tamper-resistant package with built-in firmware resiliency.

HPE Compute Scale-up Server 3200 also works with **HPE Serviceguard for Linux®**, a high availability and disaster recovery clustering solution that further strengthens business continuity by protecting applications from faults across physical and virtual environments over any distance.

Consume and finance on your terms

For customers who want the agility and economics of a hybrid cloud approach while staying in control of critical applications and data, HPE Compute Scale-up Server 3200 is available via the HPE GreenLake edge-to-cloud platform.

HPE GreenLake enables you to scale easily, adding capacity in minutes instead of months. You pay only for what you actually require, achieving faster time to value with solutions that are deployed quickly and can evolve ahead of your needs. Its flexible, pay-per-use* model offers simplicity, financial clarity, and proper control over compliance, as well as the performance and security of an on-premises deployment. With HPE GreenLake you benefit from metered usage that is billed monthly as a service, no up-front payment requirements, flexibility to add capacity in minutes from an on-site buffer, and enterprise-quality support.

In summary

Stay ahead of the competition by powering your mission-critical workloads with the reliable, scalable and secure HPE Compute Scale-up Server 3200 powered by Intel. Get real-time insights no matter how much data you have or how fast it grows, harness the compute power needed for your most demanding workloads, and achieve the highest service levels with proven reliability and trusted security. Deploying your solution with HPE GreenLake also provides the flexibility and economics of the cloud, while maintaining security and control of your critical data. For more information, contact your Hewlett Packard Enterprise representative or visit our website hpe.com/superdome.

* May be subject to minimums or reserve capacity may apply

Chat now (sales)

Hewlett Packard Enterprise

a50008800enw