

Overview

HPE Scalable Compute Software

HPE Scalable Compute Software is a software solution that aggregates the memory, CPUs, and I/O of multiple physical servers to create a single software-defined server. At the core of the HPE Scalable Compute Software software is a proprietary bare-metal hypervisor, called the HyperKernel, that leverages virtualization extensions on modern processors to virtualize the physical resources of a server. HPE Scalable Compute Software runs a copy of this hypervisor on each physical machine (worker node) that is aggregated into a single software-defined server.

HPE Scalable Compute Software's patented technology communicates over a high-speed interconnect network using standard Ethernet between all the worker nodes running the HyperKernel and presents an aggregated, unified view all the physical resources from every server to the Operating System and application stack. From the perspective of the Operating System, it is running on a single large server that sees the memory, CPU, and I/O capabilities as the sum of the individual server resources composing the software-defined server.

Physical servers may be added or removed as the workload demands, increasing or decreasing the overall available resources of the software-defined server and thus ensuring that the software-defined server is always appropriately sized. HPE Scalable Compute Software eliminates the need for costly and complex sizing, migration, and procurement exercises. Furthermore, HPE Scalable Compute Software avoids the long-term commitments on expensive hardware, giving users the flexibility to resize server capacity at any time based on current workload demand.

While the software-defined server is in operation, HPE Scalable Compute Software's technology incorporates real-time machine learning to continually update CPU and memory locations between physical servers to produce optimal performance regardless of shifts in the application workload. Continually adapting based on workload changes, the software predicts optimal placement of memory and CPU thread locations and triggers migrations whenever necessary. HPE Scalable Compute Software can migrate memory pages to the CPU thread that will need the memory, or migrate a CPU thread to the block of memory pages it will need to provide the best possible performance.

Standard Features

HPE Scalable Compute Software

HPE Scalable Compute Software provides two licensable features types to address the needs of modern, data-centric organizations:

- **Scalability**

For organizations with workloads that require the ability to quickly scale their systems to meet growing workloads, HPE Scalable Compute Software allows customers to rapidly add or remove servers to adjust the size of their infrastructure to meet the current and future needs of their applications.

- **Availability**

Used by organizations who wish to avoid downtime, HPE Scalable Compute Software allows administrators to swap out a running server with another server without impacting the operating system or application. Unlike other technologies which create an active/passive relationship with the incoming server and then track and transmit to that server before taking a short outage to cutover to the new system, HPE Scalable Compute Software creates an active/active relationship with the incoming server and seamlessly migrates the workload to the new system without interruption.

Using this functionality, Administrators can manually swap out a server at any time to service a system with proactive (e.g., firmware) or reactive (e.g., failed power supply) issues. Additionally, the system can be configured to monitor its underlying hardware for anomalies and automatically initiate a swap without any manual interaction to heal itself before server issues can impact application availability.



Service and Support

HPE Pointnext - Service and Support

Get the most from your HPE Products. Get the expertise you need at every step of your IT journey with **HPE Pointnext Services**. We help you lower your risks and overall costs using automation and methodologies that have been tested and refined by HPE experts through thousands of deployments globally. HPE GreenLake Cloud Services, **Advisory and Professional Services**, focus on your business outcomes and goals, partnering with you to design your transformation and build a roadmap tuned to your unique challenges to help deploy AI. Our **Professional** and **Operational Services** can be leveraged to speed up time-to-production, boost performance and accelerate your business. HPE Pointnext specializes in flawless and on-time implementation, on-budget execution, and creative configurations that get the most out of software and hardware alike.

HPE Technology Services Support Credits

Offer flexible services and technical skills to meet your changing IT demands. With a menu of services that are tailored to suit your needs, you get additional resources and specialist skills to help you maintain the peak performance of your IT. Offered as annual credits, you can plan your budgets while proactively responding to your dynamic business.

HPE Support Center

The HPE Support Center is a personalized online support portal with access to information, tools and experts to support HPE business products. Submit support cases online, chat with Hewlett Packard Enterprise experts, access support resources or collaborate with peers.

Learn more at <http://www.hpe.com/support/hpesc>



Configuration Information

HPE Scalable Compute Software

Hewlett Packard Enterprise is making the following HPE Scalable Compute Software SKUs available:

Scalability Licenses

HPE Scalable Compute Software Scalability 2 Processors 1-year E-LTU	S1U05AAE
HPE Scalable Compute Software Scalability 2 Processors 3-year E-LTU	S1U06AAE
HPE Scalable Compute Software Scalability 2 Processors 5-year E-LTU	S1U07AAE

Availability Licenses

HPE Scalable Compute Software Availability 2 Processors 1-year E-LTU	S1U02AAE
HPE Scalable Compute Software Availability 2 Processors 3-year E-LTU	S1U03AAE
HPE Scalable Compute Software Availability 2 Processors 5-year E-LTU	S1U04AAE



Technical Specifications

Worker Nodes are the physical servers used by HPE Scalable Compute Software to construct software-defined servers. All Worker Nodes within a Pool – a collection of servers grouped for the purpose of creating an SDServer – must have the same technical specifications in terms of motherboard, CPUs, memory (DIMM size, speed, quantity), and network interfaces. If multiple server types are deployed, at least one Pool must be created for each hardware configuration. Identical server types may be split into separate Pools (e.g., separating development, testing, and production environments).

Each Worker Node must meet these minimum technical specifications:

Motherboard	1-socket, 2-socket or 4-socket system with: <ul style="list-style-type: none"> • LGA 1150 or C603 chipset • An integrated Baseboard Management Controller (BMC) with IPMI 2.0 support Registered ECC DDR3 (or newer) DIMMS
CPU	Intel Xeon E5v3 Processor (or newer) with VT-x and VT-d support
RAM	32GB+ Registered ECC DDR3+ memory
Storage	As required for your OS & Application needs (may be local storage or SAN)
NICs	<ul style="list-style-type: none"> • 2 - 1GbE (or higher) NICs • 1 - 10GbE (or higher) NIC





Summary of Changes

Date	Version History	Action	Description of Change
05-Jun-2023	Version 1	New	New QuickSpecs



Copyright

**Make the right purchase decision.
Contact our presales specialists.**



© Copyright 2023 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a50006998enw - 17108 - Worldwide - V1 - 05-June-2023